Graduate Texts in Mathematics

Henri Cohen

Advanced Topics in Computational Number Theory

Springer

Graduate Texts in Mathematics 193

Editorial Board
S. Axler F.W. Gehring K.A. Ribet

Springer
New York
Berlin
Heidelberg
Barcelona
Hong Kong
London
Milan
Paris
Singapore
Tokyo

Graduate Texts in Mathematics

1 TAKEUTI/ZARING. Introduction to Axiomatic Set Theory. 2nd ed.
2 OXTOBY. Measure and Category. 2nd ed.
3 SCHAEFER. Topological Vector Spaces. 2nd ed.
4 HILTON/STAMMBACH. A Course in Homological Algebra. 2nd ed.
5 MAC LANE. Categories for the Working Mathematician. 2nd ed.
6 HUGHES/PIPER. Projective Planes.
7 SERRE. A Course in Arithmetic.
8 TAKEUTI/ZARING. Axiomatic Set Theory.
9 HUMPHREYS. Introduction to Lie Algebras and Representation Theory.
10 COHEN. A Course in Simple Homotopy Theory.
11 CONWAY. Functions of One Complex Variable I. 2nd ed.
12 BEALS. Advanced Mathematical Analysis.
13 ANDERSON/FULLER. Rings and Categories of Modules. 2nd ed.
14 GOLUBITSKY/GUILLEMIN. Stable Mappings and Their Singularities.
15 BERBERIAN. Lectures in Functional Analysis and Operator Theory.
16 WINTER. The Structure of Fields.
17 ROSENBLATT. Random Processes. 2nd ed.
18 HALMOS. Measure Theory.
19 HALMOS. A Hilbert Space Problem Book. 2nd ed.
20 HUSEMOLLER. Fibre Bundles. 3rd ed.
21 HUMPHREYS. Linear Algebraic Groups.
22 BARNES/MACK. An Algebraic Introduction to Mathematical Logic.
23 GREUB. Linear Algebra. 4th ed.
24 HOLMES. Geometric Functional Analysis and Its Applications.
25 HEWITT/STROMBERG. Real and Abstract Analysis.
26 MANES. Algebraic Theories.
27 KELLEY. General Topology.
28 ZARISKI/SAMUEL. Commutative Algebra. Vol.I.
29 ZARISKI/SAMUEL. Commutative Algebra. Vol.II.
30 JACOBSON. Lectures in Abstract Algebra I. Basic Concepts.
31 JACOBSON. Lectures in Abstract Algebra II. Linear Algebra.
32 JACOBSON. Lectures in Abstract Algebra III. Theory of Fields and Galois Theory.
33 HIRSCH. Differential Topology.
34 SPITZER. Principles of Random Walk. 2nd ed.
35 ALEXANDER/WERMER. Several Complex Variables and Banach Algebras. 3rd ed.
36 KELLEY/NAMIOKA et al. Linear Topological Spaces.
37 MONK. Mathematical Logic.
38 GRAUERT/FRITZSCHE. Several Complex Variables.
39 ARVESON. An Invitation to C^*-Algebras.
40 KEMENY/SNELL/KNAPP. Denumerable Markov Chains. 2nd ed.
41 APOSTOL. Modular Functions and Dirichlet Series in Number Theory. 2nd ed.
42 SERRE. Linear Representations of Finite Groups.
43 GILLMAN/JERISON. Rings of Continuous Functions.
44 KENDIG. Elementary Algebraic Geometry.
45 LOÈVE. Probability Theory I. 4th ed.
46 LOÈVE. Probability Theory II. 4th ed.
47 MOISE. Geometric Topology in Dimensions 2 and 3.
48 SACHS/WU. General Relativity for Mathematicians.
49 GRUENBERG/WEIR. Linear Geometry. 2nd ed.
50 EDWARDS. Fermat's Last Theorem.
51 KLINGENBERG. A Course in Differential Geometry.
52 HARTSHORNE. Algebraic Geometry.
53 MANIN. A Course in Mathematical Logic.
54 GRAVER/WATKINS. Combinatorics with Emphasis on the Theory of Graphs.
55 BROWN/PEARCY. Introduction to Operator Theory I: Elements of Functional Analysis.
56 MASSEY. Algebraic Topology: An Introduction.
57 CROWELL/FOX. Introduction to Knot Theory.
58 KOBLITZ. p-adic Numbers, p-adic Analysis, and Zeta-Functions. 2nd ed.
59 LANG. Cyclotomic Fields.
60 ARNOLD. Mathematical Methods in Classical Mechanics. 2nd ed.
61 WHITEHEAD. Elements of Homotopy Theory.

(continued after index)

Henri Cohen

Advanced Topics in Computational Number Theory

Springer

Henri Cohen
Université de Bordeaux 1
Lab. Algorithmique Arithmétique Expérimentale
351, cours de la Libération
33405 Talence
France

Editorial Board

S. Axler
Mathematics Department
San Francisco State
University
San Francisco, CA 94132
USA

F.W. Gehring
Mathematics Department
East Hall
University of Michigan
Ann Arbor, MI 48109
USA

K.A. Ribet
Mathematics Department
University of California
at Berkeley
Berkeley, CA 94720-3840
USA

Mathematics Subject Classification (1991): 11-01, 11Yxx, 11Y16

Library of Congress Cataloging-in-Publication Data
Cohen, Henri
Advanced topics in computational number theory / Henri Cohen.
p. cm. — (Graduate texts in mathematics ; 193)
Includes bibliographical references and index.
ISBN 0-387-98727-4 (hardcover : alk. paper)
1. Number theory—data processing. I. Title. II. Series.
QA241.C667 1999
512'.7'0285—dc21 99-20756

Printed on acid-free paper.

© 2000 Springer-Verlag New York, Inc.
All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer-Verlag New York, Inc., 175 Fifth Avenue, New York, NY 10010, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.
The use of general descriptive names, trade names, trademarks, etc., in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

Production managed by Terry Kornak; manufacturing supervised by Jerome Basma.
Photocomposed copy prepared by the author using AMS-LaTeX and Springer's `clmono01` macros.
Printed and bound by R.R. Donnelley and Sons, Inc., Harrisonburg, VA.
Printed in the United States of America.

9 8 7 6 5 4 3 2 1

ISBN 0-387-98727-4 Springer-Verlag New York Berlin Heidelberg SPIN 10708040

Preface

The computation of invariants of algebraic number fields such as integral bases, discriminants, prime decompositions, ideal class groups, and unit groups is important both for its own sake and for its numerous applications, for example, to the solution of Diophantine equations. The practical completion of this task (sometimes known as the Dedekind program) has been one of the major achievements of computational number theory in the past ten years, thanks to the efforts of many people. Even though some practical problems still exist, one can consider the subject as solved in a satisfactory manner, and it is now routine to ask a specialized Computer Algebra System such as Kant/Kash, LiDIA, Magma, or Pari/GP, to perform number field computations that would have been unfeasible only ten years ago.The (very numerous) algorithms used are essentially all described in *A Course in Computational Algebraic Number Theory*, GTM 138, first published in 1993 (third corrected printing 1996), which is referred to here as [Coh0]. That text also treats other subjects such as elliptic curves, factoring, and primality testing.

It is important and natural to generalize these algorithms. Several generalizations can be considered, but the most important are certainly the generalizations to global function fields (finite extensions of the field of rational functions in one variable over a finite field) and to relative extensions of number fields. As in [Coh0], in the present book we will consider number fields only and not deal at all with function fields.

We will thus address some specific topics related to number fields; contrary to [Coh0], there is no attempt to be exhaustive in the choice of subjects. The topics have been chosen primarily because of my personal tastes, and of course because of their importance. Almost all of the subjects discussed in this book are quite new from the algorithmic aspect (usually post-1990), and nearly all of the algorithms have been implemented and tested in the number theory package Pari/GP (see [Coh0] and [BBBCO]). The fact that the subjects are new does not mean that they are more difficult. In fact, as the reader will see when reading this book in depth, the algorithmic treatment of certain parts of number theory which have the reputation of being "difficult" is in fact much *easier* than the theoretical treatment. A case in point is computational class field theory (see Chapters 4 to 6). I do not mean that the proofs become any simpler, but only that one gets a much better grasp on the subject by studying its algorithmic aspects.

As already mentioned, a common point to most of the subjects discussed in this book is that we deal with *relative* extensions, but we also study other subjects. We will see that most of the algorithms given in [Coh0] for the absolute case can be generalized to the relative case.

The book is organized as follows. Chapters 1 and 2 contain the theory and algorithms concerning Dedekind domains and relative extensions of number

fields, and in particular the generalization to the relative case of the round 2 and related algorithms.

Chapters 3, 4, 5, and 6 contain the theory and complete algorithms concerning class field theory over number fields. The highlights are the algorithms for computing the structure of $(\mathbb{Z}_K/\mathfrak{m})^*$, of ray class groups, and relative equations for Abelian extensions of number fields using Kummer theory, Stark's conjectures, and complex multiplication. The reader is warned that Chapter 5 is rather technical but contains a wealth of information useful both for further research and for any serious implementation. The analytic techniques using Stark's conjecture or complex multiplication described in Chapter 6 are fascinating since they construct purely algebraic objects using analytic means.

Chapters 1 through 6 together with Chapter 10 form a homogeneous subject matter that can be used for a one-semester or full-year advanced graduate course in computational number theory, omitting the most technical parts of Chapter 5.

The subsequent chapters deal with more miscellaneous subjects. In Chapter 7, we consider other variants of the notions of class and unit groups, such as relative class and unit groups or S-class and unit groups. We sketch an algorithm that allows the direct computation of relative class and unit groups and give applications of S-class and unit groups to the algorithmic solution of norm equations, due to D. Simon.

In Chapter 8, we explain in detail the correspondence between cubic fields and binary cubic forms, discovered by H. Davenport and H. Heilbronn, and examine the important algorithmic consequences discovered by K. Belabas.

In Chapter 9, we give a detailed description of known methods for constructing tables of number fields or number fields of small discriminant, either by using absolute techniques based on the geometry of numbers or by using relative techniques based either on the geometry of numbers or on class field theory.

In Appendix A, we give and prove a number of important miscellaneous results that can be found scattered in the literature but are used in the rest of the book.

In Appendix B, we give an updated but much shortened version of [Coh0, Appendix A] concerning packages for number theory and other useful electronic information.

In Appendix C, we give a number of useful tables that can be produced using the results of this book.

The book ends with an index of notation, an index of algorithms, and a general index.

The prerequisites for reading this book are essentially the basic definitions and results of algebraic number theory, as can be found in many textbooks, including [Coh0]. Apart from that, this book is almost entirely self-contained. Although numerous references are made to the algorithms con-

tained in [Coh0], these should be considered as "black boxes" and used as such. It would, however, be preferable at some point for the reader to study some of the algorithms of [Coh0]; in particular, those generalized here.

WARNINGS

(1) As usual, neither the author nor Springer-Verlag can assume any responsibility for consequences arising from the use of the algorithms given in this book.

(2) The author would like to hear about errors, typographical or otherwise. Please send e-mail to

`cohen@math.u-bordeaux.fr`

Lists of known errors, both for [Coh0] and for the present book, can be obtained by anonymous ftp from the URL

`ftp://megrez.math.u-bordeaux.fr/pub/cohenbook`

or obtained through the author's home page on the Web at the URL

`http://www.math.u-bordeaux.fr/~cohen`

(3) There is, however, another important warning that is almost irrelevant in [Coh0]. Almost all of the algorithms or the algorithmic aspects presented in this book are new, and most have never been published before or are being published while this book is going to press. Therefore, it is quite possible that major mistakes are present, although this possibility is largely diminished by the fact that almost all of the algorithms have been tested, although not always thoroughly. More likely it is possible that some algorithms can be radically improved. The contents of this book only reflect the knowledge of the author at the time of writing.

Acknowledgments

First of all, I would like to thank my colleagues Francisco Diaz y Diaz and Michel Olivier, with whom I have the pleasure of working every day and who collaborated with me on the discovery and implementation of many of the algorithms described in this book. Second, I would like to thank Jacques Martinet, head of our Laboratoire, who has enormously helped by giving me an ideal working environment and who also has tirelessly answered my numerous questions about most of the subject matter of this book. Third, I thank my former students Karim Belabas, Jean-Marc Couveignes, Denis Simon, and Emmanuel Tollis, who also contributed to part of the algorithms described here, and Xavier Roblot for everything related to Stark's conjectures.

In particular, Karim Belabas is to be thanked for the contents of Chapter 8, which are mainly due to him, for having carefully read the manuscript of this book, and not least for having taken the ungrateful job of managing the Pari software, after making thorough modifications leading to version 2.

I would like to thank several additional people who helped me in the preparation of this book. In alphabetical order, they are Claus Fieker (for Chapter 5), David Ford (for Chapter 2), Eduardo Friedman (for Chapter 7 and Appendix A), Thomas Papanikolaou (for Appendix B and for a lot of TEXnical help), and Michael Pohst (for Chapter 5).

I would also like to thank Mehpare Bilhan and the Middle East Technical University (METU) in Ankara, Turkey, for having given me an opportunity to write a first version of part of the subjects treated in this book, which appeared as an internal report of METU in 1997.

Last but not least, I thank all of our funding agencies, in particular, the C.N.R.S., the Ministry of Education and Research, the Ministry of Defense, the University of Bordeaux I, and the Région Aquitaine.

Contents

Preface v

1. Fundamental Results and Algorithms in Dedekind Domains 1
1.1 Introduction 1
1.2 Finitely Generated Modules Over Dedekind Domains 2
1.2.1 Finitely Generated Torsion-Free and Projective Modules 6
1.2.2 Torsion Modules 13
1.3 Basic Algorithms in Dedekind Domains 17
1.3.1 Extended Euclidean Algorithms in Dedekind Domains 17
1.3.2 Deterministic Algorithms for the Approximation Theorem 20
1.3.3 Probabilistic Algorithms 23
1.4 The Hermite Normal Form Algorithm in Dedekind Domains . 25
1.4.1 Pseudo-Objects 26
1.4.2 The Hermite Normal Form in Dedekind Domains 28
1.4.3 Reduction Modulo an Ideal 32
1.5 Applications of the HNF Algorithm 34
1.5.1 Modifications to the HNF Pseudo-Basis 34
1.5.2 Operations on Modules and Maps 35
1.5.3 Reduction Modulo $\mathfrak{p}$ of a Pseudo-Basis 37
1.6 The Modular HNF Algorithm in Dedekind Domains 38
1.6.1 Introduction 38
1.6.2 The Modular HNF Algorithm 38
1.6.3 Computing the Transformation Matrix 41
1.7 The Smith Normal Form Algorithm in Dedekind Domains ... 42
1.8 Exercises for Chapter 1 46

2. Basic Relative Number Field Algorithms 49
2.1 Compositum of Number Fields and Relative and Absolute Equations 49
2.1.1 Introduction 49
2.1.2 Étale Algebras 50
2.1.3 Compositum of Two Number Fields 56
2.1.4 Computing θ_1 and θ_2 59

2.1.5 Relative and Absolute Defining Polynomials 62
2.1.6 Compositum with Normal Extensions 66
2.2 Arithmetic of Relative Extensions 72
2.2.1 Relative Signatures 72
2.2.2 Relative Norm, Trace, and Characteristic Polynomial . 76
2.2.3 Integral Pseudo-Bases 76
2.2.4 Discriminants 78
2.2.5 Norms of Ideals in Relative Extensions 80
2.3 Representation and Operations on Ideals 83
2.3.1 Representation of Ideals 83
2.3.2 Representation of Prime Ideals 89
2.3.3 Computing Valuations 92
2.3.4 Operations on Ideals 94
2.3.5 Ideal Factorization and Ideal Lists 99
2.4 The Relative Round 2 Algorithm and Related Algorithms ... 102
2.4.1 The Relative Round 2 Algorithm 102
2.4.2 Relative Polynomial Reduction 110
2.4.3 Prime Ideal Decomposition 111
2.5 Relative and Absolute Representations 114
2.5.1 Relative and Absolute Discriminants 114
2.5.2 Relative and Absolute Bases 115
2.5.3 Ups and Downs for Ideals 116
2.6 Relative Quadratic Extensions and Quadratic Forms 118
2.6.1 Integral Pseudo-Basis, Discriminant 118
2.6.2 Representation of Ideals 121
2.6.3 Representation of Prime Ideals 123
2.6.4 Composition of Pseudo-Quadratic Forms 125
2.6.5 Reduction of Pseudo-Quadratic Forms 127
2.7 Exercises for Chapter 2 129

3. The Fundamental Theorems of Global Class Field Theory 133
3.1 Prologue: Hilbert Class Fields 133
3.2 Ray Class Groups 135
3.2.1 Basic Definitions and Notation 135
3.3 Congruence Subgroups: One Side of Class Field Theory 138
3.3.1 Motivation for the Equivalence Relation 138
3.3.2 Study of the Equivalence Relation 139
3.3.3 Characters of Congruence Subgroups 145
3.3.4 Conditions on the Conductor and Examples 147
3.4 Abelian Extensions: The Other Side of Class Field Theory ... 150
3.4.1 The Conductor of an Abelian Extension 150
3.4.2 The Frobenius Homomorphism 151
3.4.3 The Artin Map and the Artin Group $A_{\mathfrak{m}}(L/K)$ 152
3.4.4 The Norm Group (or Takagi Group) $T_{\mathfrak{m}}(L/K)$ 153
3.5 Putting Both Sides Together: The Takagi Existence Theorem 154

3.5.1 The Takagi Existence Theorem 154
3.5.2 Signatures, Characters, and Discriminants 156
3.6 Exercises for Chapter 3 160

4. Computational Class Field Theory 163
4.1 Algorithms on Finite Abelian groups 164
4.1.1 Algorithmic Representation of Groups 164
4.1.2 Algorithmic Representation of Subgroups 166
4.1.3 Computing Quotients 168
4.1.4 Computing Group Extensions 169
4.1.5 Right Four-Term Exact Sequences 170
4.1.6 Computing Images, Inverse Images, and Kernels 172
4.1.7 Left Four-Term Exact Sequences 174
4.1.8 Operations on Subgroups 176
4.1.9 p-Sylow Subgroups of Finite Abelian Groups 177
4.1.10 Enumeration of Subgroups 179
4.1.11 Application to the Solution of Linear Equations and Congruences 182
4.2 Computing the Structure of $(\mathbb{Z}_K/\mathfrak{m})^*$ 185
4.2.1 Standard Reductions of the Problem 186
4.2.2 The Use of $\mathfrak{p}$-adic Logarithms 190
4.2.3 Computing $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ by Induction 198
4.2.4 Representation of Elements of $(\mathbb{Z}_K/\mathfrak{m})^*$ 204
4.2.5 Computing $(\mathbb{Z}_K/\mathfrak{m})^*$ 206
4.3 Computing Ray Class Groups 209
4.3.1 The Basic Ray Class Group Algorithm 209
4.3.2 Size Reduction of Elements and Ideals 211
4.4 Computations in Class Field Theory 213
4.4.1 Computations on Congruence Subgroups 213
4.4.2 Computations on Abelian Extensions 214
4.4.3 Conductors of Characters 218
4.5 Exercises for Chapter 4 219

5. Computing Defining Polynomials Using Kummer Theory . 223
5.1 General Strategy for Using Kummer Theory 223
5.1.1 Reduction to Cyclic Extensions of Prime Power Degree 223
5.1.2 The Four Methods 226
5.2 Kummer Theory Using Hecke's Theorem When $\zeta_\ell \in K$ 227
5.2.1 Characterization of Cyclic Extensions of Conductor $\mathfrak{m}$ and Degree ℓ 227
5.2.2 Virtual Units and the ℓ-Selmer Group 229
5.2.3 Construction of Cyclic Extensions of Prime Degree and Conductor $\mathfrak{m}$ 233
5.2.4 Algorithmic Kummer Theory When $\zeta_\ell \in K$ Using Hecke236
5.3 Kummer Theory Using Hecke When $\zeta_\ell \notin K$ 242

5.3.1 Eigenspace Decomposition for the Action of τ 242
5.3.2 Lift in Characteristic 0 248
5.3.3 Action of τ on Units 254
5.3.4 Action of τ on Virtual Units 255
5.3.5 Action of τ on the Class Group 256
5.3.6 Algorithmic Kummer Theory When $\zeta_\ell \notin K$ Using Hecke 260
5.4 Explicit Use of the Artin Map in Kummer Theory When $\zeta_n \in K$ 270
5.4.1 Action of the Artin Map on Kummer Extensions 270
5.4.2 Reduction to $\alpha \in U_S(K)/U_S(K)^n$ for a Suitable S 272
5.4.3 Construction of the Extension L/K by Kummer Theory 274
5.4.4 Picking the Correct α 277
5.4.5 Algorithmic Kummer Theory When $\zeta_n \in K$ Using Artin 278
5.5 Explicit Use of the Artin Map When $\zeta_n \notin K$ 280
5.5.1 The Extension K_z/K 280
5.5.2 The Extensions L_z/K_z and L_z/K 281
5.5.3 Going Down to the Extension L/K 283
5.5.4 Algorithmic Kummer Theory When $\zeta_n \notin K$ Using Artin 284
5.5.5 Comparison of the Methods 287
5.6 Two Detailed Examples 288
5.6.1 Example 1 289
5.6.2 Example 2 290
5.7 Exercises for Chapter 5 293

6. Computing Defining Polynomials Using Analytic Methods 297
6.1 The Use of Stark Units and Stark's Conjecture 297
6.1.1 Stark's Conjecture 298
6.1.2 Computation of $\zeta'_{K,S}(0,\sigma)$ 299
6.1.3 Real Class Fields of Real Quadratic Fields 301
6.2 Algorithms for Real Class Fields of Real Quadratic Fields ... 303
6.2.1 Finding a Suitable Extension N/K 303
6.2.2 Computing the Character Values 306
6.2.3 Computation of $W(\chi)$ 307
6.2.4 Recognizing an Element of $\mathbb{Z}_K$ 309
6.2.5 Sketch of the Complete Algorithm 310
6.2.6 The Special Case of Hilbert Class Fields 311
6.3 The Use of Complex Multiplication 313
6.3.1 Introduction 314
6.3.2 Construction of Unramified Abelian Extensions 315
6.3.3 Quasi-Elliptic Functions 325
6.3.4 Construction of Ramified Abelian Extensions Using Complex Multiplication 333
6.4 Exercises for Chapter 6 344

7. Variations on Class and Unit Groups 347
7.1 Relative Class Groups 347
7.1.1 Relative Class Group for $i_{L/K}$ 348
7.1.2 Relative Class Group for $\mathcal{N}_{L/K}$ 349
7.2 Relative Units and Regulators 352
7.2.1 Relative Units and Regulators for $i_{L/K}$ 352
7.2.2 Relative Units and Regulators for $\mathcal{N}_{L/K}$ 358
7.3 Algorithms for Computing Relative Class and Unit Groups .. 360
7.3.1 Using Absolute Algorithms 360
7.3.2 Relative Ideal Reduction 365
7.3.3 Using Relative Algorithms 367
7.3.4 An Example 369
7.4 Inverting Prime Ideals 371
7.4.1 Definitions and Results 371
7.4.2 Algorithms for the S-Class Group and S-Unit Group . 373
7.5 Solving Norm Equations 377
7.5.1 Introduction 377
7.5.2 The Galois Case 378
7.5.3 The Non-Galois Case 380
7.5.4 Algorithmic Solution of Relative Norm Equations 382
7.6 Exercises for Chapter 7 386

8. Cubic Number Fields 389
8.1 General Binary Forms 389
8.2 Binary Cubic Forms and Cubic Number Fields 395
8.3 Algorithmic Characterization of the Set U 400
8.4 The Davenport–Heilbronn Theorem 404
8.5 Real Cubic Fields 409
8.6 Complex Cubic Fields 418
8.7 Implementation and Results 422
8.7.1 The Algorithms 422
8.7.2 Results 425
8.8 Exercises for Chapter 8 426

9. Number Field Table Constructions 429
9.1 Introduction 429
9.2 Using Class Field Theory 430
9.2.1 Finding Small Discriminants 430
9.2.2 Relative Quadratic Extensions 433
9.2.3 Relative Cubic Extensions 437
9.2.4 Finding the Smallest Discriminants Using Class Field Theory 444
9.3 Using the Geometry of Numbers 445
9.3.1 The General Procedure 445
9.3.2 General Inequalities 451

9.3.3 The Totally Real Case 453
9.3.4 The Use of Lagrange Multipliers 455
9.4 Construction of Tables of Quartic Fields 460
9.4.1 Easy Inequalities for All Signatures 460
9.4.2 Signature $(0,2)$: The Totally Complex Case 461
9.4.3 Signature $(2,1)$: The Mixed Case 463
9.4.4 Signature $(4,0)$: The Totally Real Case 464
9.4.5 Imprimitive Degree 4 Fields 465
9.5 Miscellaneous Methods (in Brief) 466
9.5.1 Euclidean Number Fields 467
9.5.2 Small Polynomial Discriminants 467
9.6 Exercises for Chapter 9 468

10. Appendix A: Theoretical Results 475
10.1 Ramification Groups and Applications 475
10.1.1 A Variant of Nakayama's Lemma 475
10.1.2 The Decomposition and Inertia Groups 477
10.1.3 Higher Ramification Groups 480
10.1.4 Application to Different and Conductor Computations 484
10.1.5 Application to Dihedral Extensions of Prime Degree .. 487
10.2 Kummer Theory 492
10.2.1 Basic Lemmas 492
10.2.2 The Basic Theorem of Kummer Theory 494
10.2.3 Hecke's Theorem 498
10.2.4 Algorithms for ℓth Powers 504
10.3 Dirichlet Series with Functional Equation 508
10.3.1 Computing L-Functions Using Rapidly Convergent Series 508
10.3.2 Computation of $F_i(s,x)$ 516
10.4 Exercises for Chapter 10 518

11. Appendix B: Electronic Information 523
11.1 General Computer Algebra Systems 523
11.2 Semi-general Computer Algebra Systems 524
11.3 More Specialized Packages and Programs 525
11.4 Specific Packages for Curves 526
11.5 Databases and Servers 527
11.6 Mailing Lists, Websites, and Newsgroups 529
11.7 Packages Not Directly Related to Number Theory 530

12. Appendix C: Tables 533
12.1 Hilbert Class Fields of Quadratic Fields 533
12.1.1 Hilbert Class Fields of Real Quadratic Fields 533
12.1.2 Hilbert Class Fields of Imaginary Quadratic Fields ... 538
12.2 Small Discriminants 543

12.2.1 Lower Bounds for Root Discriminants 543
12.2.2 Totally Complex Number Fields of Smallest Discriminant 545

Bibliography 549

Index of Notation 556

Index of Algorithms 564

General Index 569

1. Fundamental Results and Algorithms in Dedekind Domains

1.1 Introduction

The easiest way to start studying number fields is to consider them per se, as *absolute* extensions of $\mathbb{Q}$; this is, for example, what we have done in [Coh0]. In practice, however, number fields are frequently not given in this way. One of the most common other ways is to give a number field as a *relative* extension, in other words as an algebra L/K over some base field K that is not necessarily equal to $\mathbb{Q}$. In this case, the basic algebraic objects such as the ring of integers $\mathbb{Z}_L$ and the ideals of $\mathbb{Z}_L$ are not only $\mathbb{Z}$-modules, but also $\mathbb{Z}_K$-modules. The $\mathbb{Z}_K$-module structure is much richer and must be preserved. No matter what means are chosen to compute $\mathbb{Z}_L$, we have the problem of representing the result. Indeed, here we have a basic stumbling block: considered as $\mathbb{Z}$-modules, $\mathbb{Z}_L$ or ideals of $\mathbb{Z}_L$ are free and hence may be represented by $\mathbb{Z}$-bases, for instance using the Hermite normal form (HNF); see, for example, [Coh0, Chapter 2]. This theory can easily be generalized by replacing $\mathbb{Z}$ with any other explicitly computable Euclidean domain and, under certain additional conditions, to a principal ideal domain (PID). In general, $\mathbb{Z}_K$ is not a PID, however, and hence there is no reason for $\mathbb{Z}_L$ to be a free module over $\mathbb{Z}_K$. A simple example is given by $K = \mathbb{Q}(\sqrt{-10})$ and $L = K(\sqrt{-1})$ (see Exercise 22 of Chapter 2).

A remarkable fact, discovered independently by several authors (see [Bos-Poh] and [Coh1]) is that this stumbling block can easily be overcome, and there is no difficulty in generalizing most of the linear algebra algorithms for $\mathbb{Z}$-modules seen in [Coh0, Chapter 2] to the case of $\mathbb{Z}_K$-modules. This is the subject matter of the present chapter, which is essentially an expanded version of [Coh1].

Thus, the basic objects of study in this chapter are (finitely generated) modules over Dedekind domains, and so we will start by giving a detailed description of the main results about such modules. For further reading, I recommend [Frö-Tay] or [Bou1].

Note that, as usual, many theoretical results can be proved differently by using algorithmic methods. After finishing this chapter, and in particular after the study of the Hermite and Smith normal form algorithms over Dedekind domains, the reader is advised to try and prove the results of the next section using these algorithms.

1.2 Finitely Generated Modules Over Dedekind Domains

I would like to thank J. Martinet for his help in writing this section. For the sake of completeness, we first recall the following definitions.

Definition 1.2.1. *Let R be a domain, in other words a nonzero, commutative ring with unit, and no zero divisors.*

(1) *We say that R is* Noetherian *if every ascending chain of ideals of R is finite or, equivalently, if every ideal of R is finitely generated.*
(2) *We say that R is* integrally closed *if any x belonging to the ring of fractions of R which is a root of a* monic *polynomial in $R[X]$ belongs in fact to R.*
(3) *We say that R is a* Dedekind domain *if it is Noetherian, integrally closed, and if every nonzero prime ideal of R is a maximal ideal.*

Definition 1.2.2. *Let R be an integral domain and K its field of fractions. A* fractional ideal *is a finitely generated,* nonzero *sub-R-module of K or, equivalently, an R-module of the form I/d for some nonzero ideal I of R and nonzero $d \in R$. If we can take $d = 1$, the fractional ideal is an ordinary ideal, and we say that it is an* integral ideal.

Unless explicitly mentioned otherwise, we will *always* assume that ideals and fractional ideals are nonzero.

We recall the following basic facts about Dedekind domains, which explain their importance.

Proposition 1.2.3. *Let R be a Dedekind domain and K its field of fractions.*

(1) *Every fractional ideal of R is invertible and is equal in a unique way to a product of powers of prime ideals.*
(2) *Every fractional ideal is generated by at most two elements, and the first one can be an arbitrarily chosen nonzero element of the ideal.*
(3) *(Weak Approximation Theorem) Let S be a finite set of prime ideals of R, let $(e_{\mathfrak{p}})_{\mathfrak{p}\in S}$ be a set of integers, and let $(x_{\mathfrak{p}})_{\mathfrak{p}\in S}$ be a set of elements of K both indexed by S. There exists an element $x \in K$ such that for all $\mathfrak{p} \in S$, $v_{\mathfrak{p}}(x - x_{\mathfrak{p}}) = e_{\mathfrak{p}}$, while for all $\mathfrak{p} \notin S$, $v_{\mathfrak{p}}(x) \geq 0$, where $v_{\mathfrak{p}}(x)$ denotes the $\mathfrak{p}$-adic valuation.*
(4) *If K is a number field, the ring of integers $\mathbb{Z}_K$ of K is a Dedekind domain.*

In the context of number fields, we recall the following definitions and results.

Definition 1.2.4. *Let $|\ |$ be a map from K to the set of nonnegative real numbers.*

(1) *We say that* $|\ |$ *is a* field norm *on* K *if* $|x| = 0 \iff x = 0$, $|x+y| \leq |x| + |y|$, *and* $|xy| = |x|\,|y|$ *for all* x *and* y *in* K.
(2) *We say that the norm is* non-Archimedean *if we have the stronger condition* $|x+y| \leq \max(|x|, |y|)$ *for all* x *and* y *in* K; *otherwise, we say that the norm is* Archimedean.
(3) *We say that the norm is* trivial *if* $|x| = 1$ *for all* $x \neq 0$.
(4) *We say that two norms are* equivalent *if they define the same topology on* K.

Theorem 1.2.5 (Ostrowsky). *Let* K *be a number field and let* σ_i *be the* $n = r_1 + 2r_2$ *embeddings of* K *into* $\mathbb{C}$ *ordered in the usual way.*

(1) *Let* $\mathfrak{p}$ *be a prime ideal of* K. *Set*

$$|x|_{\mathfrak{p}} = \mathcal{N}(\mathfrak{p})^{-v_{\mathfrak{p}}(x)}$$

if $x \neq 0$, *and* $|0|_{\mathfrak{p}} = 0$ *otherwise. Then* $|x|_{\mathfrak{p}}$ *is a non-Archimedean field norm.*
(2) *Any nontrivial, non-Archimedean field norm is equivalent to* $|x|_{\mathfrak{p}}$ *for a unique prime ideal* $\mathfrak{p}$.
(3) *If* σ *is an embedding of* K *into* $\mathbb{C}$ *and if we set*

$$|x|_{\sigma} = |\sigma(x)| \ ,$$

where $|\ |$ *is the usual absolute value on* $\mathbb{C}$, *then* $|x|_{\sigma}$ *is an Archimedean field norm.*
(4) *Any Archimedean field norm is equivalent to* $|x|_{\sigma_i}$ *for a unique* σ_i *with* $1 \leq i \leq r_1 + r_2$. *(Note that* $|x|_{\sigma_{i+r_2}}$ *is equivalent to* $|x|_{\sigma_i}$ *for* $r_1 < i \leq r_1 + r_2$.*)*

Definition 1.2.6. *A* place *of a number field* K *is an equivalence class of nontrivial field norms. Thus, thanks to the above theorem, the places of* K *can be identified with the prime ideals of* K *together with the embeddings* σ_i *for* $1 \leq i \leq r_1 + r_2$.

Finally, we note the important *product formula* (see Exercise 1).

Proposition 1.2.7. *Let* $n_i = 1$ *for* $1 \leq i \leq r_1$, $n_i = 2$ *if* $r_1 < i \leq r_1 + r_2$. *Then, for all* $x \in K$ *we have*

$$\prod_{1 \leq i \leq r_1 + r_2} |x|_{\sigma_i}^{n_i} \prod_{\mathfrak{p}} |x|_{\mathfrak{p}} = 1 \ .$$

With these definitions, in the context of number fields we have a strengthening of Proposition 1.2.3 (3) to the case of *places* as follows.

Proposition 1.2.8 (Strong Approximation Theorem). *Let S be a finite set of places $|\ |_i$ of K, let $(x_i)_{i\in S}$ be a set of elements of K, and let $(\varepsilon_i)_{i\in S}$ be a set of positive real numbers both indexed by S. There exists $x \in K$ such that $|x - x_i|_i < \varepsilon_i$ for all $|\ |_i \in S$, while $|x|_i \leq 1$ for all places $|\ |_i \notin S$ except perhaps at one place not belonging to S, which can be arbitrarily chosen.*

Note that, due to the product formula, it is necessary to exclude one place, otherwise the proposition is trivially false (see Exercise 2). Clearly the weak approximation theorem is a consequence of the strong one (we choose for the excluded place any Archimedean one, since there always exists at least one). The following corollary is also important.

Corollary 1.2.9. *Let S_0 be a finite set of prime ideals of K, let $(e_{\mathfrak{p}})_{\mathfrak{p}\in S_0}$ be a set of integers indexed by S_0, and let $(s_\sigma)_{\sigma\in S_\infty}$ be a set of signs ± 1 indexed by the set S_∞ of all r_1 real embeddings of K. There exists an element $x \in K$ such that for all $\mathfrak{p} \in S_0$, $v_{\mathfrak{p}}(x) = e_{\mathfrak{p}}$, for all $\sigma \in S_\infty$, $\operatorname{sign}(\sigma(x)) = s_\sigma$, while for all $\mathfrak{p} \notin S_0$, $v_{\mathfrak{p}}(x) \geq 0$, where $v_{\mathfrak{p}}(x)$ denotes the $\mathfrak{p}$-adic valuation.*

Proof. Set $S = S_0 \cup S_\infty$ considered as a set of places of K thanks to Ostrowsky's theorem. For $\mathfrak{p} \in S_0$, we choose

$$y_{\mathfrak{p}} \in \mathfrak{p}^{e_{\mathfrak{p}}} \smallsetminus \mathfrak{p}^{e_{\mathfrak{p}}+1} \quad \text{and} \quad \varepsilon_{\mathfrak{p}} = \mathcal{N}(\mathfrak{p})^{-e_{\mathfrak{p}}}\ ,$$

while for $\sigma \in S_\infty$, we choose

$$y_\sigma = s_\sigma \quad \text{and} \quad \varepsilon_\sigma = \frac{1}{2}\ .$$

The strong approximation theorem implies that there exists $y \in K$ such that $|y - y_{\mathfrak{p}}|_{\mathfrak{p}} < \varepsilon_{\mathfrak{p}}$ for $\mathfrak{p} \in S_0$ and $|y - y_\sigma|_\sigma < \varepsilon_\sigma$ for $\sigma \in S_\infty$, and $|y|_{\mathfrak{p}} \leq 1$ for all $\mathfrak{p} \notin S$ except at most one such $\mathfrak{p}$.

The condition $|y - y_{\mathfrak{p}}|_{\mathfrak{p}} < \varepsilon_{\mathfrak{p}}$ is equivalent to $y - y_{\mathfrak{p}} \in \mathfrak{p}^{e_{\mathfrak{p}}+1}$; hence $v_{\mathfrak{p}}(y) = e_{\mathfrak{p}}$ by our choice of $y_{\mathfrak{p}}$.

Since $s_\sigma = \pm 1$, the condition $|y - y_\sigma|_\sigma < 1/2$ implies in particular that the sign of y is equal to s_σ.

Finally, if $\mathfrak{p} \notin S$, the condition $|y|_{\mathfrak{p}} \leq 1$ is evidently equivalent to $v_{\mathfrak{p}}(y) \geq 0$.

Thus y is almost the element that we need, except that we may have $v_{\mathfrak{p}_0}(y) < 0$ for some $\mathfrak{p}_0 \notin S$. Assume that this is the case (otherwise we simply take $x = y$), and set $v = -v_{\mathfrak{p}_0}(y) > 0$. By the *weak* approximation theorem, we can find an element π such that $v_{\mathfrak{p}_0}(\pi) = v$, $v_{\mathfrak{p}}(\pi) = 0$ for all $\mathfrak{p} \in S_0$, and $v_{\mathfrak{p}}(\pi) \geq 0$ for $\mathfrak{p} \notin S_0 \cup \{\mathfrak{p}_0\}$ (we can use the weak approximation theorem since we do not need to impose any Archimedean conditions on π). Since a square is positive, it is immediately checked that $x = \pi^2 y$ satisfies the desired properties. □

Corollary 1.2.10. *Let* $\mathfrak{m}$ *be any nonzero ideal. There exists* $\alpha \in \mathfrak{m}$ *such that for every prime ideal* $\mathfrak{p}$ *such that* $v_{\mathfrak{p}}(\mathfrak{m}) \neq 0$ *we have* $v_{\mathfrak{p}}(\alpha) = v_{\mathfrak{p}}(\mathfrak{m})$. *Such an element* α *will be called a* uniformizer *of the ideal* $\mathfrak{m}$.

Proof. This is an immediate consequence of Corollary 1.2.9. □

The two most important examples are the following: if $\mathfrak{m} = \mathfrak{p}$ is a prime ideal, then α is a uniformizer of $\mathfrak{p}$ if and only if $\alpha \in \mathfrak{p} \setminus \mathfrak{p}^2$; if $\mathfrak{m} = \mathfrak{p}^{-1}$ is the inverse of a prime ideal, then α is a uniformizer of $\mathfrak{p}^{-1}$ if and only if $\alpha \in \mathfrak{p}^{-1} \setminus \mathbb{Z}_K$.

Corollary 1.2.11. *Let* $\mathfrak{m}$ *be any (nonzero) integral ideal, and let* $\mathfrak{a}$ *be an ideal of* R. *There exists* $\alpha \in K^*$ *such that* $\alpha\mathfrak{a}$ *is an integral ideal coprime to* $\mathfrak{m}$; *in other words, in any ideal class there exists an integral ideal coprime to any fixed integral ideal.*

Proof. Indeed, apply the weak approximation theorem to the set of prime ideals $\mathfrak{p}$ that divide $\mathfrak{m}$ or such that $v_{\mathfrak{p}}(\mathfrak{a}) < 0$, taking $e_{\mathfrak{p}} = -v_{\mathfrak{p}}(\mathfrak{a})$. Then, if α is such that $v_{\mathfrak{p}}(\alpha) = e_{\mathfrak{p}}$ for all such $\mathfrak{p}$ and nonnegative for all other $\mathfrak{p}$, it is clear that $\alpha\mathfrak{a}$ is an integral ideal coprime to $\mathfrak{m}$. □

In this chapter, R will always denote a Dedekind domain and K its field of fractions. In the following sections, we will also assume that we can compute explicitly in R (this is, for example, the case if K is a number field), but for the theoretical part, we do not need this.

The main goal of this section is to prove the following results, which summarize the main properties of finitely generated modules over Dedekind domains (see below for definitions).

Theorem 1.2.12. *Let* M *be a finitely generated module over a Dedekind domain* R.

(1) *The* R*-module* M *is torsion-free if and only if* M *is a projective module.*
(2) *There exists a torsion-free submodule* N *of* M *such that*

$$M = M_{\text{tors}} \oplus N \qquad \text{and} \qquad N \simeq M/M_{\text{tors}} .$$

(3) *If* M *is a torsion-free* R*-module and* $V = KM$, *there exist (fractional) ideals* $\mathfrak{a}_i$ *and elements* $\omega_i \in V$ *such that*

$$M = \mathfrak{a}_1\omega_1 \oplus \mathfrak{a}_2\omega_2 \oplus \cdots \oplus \mathfrak{a}_n\omega_n .$$

The ideal class of the product $\mathfrak{a} = \mathfrak{a}_1\mathfrak{a}_2\cdots\mathfrak{a}_n$ *in the class group of* R *depends only on the module* M *and is called the* Steinitz class *of* M.
(4) *The module* M *is a free* R*-module if and only if its Steinitz class is equal to the trivial class, in other words if and only if* $\mathfrak{a}$ *is a principal ideal.*

(5) *If M is a torsion module, there exist unique nonzero integral ideals $\mathfrak{d}_i$ of R and (nonunique) elements $\omega_i \in M$ such that*

$$M = (R/\mathfrak{d}_1)\omega_1 \oplus \cdots \oplus (R/\mathfrak{d}_n)\omega_n$$

and $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$ for $2 \leq i \leq n$.

Corollary 1.2.13. *Let M be a finitely generated module over R of rank r. There exist fractional ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_n$, unique integral ideals $\mathfrak{d}_1, \ldots, \mathfrak{d}_n$ (possibly equal to zero), and elements $\omega_1, \ldots, \omega_n$ in M such that*

(1) $M = (\mathfrak{a}_1/\mathfrak{d}_1\mathfrak{a}_1)\omega_1 \oplus \cdots \oplus (\mathfrak{a}_n/\mathfrak{d}_n\mathfrak{a}_n)\omega_n$,
(2) $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$ *for* $2 \leq i \leq n$,
(3) $\mathfrak{d}_i = \{0\}$ *if and only if* $1 \leq i \leq r$.

We will prove these results completely in this section, and in passing we will also prove a number of important auxiliary results.

1.2.1 Finitely Generated Torsion-Free and Projective Modules

Definition and Proposition 1.2.14. *Let M be an R-module.*

(1) *We say that M is* finitely generated *if there exist $\alpha_1, \ldots, \alpha_n$ belonging to M such that any element x of M can be written (not necessarily uniquely) as $x = \sum_{i=1}^{n} x_i\alpha_i$ with $x_i \in R$.*
(2) *We define $KM = K \otimes_R M$; in other words,*

$$KM = (K \times M)/\mathcal{R} \ ,$$

where $\mathcal{R}$ is the equivalence relation defined by

$$\frac{a_1}{a_2}\alpha \ \mathcal{R} \ \frac{b_1}{b_2}\beta \iff \exists d \in R \setminus \{0\} \text{ such that } \ d(b_2a_1\alpha - a_2b_1\beta) = 0 \ ,$$

and with a natural definition of addition and multiplication.
(3) *If M is finitely generated, then KM is a finite-dimensional K-vector space, whose dimension is called the* rank *of the R-module M.*

Proof. All the assertions are clear, except perhaps for the fact that $\mathcal{R}$ is a transitive relation.

Assume that $(a_1/a_2)\alpha \ \mathcal{R} \ (b_1/b_2)\beta$ and $(b_1/b_2)\beta \ \mathcal{R} \ (c_1/c_2)\gamma$. Then, by definition, there exist nonzero elements d_1 and d_2 of R such that

$$d_1(b_2a_1\alpha - a_2b_1\beta) = d_2(c_2b_1\beta - b_2c_1\gamma) = 0 \in M \ .$$

Set $z = c_2a_1\alpha - a_2c_1\gamma$. We have

$$\begin{aligned} d_1d_2b_2z &= d_2c_2(d_1b_2a_1\alpha) - d_1a_2(d_2b_2c_1\gamma) \\ &= d_2c_2(d_1a_2b_1\beta) - d_1a_2(d_2c_2b_1\beta) = 0 \ . \end{aligned}$$

Since $d_1 \neq 0$, $d_2 \neq 0$, $b_2 \neq 0$, and R is an integral domain, it follows that $\mathcal{R}$ is an equivalence relation, as desired. □

Remark. It is easy to see that if we had defined $(a_1/a_2)\alpha \,\mathcal{R}\, (b_1/b_2)\beta \iff b_2a_1\alpha - a_2b_1\beta = 0$, this would in general not have been an equivalence relation (see Exercise 3).

Definition 1.2.15. *Let M be an R-module.*

(1) *The* torsion submodule *of M is defined by*

$$M_{\text{tors}} = \{x \in M \;/\; \exists a \in R \smallsetminus \{0\}, ax = 0\} \ .$$

An element of M_{tors} is called a torsion element.

(2) *We say that M is* torsion-free *if 0 is the only torsion element; in other words, if $M_{\text{tors}} = \{0\}$.*

(3) *We say that M is a* torsion module *if all the elements of M are torsion elements or, equivalently, if $M = M_{\text{tors}}$.*

Thus, the equivalence relation $\mathcal{R}$ defined above can also be given by saying that $(a_1/a_2)\alpha \;\mathcal{R}\; (b_1/b_2)\beta$ if and only if $b_2a_1\alpha - a_2b_1\beta$ is a torsion element. In particular, if $\lambda = a_1/a_2$, an element (λ, α) of KM is equal to zero if and only if $a_1\alpha$ is a torsion element, hence either if $\lambda = 0$ or if α itself is a torsion element.

For notational convenience, the equivalence class $\overline{(\lambda, \alpha)}$ in KM of a pair (λ, α) will be denoted $\lambda\alpha$. Note that when $\lambda \in R$, this is equal (modulo the equivalence relation) to the pair $(1, \lambda\alpha)$, and hence the two notations are compatible.

Note also that when M is torsion-free, the map $\alpha \mapsto \overline{(1, \alpha)}$ is injective, and hence in this case M can be considered as a sub-R-module of KM, and KM is simply the K-vector space spanned by M.

Definition and Proposition 1.2.16. *A module P is* projective *if it satisfies one of the following three equivalent conditions.*

(1) *Let f be a surjective map from a module F onto a module G. Then for any linear map g from P to G there exists a linear map h from P to F such that $g = f \circ h$ (see diagram below).*

(2) *If f is a surjective linear map from a module F onto P, there exists a* section *h of f, in other words a linear map from P to F such that $f \circ h = id_P$ (where id_P denotes the identity map on P).*

(3) *There exists a module P' such that $P \oplus P'$ is a free module.*

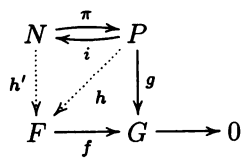

Proof. Let us prove that these conditions are equivalent. (1) implies (2) is obvious by taking $G = P$ and $g = id_P$. Assume (2), and let $(g_i)_{i\in I}$ be a (not necessarily finite) system of generators of P. Let $F = R^{(I)}$ be the set of maps v from I to R such that $v(i) = 0$ for all but a finite number of i. Then F is a free R-module with basis v_i such that $v_i(i) = 1$ and $v_i(j) = 0$ for $j \neq i$. Finally, let f be the map from F to P such that $f(v_i) = g_i$. By definition, f is a surjective linear map. By (2), we deduce that there exists a section h of f from P to F.

Set $P_1 = h(P)$. Since $f \circ h = id_P$, the map h is injective; hence P_1 is isomorphic to P. In addition, I claim that $F = P_1 \oplus \mathrm{Ker}(f)$. Indeed, for future reference, we isolate this as a lemma:

Lemma 1.2.17. *If f is a surjective map from any module F onto a projective module P and if h is a* section *of f (so that $f \circ h = id_P$), then $F = h(P) \oplus \mathrm{Ker}(f)$.*

Proof. Indeed, if $x \in F$, then $y = x - h(f(x))$ is clearly in $\mathrm{Ker}(f)$ since $f \circ h = id_P$; hence $x \in h(P) + \mathrm{Ker}(f)$, so $F = h(P) + \mathrm{Ker}(f)$. Furthermore, if $x \in h(P) \cap \mathrm{Ker}(f)$, then since $x \in h(P)$, $x = h(z)$ for some $z \in P$; hence since $x \in \mathrm{Ker}(f)$, $0 = f(x) = f(h(z)) = z$, hence $x = h(0) = 0$, so we have a direct sum, proving the lemma. □

This lemma implies Proposition 1.2.16 (3).

Finally, assume that $N = P \oplus P'$ is a free module, and let F, G, f, g be as in (1). Denote by π the projection from N to P defined by $\pi(p+p') = p$ if $p \in P$ and $p' \in P'$, denote by i the injection from P to N so that $\pi \circ i = id_P$, let $(u_i)_i$ be a basis of N, and set $g' = g \circ \pi$ (see preceding diagram).

Since f is surjective, we can find elements $v_i \in F$ such that $f(v_i) = g'(u_i)$. We arbitrarily fix such elements and set $h'\left(\sum_i x_i u_i\right) = \sum_i x_i v_i$. Since N is free, this is a well-defined linear map from N to F which clearly satisfies $g' = f \circ h'$; hence $g = g' \circ i = f \circ h' \circ i$, and so $h = h' \circ i$ satisfies (1). □

Note that the classical proof above is valid in any (commutative) ring, and not only in a Dedekind domain, and does not need the condition that the modules be finitely generated. Note also that the proof of (3) is essentially the proof that a free module is projective.

Corollary 1.2.18. *A projective module is torsion-free.*

Proof. Indeed, the third characterization of projective modules shows that a projective module is isomorphic to a submodule of a free module and hence is torsion-free since a free module is evidently torsion-free. □

The first important result of this section is the converse of this corollary for finitely generated modules over Dedekind domains.

Theorem 1.2.19. *Let M be a finitely generated, torsion-free module of rank n over a Dedekind domain R. Then M is a projective module. In addition, there exists an ideal I of R such that*

$$M \simeq R^{n-1} \oplus I \ .$$

Before proving this theorem we prove some lemmas.

Lemma 1.2.20. *If I and J are any fractional ideals of R, we have an isomorphism of R-modules:*

$$I \oplus J \simeq R \oplus IJ \ .$$

Proof. Since $I \simeq kI$ for any $k \in R$, we can always reduce to the case where I and J are integral ideals. By Corollary 1.2.11, in the ideal class of J there exists an integral ideal J_1 coprime to I. Thus, there exists $\alpha \in K^*$ such that $J_1 = \alpha J$, and it follows that $J_1 \simeq J$ and $IJ_1 \simeq IJ$, so we may replace J by J_1; in other words, we may assume that I and J are coprime integral ideals.

Let f be the map from $I \oplus J$ to R defined by $f(x,y) = x + y$. Since R is free, hence projective, and since $I + J = R$, f is surjective, so there exists a map g from R to $I \oplus J$ such that $f \circ g = id$. Lemma 1.2.17 says that $I \oplus J = g(R) \oplus \mathrm{Ker}(f)$. Since $f \circ g = id$, g is injective; hence $g(R) \simeq R$. Finally,

$$\mathrm{Ker}(f) = \{(x,-x)/x \in I, -x \in J\} = \{(x,-x)/x \in I \cap J\} \simeq I \cap J = IJ$$

since I and J are coprime, proving the lemma. □

Remark. We will see later how to transform this important isomorphism into an algorithmic equality (Corollary 1.3.6 and Proposition 1.3.12).

Corollary 1.2.21. *Every fractional ideal is a projective R-module.*

Proof. Simply apply the preceding lemma to $J = I^{-1}$ and use Proposition 1.2.16 (3). □

Lemma 1.2.22. *Let M be a finitely generated, torsion-free module of rank n, set $V = KM$, which is a K-vector space of dimension n, and let e be a nonzero element of V. Finally, set*

$$I = \{\lambda \in K/\lambda e \in M\} \ .$$

Then

(1) *I is a fractional ideal of R,*
(2) *M/Ie is a torsion-free R-module of rank $n-1$.*

Proof. (1). It is clear that I is a nonzero R-module. Since M is torsion-free, as an R-module, I is isomorphic to Ie (send x to xe), which is a submodule of the finitely generated module M. Since R is a Noetherian ring, a submodule of a finitely generated module is still finitely generated, hence I is finitely generated. It follows that I is a fractional ideal (take as denominator for I the product of denominators of generating elements of I).

(2). Let $\overline{x} \in M/Ie$ be a torsion element. Thus, there exists $a \in R \setminus \{0\}$ such that $ax \in Ie \subset Ke$, so $x \in Ke \cap M$. It follows that $x = \lambda e \in M$, hence $\lambda \in I$, so $x \in Ie$ or, equivalently, $\overline{x} = \overline{0}$, so M/Ie is torsion-free. We have $(M/Ie)K = (MK)/(Ke)$ and Ke is of dimension 1; hence M/Ie is of rank $n-1$. □

Proof of Theorem 1.2.19. We prove the theorem by induction on the rank of M. If the rank of M is zero, then M is torsion, and since M is torsion-free, $M = \{0\}$. Assume the theorem proved up to rank $n-1$, and let M be a torsion-free module of rank n. Let e be a nonzero element of M. By Lemma 1.2.22 above, M/Ie is a torsion-free module of rank $n-1$; hence by our induction hypothesis, M/Ie is a projective module and isomorphic to $R^{n-2} \oplus J$ for some ideal J (or is zero if $n = 1$). Lemma 1.2.17 implies that $M = g(M/Ie) \oplus Ie$ for a section g of the canonical surjective map from M to M/Ie, and since g is injective, $M \simeq M/Ie \oplus Ie$. Since M/Ie is projective by induction and $Ie \simeq I$ is also projective by Corollary 1.2.21, we deduce that M is projective. In addition, we have $M \simeq R^{n-2} \oplus J \oplus I \simeq R^{n-1} \oplus IJ$ by Lemma 1.2.20, thus showing our induction hypothesis and finishing the proof of Theorem 1.2.19. □

Before finishing this section, we must study in more detail the relationship between the module M and the ideal I such that $M \simeq R^{n-1} \oplus I$.

Theorem 1.2.23. *Let I be a fractional ideal of R. Then $R^{n-1} \oplus I$ is a free R-module if and only if I is a principal ideal.*

Proof. If I is a principal ideal, then $I \simeq R$; hence $R^{n-1} \oplus I \simeq R^n$ is free. Conversely, assume that $R^{n-1} \oplus I$ is free. Since I is of rank 1, we have $R^{n-1} \oplus I \simeq R^n$. Let f be an isomorphism from R^n to $R^{n-1} \oplus I$. Let $(e_i)_{1 \le i \le n}$ be the canonical basis of R^n. Any element $x \in K^n$ can be written uniquely as $x = \sum_{1 \le i \le n} x_i e_i$ for some $x_i \in K$. If we set $g(x) = \sum_{1 \le i \le n} x_i f(e_i)$, it is clear that $g(x) \in (R^{n-1} \oplus I)K = K^n$, that g is a well-defined isomorphism from K^n into itself such that the restriction of g to R^n is equal to f. In other words, g can be considered as an element of $\mathrm{GL}_n(K)$. Let $M = (a_{i,j})$ be the matrix of g on the canonical basis, so that $g(e_j) = f(e_j) = \sum_{1 \le i \le n} a_{i,j} e_i$ for all j. I claim that I is the principal ideal generated by $\det(M)$ — in other words, that $I = \det(M)R$.

Note first that by definition, for all j such that $1 \leq j \leq n$ we have $a_{i,j} \in R$ for $i < n$ and $a_{n,j} \in I$. If we expand $\det(M)$ along the bottom row, it immediately follows that $\det(M) \in I$, hence that $\det(M)R \subset I$.

Conversely, since f is surjective, it follows that for all $u \in I$ there exists $v = \sum_{1 \leq j \leq n} v_j e_j \in R^n$ such that $f(v) = ue_n$, which implies that $u = \sum_{1 \leq j \leq n} a_{n,j} v_j$; hence the $a_{n,j}$ generate the ideal I. Moreover, for any $i < n$, there exists $y_i = \sum_{1 \leq j \leq n} y_{i,j} e_j$ such that $f(y_i) = e_i$. Fix an index i_0, and let $X = (x_{i,j})$ be the $n \times n$ matrix defined by $x_{i,j} = y_{i,j}$ for $j < n$, $x_{i,n} = 0$ for $i \neq i_0$, and $x_{i_0,n} = 1$. It is clear that we have the block matrix equality

$$MX = \begin{pmatrix} I_{n-1} & C \\ 0 & a_{n,i_0} \end{pmatrix} ,$$

where I_{n-1} is the $(n-1) \times (n-1)$ identity matrix and C an $(n-1) \times 1$ column matrix. Taking determinants, we deduce that $a_{n,i_0} \in \det(M)R$. Since this is true for all i_0 and since the a_{n,i_0} generate the ideals I, it follows that $I \subset \det(M)R$; hence $I = \det(M)R$, as was to be proved.

Note that this proof is valid over any integral domain, not only over a Dedekind domain (I thank D. Bernardi for simplifying my initial proof). □

Corollary 1.2.24. *If I and J are two (fractional) ideals of R and $R^{m-1} \oplus I \simeq R^{n-1} \oplus J$, then $m = n$ and J and I are in the same ideal class (in other words, there exists $\alpha \in K^*$ such that $J = \alpha I$).*

Proof. Since I and J are of rank 1, it is clear that $m = n$. From the given isomorphism, we deduce that

$$R^{n-1} \oplus I \oplus I^{-1} \simeq R^{n-1} \oplus J \oplus I^{-1} .$$

Using Lemma 1.2.20, we obtain

$$R^{n+1} \simeq R^n \oplus JI^{-1} .$$

Thus Theorem 1.2.23 implies that JI^{-1} is a principal ideal, whence the corollary. □

This corollary shows that, if $M \simeq R^{n-1} \oplus I$ as in Theorem 1.2.19, the ideal class of I is well-defined and depends only on M. We will call it the *Steinitz class* of M and denote it by $\mathrm{St}(M)$.

We can restate the above results by saying that the isomorphism class of a finitely generated, torsion-free (or projective) module is completely classified by its rank and its Steinitz class.

Corollary 1.2.25. *Let M be a finitely generated, torsion-free module. There exist elements $\omega_1, \ldots, \omega_n$ in M and fractional ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_n$ of R such that*

$$M = \mathfrak{a}_1\omega_1 \oplus \cdots \oplus \mathfrak{a}_n\omega_n .$$

The Steinitz class of M is the ideal class of the product $\mathfrak{a}_1 \cdots \mathfrak{a}_n$.

Proof. From Theorem 1.2.19, we know that M is isomorphic to $R^{n-1} \oplus I$ for some ideal I whose ideal class is the Steinitz class of M. Replacing if necessary I by I/α for some nonzero element α of I, we may assume that $1 \in I$. Let f be the isomorphism from $R^{n-1} \oplus I$ to M, let $e_i = (0, \ldots, 1, \ldots, 0) \in R^{n-1} \oplus I$ (with 1 at the ith component), and let $\omega_i = f(e_i) \in M$. Since f is an isomorphism, we have $M = \mathfrak{a}_1\omega_1 \oplus \cdots \oplus \mathfrak{a}_n\omega_n$, with $\mathfrak{a}_i = R$ for $1 \leq i \leq n-1$ and $\mathfrak{a}_n = I$.

By Lemma 1.2.20 we have

$$\mathfrak{a}_1\omega_1 \oplus \cdots \oplus \mathfrak{a}_n\omega_n \simeq R^{n-1} \oplus (\mathfrak{a}_1 \cdots \mathfrak{a}_n) \ ,$$

so the corollary follows. □

Corollary 1.2.26. *Let M, N, and P be three finitely generated, torsion-free modules. Assume that $P \oplus M \simeq P \oplus N$. Then $M \simeq N$.*

Proof. Using Theorem 1.2.19, we have $M \simeq R^{m-1} \oplus \mathrm{St}(M)$, $N \simeq R^{n-1} \oplus \mathrm{St}(N)$, $P \simeq R^{p-1} \oplus \mathrm{St}(P)$, so that

$$R^{p+m-2} \oplus \mathrm{St}(P) \oplus \mathrm{St}(M) \simeq R^{p+n-2} \oplus \mathrm{St}(P) \oplus \mathrm{St}(N)$$

or, in other words, by Lemma 1.2.20,

$$R^{p+m-1} \oplus \mathrm{St}(P)\,\mathrm{St}(M) \simeq R^{p+n-1} \oplus \mathrm{St}(P)\,\mathrm{St}(N) \ .$$

We deduce from Corollary 1.2.24 that $m = n$ and that there exists $\alpha \in K$ such that $\mathrm{St}(P)\,\mathrm{St}(M) = \alpha\,\mathrm{St}(P)\,\mathrm{St}(N)$; hence $\mathrm{St}(M) = \alpha\,\mathrm{St}(N) \simeq \mathrm{St}(N)$ since $\mathrm{St}(P)$ is invertible, so $M \simeq N$. □

We end this section with the following two propositions.

Proposition 1.2.27. *Let*

$$0 \longrightarrow M' \longrightarrow M \longrightarrow M'' \longrightarrow 0$$

be an exact sequence of finitely generated, torsion-free modules. Then

$$M \simeq M' \oplus M'' \qquad \text{and} \qquad \mathrm{St}(M) = \mathrm{St}(M')\,\mathrm{St}(M'') \ .$$

Proof. The isomorphism follows immediately from Lemma 1.2.17: if f is the map from M to M'', there exists a map h from M'' to M such that $f \circ h = id_{M''}$ and $M = h(M'') \oplus \mathrm{Ker}(f) \simeq M'' \oplus M'$ since the sequence is exact. The required equality of Steinitz classes now follows immediately from Theorem 1.2.19 and Lemma 1.2.20. □

Proposition 1.2.28. *If R is a Dedekind domain with only a finite number of prime ideals, then R is a principal ideal domain.*

Proof. Let $\mathfrak{b}$ be the product of the (nonzero) prime ideals of R, which are finite in number. If $\mathfrak{c}$ is an ideal of R, by Corollary 1.2.11 we can find an $x \in K^*$ such that $x\mathfrak{c}$ is an integral ideal coprime to $\mathfrak{b}$. But this means that $x\mathfrak{c}$ is not divisible by any prime ideal of R, hence $x\mathfrak{c} = R$, and so $\mathfrak{c} = (1/x)R$ is a principal ideal, hence R is a principal ideal domain. □

1.2.2 Torsion Modules

We first show that one can split the study of finitely generated modules over a Dedekind domain into two essentially nonoverlapping parts: the torsion-free modules we have just studied (Corollary 1.2.25 in particular) and the torsion modules.

Proposition 1.2.29. *Let M be a finitely generated R-module, and let M_{tors} be the torsion submodule of M. Then there exists a torsion-free submodule N of M such that*

$$M = M_{\mathrm{tors}} \oplus N \ .$$

Proof. If $P = M/M_{\mathrm{tors}}$, then P is torsion-free. Indeed, if $\overline{y} \in P_{\mathrm{tors}}$, there exists $a \in R \setminus \{0\}$ such that $ay \in M_{\mathrm{tors}}$, and hence there exists $b \in R \setminus \{0\}$ such that $bay = 0$, so $y \in M_{\mathrm{tors}}$ since R is an integral domain, and so $\overline{y} = \overline{0}$. From Theorem 1.2.19, we deduce that P is a projective R-module. It follows that there exists a linear map h from P to M such that $f \circ h = id_P$, where we denote by f the canonical surjection from M onto $P = M/M_{\mathrm{tors}}$. From Lemma 1.2.17 we deduce that $M = h(P) \oplus M_{\mathrm{tors}}$, and, since h is injective, $N = h(P)$ is isomorphic to P, hence is projective (or torsion-free), thus proving the proposition. □

Thus, to finish our study of the structure of finitely generated modules over Dedekind domains, it remains only to study torsion modules. The main result is the following theorem.

Theorem 1.2.30. *Let M be a finitely generated torsion module over a Dedekind domain R. There exist nonzero integral ideals $\mathfrak{d}_1, \ldots, \mathfrak{d}_r$, different from R, and elements $\omega_j \in M$ such that*

(1) $M = (R/\mathfrak{d}_1)\omega_1 \oplus \cdots \oplus (R/\mathfrak{d}_r)\omega_r$,
(2) $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$ *for* $2 \leq i \leq r$.

The ideals $\mathfrak{d}_i$ are unique and depend only on the isomorphism class of M.

We first prove two lemmas that are of independent interest.

Lemma 1.2.31. *Let S be a finite set of prime ideals of R and let $x \in K^*$ such that $v_{\mathfrak{p}}(x) \geq 0$ for all $\mathfrak{p} \in S$. There exist n and d in R such that $x = n/d$ and d not divisible by any $\mathfrak{p}$ in S.*

Proof. Let $x = n/d$ with n and d in R, for the moment arbitrary. By the approximation theorem, there exists $b \in K$ such that

$$\forall \mathfrak{p} \in S,\ v_{\mathfrak{p}}(b) = -v_{\mathfrak{p}}(d) \qquad \text{and} \qquad \forall \mathfrak{p} \notin S,\ v_{\mathfrak{p}}(b) \geq 0\ .$$

It follows that for $\mathfrak{p} \in S$, $v_{\mathfrak{p}}(db) = 0$ and for $\mathfrak{p} \notin S$, $v_{\mathfrak{p}}(db) \geq 0$, so $db \in R$ and is not divisible by any $\mathfrak{p}$ in S. Since for all $\mathfrak{p} \in S$, $v_{\mathfrak{p}}(x) \geq 0$ or, equivalently, $v_{\mathfrak{p}}(n) \geq v_{\mathfrak{p}}(d)$, it follows that $v_{\mathfrak{p}}(nb) \geq v_{\mathfrak{p}}(db) = 0$ for $\mathfrak{p} \in S$ and $v_{\mathfrak{p}}(nb) \geq 0$ for $\mathfrak{p} \notin S$, hence $nb \in R$, so $x = (nb)/(db)$ is a suitable representation of x. □

Lemma 1.2.32. *Let $\mathfrak{a}$ be a nonzero integral ideal of R and set*

$$B = \{x \in K/\ \forall \mathfrak{p} \mid \mathfrak{a},\ v_{\mathfrak{p}}(x) \geq 0\}\ .$$

Then

(1)

$$B = \left\{x = \frac{n}{d}\ /\ n, d \in R, (dR, \mathfrak{a}) = 1\right\}\ ;$$

in other words, $B = S^{-1}R$, where S is the multiplicative set of elements of R coprime to $\mathfrak{a}$. (We write $(I, J) = 1$ for two integral ideals I and J to mean that they are coprime — in other words, that $I + J = R$.)

(2) *B is a principal ideal domain.*

Proof. (1). It is clear that if $(dR, \mathfrak{a}) = 1$, then $v_{\mathfrak{p}}(n/d) = v_{\mathfrak{p}}(n) \geq 0$ for all $\mathfrak{p} \mid \mathfrak{a}$, and hence $n/d \in B$. Conversely, let $x \in B$. Taking for S the set of prime ideals dividing $\mathfrak{a}$, it follows from Lemma 1.2.31 that one can write $x = n/d$ with n and d in R and d coprime to $\mathfrak{a}$, proving (1).

(2). It is clear that B is a ring, and it is also a domain since $B \subset K$. By general properties of rings of fractions $S^{-1}R$, we know that the prime ideals of B are exactly the ideals $S^{-1}\mathfrak{p}$ for the prime ideals $\mathfrak{p}$ such that $\mathfrak{p} \cap S = \varnothing$, hence in our case the prime ideals dividing $\mathfrak{a}$, which are finite in number. Since $B = S^{-1}R$ is also a Dedekind domain, it follows from Proposition 1.2.28 that B is a principal ideal domain. □

Proof of Theorem 1.2.30. Let $\mathfrak{a}$ be the annihilator of M in R, so that

$$\mathfrak{a} = \{x \in R/\ xM = \{0\}\}\ .$$

Clearly, $\mathfrak{a}$ is an R-module contained in R, hence is an integral ideal, and it is nonzero since M is a finitely generated torsion module (it is the intersection of the annihilators of some generators of M, hence a finite intersection of nonzero ideals). Call B the ring defined in Lemma 1.2.32 above. Then B is a principal ideal domain. Furthermore, if $x \in B$, then $x = n/d$ with $(dR, \mathfrak{a}) = 1$; hence $dR + \mathfrak{a} = R$. Multiplying by M, we obtain $dM = M$, hence $M = M/d$, and so $xM = nM/d \subset M$; hence $BM \subset M$, and so $BM = M$ since $1 \in B$. It

follows that M can be considered as a B-module instead of as an R-module. The main advantage is that B is a principal ideal domain. Since $R \subset B$, M is still a torsion module. Hence the structure theorem for modules over principal ideal domain applies and we deduce that

$$M \simeq B/\mathfrak{b}_1 \oplus \cdots \oplus B/\mathfrak{b}_r$$

for some integral ideals $\mathfrak{b}_i$ of B, not equal to $\{0\}$ or B, and such that $\mathfrak{b}_{i-1} \subset \mathfrak{b}_i$ for $2 \leq i \leq r$.

Since $B = S^{-1}R$ and $\mathfrak{b}_i = S^{-1}\mathfrak{d}_i$ for some ideal $\mathfrak{d}_i$ divisible only by prime ideals dividing $\mathfrak{a}$, we have $B/\mathfrak{b}_i \simeq R/\mathfrak{d}_i$, showing the existence of ideals $\mathfrak{d}_i$ such that $M \simeq \bigoplus R/\mathfrak{d}_i$. Let f be the isomorphism from $\bigoplus R/\mathfrak{d}_i$ to M. Then, if we let $\omega_i = f(0,\ldots,1,\ldots,0)$ (with 1 at the ith component), we have $M = \bigoplus(R/\mathfrak{d}_i)\omega_i$ as desired. The uniqueness statement follows from the uniqueness of the $\mathfrak{b}_i$. □

Thanks to Theorem 1.2.30, we can give the following definition.

Definition 1.2.33. (1) *Let M be a finitely generated torsion module over a Dedekind domain R, and let $\mathfrak{d}_i$ be the ideals given by Theorem 1.2.30. We will say that the $\mathfrak{d}_i$ are the* invariant factors *or the* elementary divisors *of M, and the ideal product $\mathfrak{a} = \mathfrak{d}_1 \cdots \mathfrak{d}_r$ will be called the* order-ideal *of the torsion module M.*

(2) *Let P and Q be two finitely generated, torsion-free R-modules having the same rank and such that $P \subset Q$. The order-ideal of the torsion module Q/P will be called the* index-ideal *of P into Q and denoted $[Q : P]$.*

(3) *More generally, if P and Q are two finitely generated, torsion-free R-modules having the same rank and such that $P \cap Q$ is also of the same rank, then the (fractional) index-ideal of P into Q is defined by the formula $[Q : P] = [Q : P \cap Q] \cdot [P : P \cap Q]^{-1}$.*

It is easy to see that the definition of the fractional index-ideal does not depend on the common submodule of P and Q that is chosen, as long as it is of maximal rank.

When $R = \mathbb{Z}$, the unique positive generator of the order-ideal of a finite $\mathbb{Z}$-module M is clearly equal to the order of M. When $R = \mathbb{Z}_K$ for some number field K, the order-ideal of a $\mathbb{Z}_K$-module M is a nonzero ideal $\mathfrak{a}$ of $\mathbb{Z}_K$, and by the multiplicativity of the norm, we can recover the order itself by the formula $|M| = |\mathbb{Z}_K/\mathfrak{a}| = \mathcal{N}(\mathfrak{a})$. Thus, the order-ideal is a richer invariant than the order.

We also have the following simple proposition.

Proposition 1.2.34. *Assume that there exist nonzero ideals $\mathfrak{a}_i$ such that an R-module M satisfies $M \simeq \bigoplus_{1\leq i\leq k} R/\mathfrak{a}_i$. Then the order-ideal of M is equal to $\prod_{1\leq i\leq k} \mathfrak{a}_i$.*

Proof. This immediately follows from the fact that the order-ideal is unchanged by module isomorphism, and that the order-ideal of a product of two modules is equal to the product of the order-ideals. □

We end this section with the elementary divisor theorem for torsion-free modules, which is now easy to prove using the above techniques.

Theorem 1.2.35. *Let M and N be two torsion-free (or projective) modules of rank m and n, respectively, such that $N \subset M$ (so $n \leq m$). There exist fractional ideals $\mathfrak{b}_1, \ldots, \mathfrak{b}_m$ of R, a basis $(e_1, \ldots, e_m)$ of $V = KM$, and integral ideals $\mathfrak{d}_1, \ldots, \mathfrak{d}_n$ such that*

$$M = \mathfrak{b}_1 e_1 \oplus \cdots \oplus \mathfrak{b}_m e_m, \ N = \mathfrak{d}_1\mathfrak{b}_1 e_1 \oplus \cdots \oplus \mathfrak{d}_n\mathfrak{b}_n e_n$$

and $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$ for $2 \leq i \leq n$.

The ideals $\mathfrak{d}_i$ (for $1 \leq i \leq n$) and the ideal classes *of the ideal products $\mathfrak{b}_1 \cdots \mathfrak{b}_n$ and $\mathfrak{b}_{n+1} \cdots \mathfrak{b}_m$ depend only on M and N.*

Proof. Let us first prove uniqueness, so let $\mathfrak{d}_i$ and $\mathfrak{b}_i$ be ideals as in the theorem. Since $\mathfrak{b}_i/\mathfrak{d}_i\mathfrak{b}_i \simeq R/\mathfrak{d}_i$, we have

$$M/N \simeq R/\mathfrak{d}_1 \oplus \cdots R/\mathfrak{d}_n \oplus R^{m-n} \ ,$$

hence $(M/N)_{\text{tors}} \simeq R/\mathfrak{d}_1 \oplus \cdots R/\mathfrak{d}_n$, so the uniqueness statement for the $\mathfrak{d}_i$ follows from the uniqueness statement of Theorem 1.2.30. Furthermore, $M \simeq \mathfrak{b}_1 \oplus \cdots \oplus \mathfrak{b}_m \simeq R^{m-1} \oplus \mathfrak{b}_1 \cdots \mathfrak{b}_m$ by Lemma 1.2.20, and similarly $N \simeq R^{n-1} \oplus \mathfrak{d}_1 \cdots \mathfrak{d}_n\mathfrak{b}_1 \cdots \mathfrak{b}_n$. By Corollary 1.2.24, the ideal class of $\mathfrak{d}_1 \cdots \mathfrak{d}_n\mathfrak{b}_1 \cdots \mathfrak{b}_n$ is well-defined, hence also that of $\mathfrak{b}_1 \cdots \mathfrak{b}_n$ since the $\mathfrak{d}_i$ are unique. Finally, the ideal class of $\mathfrak{b}_1 \cdots \mathfrak{b}_m$ is well-defined, hence also that of $\mathfrak{b}_{n+1} \cdots \mathfrak{b}_m$.

To prove the existence statement, we first reduce to the case where $m = n$ by writing $M/N = (M/N)_{\text{tors}} \oplus M'$ for some torsion-free module M', which can be done using Proposition 1.2.29. If we set $M'' = \{x \in M/\ x \bmod N \in (M/N)_{\text{tors}}\}$, then $M''/N = (M/N)_{\text{tors}}$. Hence, once suitable ideals $\mathfrak{d}_i$ and $\mathfrak{b}_i$ are found for the pair (M'', N), we add some extra ideals $\mathfrak{b}_i$ by using Theorem 1.2.19 applied to the torsion-free module M'.

Hence, we now assume that $m = n$, so M/N is a finitely generated torsion module. We prove the result by induction on n. Assume that $n \geq 1$ and that it is true for $n-1$. By Theorem 1.2.30, we have $M/N = \bigoplus_{1 \leq i \leq r} \mathfrak{d}_i\overline{\omega_i}$ for certain ideals $\mathfrak{d}_i$. Using the same method as in the proof of Theorem 1.2.19, we see that if $\mathfrak{b}_1 = \{x \in K/\ x\omega_1 \in M\}$, then $M = \mathfrak{b}_1\omega_1 \oplus g(M/\mathfrak{b}_1\omega_1)$, where g is a section of the canonical projection of M onto $M/\mathfrak{b}_1\omega_1$. Similarly, if $\mathfrak{c}_1 = \{x \in K/\ x\omega_1 \in N\}$, then $N = \mathfrak{c}_1\omega_1 \oplus g'(N/\mathfrak{c}_1\omega_1)$. Since $N \subset M$, we have $\mathfrak{c}_1 \subset \mathfrak{b}_1$, and in fact $\mathfrak{c}_1 = \mathfrak{b}_1\mathfrak{d}_1$, and in addition g' can be taken to be the restriction of g to $N/\mathfrak{c}_1\omega_1$. Thus, we apply our induction hypothesis to the modules $N\mathfrak{c}_1\omega_1 \subset M/\mathfrak{b}_1\omega_1$ of rank $n-1$, and we obtain the desired result. □

Remark. The reader will have noted that in many cases we have tried as much as possible to give *equalities* between modules, and not simply isomorphisms, even if the isomorphisms are canonical. This is essential in algorithmic practice.

We now have at our disposal the main theoretical results we will need about finitely generated modules over Dedekind domains. We will always implicitly assume that all R-modules are finitely generated.

In the next section, we will study the algorithmic aspects. The reader will notice that many of the algorithms that will be described give alternate proofs of the theoretical results.

1.3 Basic Algorithms in Dedekind Domains

From now on, R will denote a Dedekind domain in which it is possible to compute efficiently. The reader can think of $R = \mathbb{Z}_K$, since this is the only application that we have in mind (see [Coh0, Sections 4.6.1 and 4.6.2] for a brief overview). However, the ring R could also be a maximal order in a global field of positive characteristic, for example.

1.3.1 Extended Euclidean Algorithms in Dedekind Domains

Proposition 1.3.1. *Given two coprime integral ideals* $\mathfrak{a}$ *and* $\mathfrak{b}$ *in* R*, we can find in polynomial time elements* $a \in \mathfrak{a}$ *and* $b \in \mathfrak{b}$ *such that* $a + b = 1$.

Proof. Since this is a very simple but basic proposition, we give the proof as an algorithm.

Algorithm 1.3.2 (Extended Euclid in Dedekind Domains). Let R be a Dedekind domain in which one can compute, and let $(\omega_i)_{1 \le i \le n}$ be an integral basis chosen so that $\omega_1 = 1$ (it is easy to reduce to this case, and in practice it is always so). Given two coprime ideals $\mathfrak{a}$ and $\mathfrak{b}$ given by their HNF matrices A and B on this integral basis, this algorithm computes $a \in \mathfrak{a}$ and $b \in \mathfrak{b}$ such that $a + b = 1$.

1. [Apply Hermite] Let C be the $n \times 2n$ matrix obtained by concatenating A and B (we will denote this by $C \leftarrow (A|B)$). Using one of the polynomial-time algorithms for HNF reduction (see, for example, [Coh0, Section 2.4.2]), compute an HNF matrix H and a $2n \times 2n$ unimodular matrix U such that $CU = (0|H)$.
2. [Check if coprime] If H is not equal to the $n \times n$ identity matrix, output an error message stating that $\mathfrak{a}$ and $\mathfrak{b}$ are not coprime, and terminate the algorithm.
3. [Compute coordinates] Set $Z \leftarrow U_{n+1}$, the $(n+1)$st column of the matrix U, and let X be the n-component column vector formed by the top n components of Z.

4. [Terminate] Let a be the element of K whose coordinate vector on the integral basis is AX, and set $b \leftarrow 1-a$. Output a and b, and terminate the algorithm.

Indeed, the HNF of the matrix C is the HNF of the ideal $\mathfrak{a}+\mathfrak{b}$. Since $\mathfrak{a}$ and $\mathfrak{b}$ are coprime, it is the identity matrix. It follows that $CZ = (1,0,\ldots,0)^t$. If we split Z into its upper half X and its lower half Y, it is clear that AX and BY represent on the integral basis elements $a \in \mathfrak{a}$ and $b \in \mathfrak{b}$ such that $a+b=1$, and hence the algorithm is valid. □

Implementation Remarks

(1) It was, of course, not really necessary in the proof that the ideals be given by HNF matrices, but only by $\mathbb{Z}$-bases. If we really do have HNF bases, the first column of the matrix A of $\mathfrak{a}$ will correspond to a generator z_a of $\mathfrak{a} \cap \mathbb{Z}$, and similarly the first column of B will correspond to a generator z_b of $\mathfrak{b} \cap \mathbb{Z}$. Frequently, z_a and z_b will be coprime. In that case, the usual extended Euclidean algorithm will easily find u and v such that $uz_a + vz_b = 1$, and we can take $a = uz_a$ and $b = vz_b$.

(2) Since the algorithm underlying this proposition will be absolutely basic to all our algorithms on Dedekind domains, we must ensure that it gives results that are as reasonable as possible. Indeed, the elements a and b are not unique and can be modified by adding and subtracting from a and b, respectively, some element of the ideal product $\mathfrak{ab}$. Hence it would be nice to have an element $r \in \mathfrak{ab}$ such that $a-r$ is "small" (and then we replace a by $a-r$ and b by $b+r = 1-(a-r)$, which will also be "small"). In Algorithm 1.4.13 we will see how this can be done reasonably well.

(3) This is the most important part of this chapter, where we specifically use the fact that the Dedekind domain R is the ring of integers of a number field, so as to be able to compute a and b in polynomial time.

We now come to a theorem that is trivial to prove but is the basic tool for our algorithms. It is a generalization to Dedekind domains of the extended Euclidean algorithm, as follows.

Theorem 1.3.3. *Let $\mathfrak{a}$ and $\mathfrak{b}$ be two (fractional) ideals in R, let a and b be two elements of K not both equal to zero, and set $\mathfrak{d} = a\mathfrak{a} + b\mathfrak{b}$. There exist $u \in \mathfrak{a}\mathfrak{d}^{-1}$ and $v \in \mathfrak{b}\mathfrak{d}^{-1}$ such that $au + bv = 1$, and these elements can be found in polynomial time.*

Proof. If a (resp., b) is equal to zero, we can take $(u,v) = (0,1/b)$ (resp., $(u,v) = (1/a,0)$), since in that case we have $1/b \in \mathfrak{b}\mathfrak{d}^{-1} = R/b$ (resp., $1/a \in \mathfrak{a}\mathfrak{d}^{-1} = R/a$). So assume a and b are nonzero.

Set $I = a\mathfrak{a}\mathfrak{d}^{-1}$ and $J = b\mathfrak{b}\mathfrak{d}^{-1}$. By the definition of $\mathfrak{d}^{-1}$, I and J are integral ideals and we have $I+J=R$. By Proposition 1.3.1, we can thus find in polynomial time $e \in I$ and $f \in J$ such that $e+f=1$, and clearly $u = e/a$ and $v = f/b$ satisfy the conditions of the lemma. □

Remark. Although this proposition is very simple, we will see that the essential conditions $u \in \mathfrak{a}\mathfrak{d}^{-1}$ and $v \in \mathfrak{b}\mathfrak{d}^{-1}$ bring as much rigidity into the problem as in the case of Euclidean domains, and this proposition will be regularly used instead of the extended Euclidean algorithm. It is, in fact, clear that it *is* an exact generalization of the extended Euclidean algorithm. Note that this lemma is useful even when R is a principal ideal domain, since R is not necessarily Euclidean.

We also need the following.

Proposition 1.3.4. *Let* $\mathfrak{a}$, $\mathfrak{b}$, $\mathfrak{c}$, $\mathfrak{d}$ *be fractional ideals of* R, *and let* a, b, c, d *be elements of* K. *Set* $e = ad - bc$, *and assume that*

$$\mathfrak{a}\mathfrak{b} = e\mathfrak{c}\mathfrak{d}, \quad a \in \mathfrak{a}\mathfrak{c}^{-1}, \quad b \in \mathfrak{b}\mathfrak{c}^{-1}, \quad c \in \mathfrak{a}\mathfrak{d}^{-1}, \quad d \in \mathfrak{b}\mathfrak{d}^{-1} \ .$$

Finally, let x *and* y *be two elements of an* R*-module* M, *and set*

$$\begin{pmatrix} x' & y' \end{pmatrix} = \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} a & c \\ b & d \end{pmatrix} \ .$$

Then

$$\mathfrak{a}x + \mathfrak{b}y = \mathfrak{c}x' + \mathfrak{d}y' \ .$$

Proof. We have $x' = ax + by$ and $y' = cx + dy$; hence

$$\mathfrak{c}x' + \mathfrak{d}y' \subset (a\mathfrak{c} + c\mathfrak{d})x + (b\mathfrak{c} + d\mathfrak{d})y \subset \mathfrak{a}x + \mathfrak{b}y \ .$$

Conversely, we have $x = (dx' - by')/e$ and $y = (-cx' + ay')/e$; hence

$$\mathfrak{a}x + \mathfrak{b}y \subset \frac{1}{e}(a\mathfrak{b}\mathfrak{d}^{-1}x' + a\mathfrak{b}\mathfrak{c}^{-1}y') \ ,$$

and since $\mathfrak{a}\mathfrak{b} \subset e\mathfrak{c}\mathfrak{d}$,

$$\mathfrak{a}x + \mathfrak{b}y \subset \mathfrak{c}\mathfrak{d}(\mathfrak{d}^{-1}x' + \mathfrak{c}^{-1}y') = \mathfrak{c}x' + \mathfrak{d}y' \ ,$$

thus showing the double inclusion.

Note that, although we have used only the inclusion $\mathfrak{a}\mathfrak{b} \subset e\mathfrak{c}\mathfrak{d}$ in the proof, the hypotheses on a, b, c, and d imply that $e\mathfrak{c}\mathfrak{d} \subset \mathfrak{a}\mathfrak{b}$, so we must have equality. □

Corollary 1.3.5. *Let* $\mathfrak{a}$ *and* $\mathfrak{b}$ *be two ideals,* a *and* b *be two elements of* K *not both zero,* $\mathfrak{d} = a\mathfrak{a} + b\mathfrak{b}$, *and* $u \in \mathfrak{a}\mathfrak{d}^{-1}$, $v \in \mathfrak{b}\mathfrak{d}^{-1}$ *such that* $au + bv = 1$ *as given by Theorem 1.3.3.*

Let x *and* y *be two elements of an* R*-module* M, *and set*

$$\begin{pmatrix} x' & y' \end{pmatrix} = \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} b & u \\ -a & v \end{pmatrix} \ .$$

Then

$$\mathfrak{a}x + \mathfrak{b}y = a\mathfrak{b}\mathfrak{d}^{-1}x' + \mathfrak{d}y' \ .$$

Proof. Since $b \in \mathfrak{b}^{-1}\mathfrak{d}$ and $a \in \mathfrak{a}^{-1}\mathfrak{d}$, this is clearly a special case of Proposition 1.3.4 with $\mathfrak{c} = \mathfrak{a}\mathfrak{b}\mathfrak{d}^{-1}$. □

Corollary 1.3.6. *Let $\mathfrak{a}$, $\mathfrak{b}$ be two ideals. Assume that a, b, c, and d are four elements of K such that*

$$ad - bc = 1, \quad a \in \mathfrak{a}, \quad b \in \mathfrak{b}, \quad c \in \mathfrak{b}^{-1}, \quad d \in \mathfrak{a}^{-1} .$$

Let x and y be two elements of an R-module M, and set

$$\begin{pmatrix} x' & y' \end{pmatrix} = \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} a & c \\ b & d \end{pmatrix} .$$

Then

$$\mathfrak{a}x + \mathfrak{b}y = Rx' + \mathfrak{a}\mathfrak{b}y' .$$

Proof. This is also a special case of Proposition 1.3.4 with $\mathfrak{c} = R$ and $\mathfrak{d} = \mathfrak{a}\mathfrak{b}$. We will see in Proposition 1.3.12 how to find a, b, c, and d, given $\mathfrak{a}$ and $\mathfrak{b}$. □

Remarks

(1) The type of elementary transformation described in Proposition 1.3.4, particularly in its two corollaries above, will be the only one we are allowed to use. For example, if we want simply to replace x by $x - qy$ for some q in the field K (which is the usual elementary transformation), we must have $q \in \mathfrak{b}\mathfrak{a}^{-1}$, as can easily be checked.

(2) With the notation of Proposition 1.3.4, note that we also have the formal equality

$$\begin{pmatrix} \mathfrak{c}^{-1} & \mathfrak{d}^{-1} \end{pmatrix} = \begin{pmatrix} \mathfrak{a}^{-1} & \mathfrak{b}^{-1} \end{pmatrix} \begin{pmatrix} a & c \\ b & d \end{pmatrix} .$$

Indeed, since $a \in \mathfrak{a}\mathfrak{c}^{-1}$ and $b \in \mathfrak{b}\mathfrak{c}^{-1}$, it is clear that $a\mathfrak{a}^{-1} + b\mathfrak{b}^{-1} \subset \mathfrak{c}^{-1}$. Conversely, since $e = ad - bc$, we have $e \in a\mathfrak{b}\mathfrak{d}^{-1} + b\mathfrak{a}\mathfrak{d}^{-1}$, hence $e\mathfrak{c}\mathfrak{d} \subset a\mathfrak{b}\mathfrak{c} + b\mathfrak{a}\mathfrak{c}$, and since $\mathfrak{a}\mathfrak{b} = e\mathfrak{c}\mathfrak{d}$, we obtain the reverse inclusion $\mathfrak{c}^{-1} \subset a\mathfrak{a}^{-1} + b\mathfrak{b}^{-1}$. The second equality $\mathfrak{d}^{-1} = c\mathfrak{a}^{-1} + d\mathfrak{b}^{-1}$ is proved in a similar manner. We will see in Section 1.7 that the "real" reason for these identities is that for any nonzero ideal $\mathfrak{a}$, the ideal $\mathfrak{a}^{-1}$ can be canonically identified with the set of R-linear maps from $\mathfrak{a}$ to R (see Exercise 6).

1.3.2 Deterministic Algorithms for the Approximation Theorem

It will also be useful (although not essential) to have some algorithms linked to the approximation theorem in Dedekind domains. In this section, we give straightforward deterministic versions, but in practice it is much better to use the randomized methods that we explain in the next section.

Proposition 1.3.7. *Given ideals $\mathfrak{a}_i$ for $1 \leq i \leq k$ whose sum is equal to R, we can in polynomial time find elements $a_i \in \mathfrak{a}_i$ such that $\sum_i a_i = 1$.*

Proof. Same proof as for Proposition 1.3.1, except that we concatenate the k HNF matrices of the ideals and we split Z into k pieces at the end. Note that the matrix U will be an $nk \times nk$ unimodular matrix, which can become quite large. □

Proposition 1.3.8. *Let S be a finite set of prime ideals of R and let $(e_\mathfrak{p})_{\mathfrak{p}\in S} \in \mathbb{Z}^S$. There exists a polynomial-time algorithm that finds $a \in K$ such that $v_\mathfrak{p}(a) = e_\mathfrak{p}$ for $\mathfrak{p} \in S$ and $v_\mathfrak{p}(a) \geq 0$ for $\mathfrak{p} \notin S$.*

Proof. We can write $e_\mathfrak{p} = f_\mathfrak{p} - g_\mathfrak{p}$ with $f_\mathfrak{p} \geq 0$ and $g_\mathfrak{p} \geq 0$. If we can find n (resp., d) such that the conditions are satisfied with $e_\mathfrak{p}$ replaced by $f_\mathfrak{p}$ (resp., $g_\mathfrak{p}$), it is clear that $a = n/d$ satisfies our conditions. Thus, we may assume that $e_\mathfrak{p} \geq 0$ for $\mathfrak{p} \in S$. Following the classical proof (see, for example, [Coh0, Proposition 4.7.8]), we compute the ideal product

$$I = \prod_{\mathfrak{p}\in S} \mathfrak{p}^{e_\mathfrak{p}+1}$$

and we set for each $\mathfrak{p} \in S$

$$\mathfrak{a}_\mathfrak{p} = I \cdot \mathfrak{p}^{-e_\mathfrak{p}-1} \ .$$

Then the $\mathfrak{a}_\mathfrak{p}$ are integral ideals that sum to R, so by Proposition 1.3.7, we can in polynomial time find $a_\mathfrak{p} \in \mathfrak{a}_\mathfrak{p}$ whose sum is equal to 1. Furthermore, we can find $b_\mathfrak{p} \in \mathfrak{p}^{e_\mathfrak{p}} \setminus \mathfrak{p}^{e_\mathfrak{p}+1}$ (for example, by taking the $e_\mathfrak{p}$th power of an element of $\mathfrak{p} \setminus \mathfrak{p}^2$ which can be found in polynomial time). Then $a = \sum_{\mathfrak{p}\in S} a_\mathfrak{p} b_\mathfrak{p}$ is a solution to our problem. □

Corollary 1.3.9. *Given two integral ideals $\mathfrak{a}$ and $\mathfrak{b}$ of R such that the factorization of the norm of $\mathfrak{b}$ is known, there exists a polynomial-time algorithm that finds $x \in K$ such that $x\mathfrak{a}$ is an integral ideal coprime to $\mathfrak{b}$, and similarly finds $y \in K$ such that $y\mathfrak{a}^{-1}$ is an integral ideal coprime to $\mathfrak{b}$.*

Proof. For x, apply Proposition 1.3.8 to S equal to the prime ideal factors of $\mathfrak{b}$ and to $e_\mathfrak{p} = -v_\mathfrak{p}(\mathfrak{a})$ for all $\mathfrak{p} \in S$. For y, apply Proposition 1.3.8 to S equal to the prime ideal factors of $\mathfrak{a}$ and $\mathfrak{b}$ and to $e_\mathfrak{p} = v_\mathfrak{p}(\mathfrak{a})$ for all $\mathfrak{p} \in S$. □

Proposition 1.3.10. *Let $\mathfrak{a}$ be an integral ideal of R and $a \in \mathfrak{a}$, $a \neq 0$. Assume that the prime ideal factorization of a is known. Then there exists a polynomial-time algorithm that finds $b \in \mathfrak{a}$ such that $\mathfrak{a} = aR + bR$.*

Proof. Write $aR = \prod_\mathfrak{p} \mathfrak{p}^{e_\mathfrak{p}}$ with $e_\mathfrak{p} \geq 0$. Thus, $\mathfrak{a} = \prod_\mathfrak{p} \mathfrak{p}^{v_\mathfrak{p}(\mathfrak{a})}$ with $0 \leq v_\mathfrak{p}(\mathfrak{a}) \leq e_\mathfrak{p}$. By Proposition 1.3.8 we can, in polynomial time, find $b \in R$ such that $v_\mathfrak{p}(b) = v_\mathfrak{p}(\mathfrak{a})$ for all $\mathfrak{p} \mid a$; by looking at $\mathfrak{p}$-adic valuations, it is clear that $\mathfrak{a} = aR + bR$. □

Remarks

Recall that R is the ring of integers of a number field.

(1) If $\mathfrak{p}$ is a prime ideal given by a $\mathbb{Z}$-basis, the above proposition shows that we can, in polynomial time, find a two-element generating system for $\mathfrak{p}$. Indeed, we take $a = p$, and using the polynomial-time algorithm of Buchmann and Lenstra (see [Coh0, Algorithm 6.2.9]), we can factor pR into prime ideals so the condition is satisfied.

(2) To factor a it is enough to factor the absolute norm $\mathcal{N}(a) \in \mathbb{Z}$ of a, since we can use the Buchmann–Lenstra algorithm to factor into prime ideals the prime factors of $\mathcal{N}(a)$, then use [Coh0, Algorithm 4.8.17] for computing $\mathfrak{p}$-adic valuations, which is also polynomial-time as soon as a two-element generating set is known for every prime ideal $\mathfrak{p}$, which is the case by (1).

(3) As mentioned earlier, it is much faster in practice to perform a search for the elements that we need in Corollary 1.2.11 and Proposition 1.3.10. Of course, the time to perform this search is a priori exponential, but in practice it will always be very fast (see Algorithms 1.3.14 and 1.3.15 below).

The strong form of the approximation theorem can be dealt with in the same manner:

Proposition 1.3.11. *Let S be a finite set of prime ideals of R, let $(e_{\mathfrak{p}})_{\mathfrak{p}\in S} \in \mathbb{Z}^S$, and let $(x_{\mathfrak{p}})_{\mathfrak{p}\in S} \in K^S$. Then there exists a polynomial-time algorithm that finds $x \in K$ such that $v_{\mathfrak{p}}(x - x_{\mathfrak{p}}) = e_{\mathfrak{p}}$ for $\mathfrak{p} \in S$ and $v_{\mathfrak{p}}(x) \geq 0$ for $\mathfrak{p} \notin S$.*

Proof. Assume first that the $e_{\mathfrak{p}}$ are nonnegative and $x_{\mathfrak{p}} \in R$. Then we introduce the same ideals I and $\mathfrak{a}_{\mathfrak{p}}$ and elements $a_{\mathfrak{p}}$ as in the proof of Proposition 1.3.8. If we set

$$x = \sum_{\mathfrak{p}\in S} a_{\mathfrak{p}} x_{\mathfrak{p}} \ ,$$

it is easy to see that x satisfies the required conditions.

Consider now the general case. Let $d \in R$ be a common denominator for the $x_{\mathfrak{p}}$, and multiply d by suitable elements of R so that $e_{\mathfrak{p}} + v_{\mathfrak{p}}(d) \geq 0$ for all $\mathfrak{p} \in S$. According to what we have just proved, there exists $y \in R$ such that

$$\begin{aligned} \forall \mathfrak{p} \in S, \quad & v_{\mathfrak{p}}(y - dx_{\mathfrak{p}}) = e_{\mathfrak{p}} + v_{\mathfrak{p}}(d) \quad \text{and} \\ \forall \mathfrak{p} \mid d,\ \mathfrak{p} \notin S, \quad & v_{\mathfrak{p}}(y - dx_{\mathfrak{p}}) = v_{\mathfrak{p}}(d) \ . \end{aligned}$$

It follows that $x = y/d$ satisfies the given conditions. □

Finally, we show how to find elements satisfying Corollary 1.3.6.

Proposition 1.3.12. *Let $\mathfrak{a}$ and $\mathfrak{b}$ be two (fractional) ideals in R. Assume that the prime ideal factorization of $\mathfrak{a}$ or of $\mathfrak{b}$ is known. Then it is possible to find in polynomial time elements $a \in \mathfrak{a}$, $b \in \mathfrak{b}$, $c \in \mathfrak{b}^{-1}$, and $d \in \mathfrak{a}^{-1}$ such that $ad - bc = 1$.*

Proof. Multiplying if necessary $\mathfrak{a}$ and $\mathfrak{b}$ by an element of $\mathbb{Q}^*$, we can reduce to the case where $\mathfrak{a}$ and $\mathfrak{b}$ are integral ideals. Assume, for example, that the factorization of $\mathfrak{b}$ is known. According to Corollary 1.2.11, we can, in polynomial time, find $a \in R$ such that $a\mathfrak{a}^{-1}$ is an integral ideal (or, equivalently, $a \in \mathfrak{a}$) and coprime to $\mathfrak{b}$. According to Proposition 1.3.1, we can thus find $e \in a\mathfrak{a}^{-1}$ and $f \in \mathfrak{b}$ such that $e + f = 1$. Clearly, $b = f$, $c = -1$, and $d = e/a$ satisfy the required conditions. □

Remark. All of the above can also be done in polynomial time without knowing any prime ideal factorizations by using *factor refinement*, which we will not explain here (see [Bac-Sha1]).

1.3.3 Probabilistic Algorithms

The algorithms given above suffer from two defects. First, although they are polynomial-time, they are rather slow; second, the size of the computed objects will usually be unreasonably large. We have given the algorithms just to show their existence (in any case, they are all very easy), but in practice it is much better to use randomized algorithms, as is usually the case in computational problems. Although we have already done so, we explicitly specialize to $R = \mathbb{Z}_K$.

In all these randomized algorithms, we will have to pick at random elements from a given fractional ideal. This can be done in the following simple way.

Algorithm 1.3.13 (Random Element in an Ideal). Let $\mathfrak{a}$ be an ideal of a number field K of degree m over $\mathbb{Q}$ given by some generating system over $\mathbb{Z}$. This algorithm outputs a small random element of $\mathfrak{a}$.

1. [LLL-reduce] Using an algorithm for LLL-reduction, compute an LLL-reduced basis $(\alpha_i)_{1 \le i \le m}$ for the ideal $\mathfrak{a}$.
2. [Output random element] For $1 \le i \le m$, let x_i be randomly chosen integers such that $|x_i| \le 3$. Output $\sum_{1 \le i \le m} x_i \alpha_i$ and terminate the algorithm.

Remarks

(1) On the one hand, it is essential to do an LLL-reduction in the first step so as to have small elements. On the other hand, in practice this algorithm is not used as written since we will need several random elements from the same ideal $\mathfrak{a}$. Hence, we compute once and for all an LLL-reduced basis of $\mathfrak{a}$, and then execute step 2 as many times as necessary.

(2) The constant 3 used in step 2 is arbitrary but is more than sufficient for essentially all purposes. Probably the constant 2 would also be more than enough, and perhaps even the constant 1 for most applications. Since a factor of 3 in the size of the coefficients is usually not too costly, the constant 3 seems a good choice.

We now give simple-minded but efficient randomized versions of the algorithms implicit in Corollary 1.3.9, Proposition 1.3.10, and Proposition 1.3.12.

Algorithm 1.3.14 (Coprime Ideal Class). Given two integral ideals $\mathfrak{a}$ and $\mathfrak{b}$ of a number field K of degree m over $\mathbb{Q}$, this algorithm computes $\alpha \in K$ such that $\alpha\mathfrak{a}$ is an integral ideal coprime to $\mathfrak{b}$.

1. [Compute $\mathfrak{a}^{-1}$] Using [Coh0, Algorithm 4.8.21], compute the HNF of the ideal $\mathfrak{a}^{-1}$ on some fixed integral basis, then an LLL-reduced basis (α_i) of $\mathfrak{a}^{-1}$.
2. [Pick random element] Using the (α_i) and step 2 of Algorithm 1.3.13, pick a small random element $\alpha \in \mathfrak{a}^{-1}$.
3. [Check if OK] Form the $m \times 2m$ matrix M whose first m columns give the product of α by the basis elements of $\mathfrak{a}$, and the last m columns gives a $\mathbb{Z}$-basis of $\mathfrak{b}$ on the fixed integral basis. Compute the HNF of the ideal sum $\alpha\mathfrak{a} + \mathfrak{b}$ by computing the HNF of the matrix M. If this HNF is not equal to the identity matrix, go to step 2. Otherwise, output α and terminate the algorithm.

Since α is chosen in $\mathfrak{a}^{-1}$, we have $\alpha\mathfrak{a}+\mathfrak{b} = \mathbb{Z}_K$ if and only if $v_{\mathfrak{p}}(\alpha) = -v_{\mathfrak{p}}(\mathfrak{a})$ for every prime ideal $\mathfrak{p}$ dividing $\mathfrak{b}$. This occurs with probability $\prod_{\mathfrak{p}|\mathfrak{b}}(1 - 1/(\mathcal{N}(\mathfrak{p})))$, so the algorithm should be successful quite rapidly. □

We leave as a (trivial) exercise for the reader to write the corresponding algorithm for computing $\beta \in K$ such that $\beta\mathfrak{a}^{-1}$ is coprime to $\mathfrak{b}$ (Exercise 10). In fact, we will use it implicitly in Algorithm 1.3.16.

Remark. In this algorithm as well as in the following two, it is not really necessary to compute the full HNF of the matrix M, only the determinant of this HNF, which usually can be done much faster.

Algorithm 1.3.15 (Two-Element Representation). Given a fractional ideal $\mathfrak{a}$ in a number field K and a nonzero element $a \in \mathfrak{a}$, this algorithm computes $b \in \mathfrak{a}$ such that $\mathfrak{a} = a\mathbb{Z}_K + b\mathbb{Z}_K$.

1. [Compute an LLL-reduced basis] If not given in this form, compute first the HNF matrix A of the ideal $\mathfrak{a}$ on a fixed integral basis. Then, using an LLL algorithm, compute an LLL-reduced basis $(\alpha_i)_{1\le i\le m}$ of $\mathfrak{a}$.
2. [Compute matrix M_a] Compute the matrix M_a whose columns give on a fixed integral basis the product of a by the elements of the integral basis (thus M_a will be equal to aI_m if $a \in \mathbb{Q}$).
3. [Pick random element] Using the (α_i) and step 2 of Algorithm 1.3.13, pick a small random element $b \in \mathfrak{a}$, and compute the matrix M_b in a similar way as the matrix M_a.

4. [Check if OK] Compute the HNF of the matrix $(M_a|M_b)$ obtained by concatenating the matrices M_a and M_b. If it is not equal to A, go to step 3. Otherwise, output b and terminate the algorithm.

A similar analysis to the one made above shows that even though the algorithm may seem simple-minded, it is in fact rather efficient. □

Algorithm 1.3.16 ($ad-bc=1$ Algorithm). Given two fractional ideals $\mathfrak{a}$ and $\mathfrak{b}$, this algorithm outputs four elements a, b, c, and d such that $a \in \mathfrak{a}$, $b \in \mathfrak{b}$, $c \in \mathfrak{b}^{-1}$, $d \in \mathfrak{a}^{-1}$, and $ad - bc = 1$.

1. [Remove denominators] Let $d_1 \in \mathbb{Q}$ (or even in K) be a common denominator for the generators of $\mathfrak{a}$, and similarly d_2 for $\mathfrak{b}$, and set $\mathfrak{a} \leftarrow d_1\mathfrak{a}$, $\mathfrak{b} \leftarrow d_2\mathfrak{b}$.
2. [LLL-reduce] Using an LLL-algorithm, compute an LLL-reduced basis (α_i) of $\mathfrak{a}$.
3. [Compute $\mathfrak{a}^{-1}$] Using [Coh0, Algorithm 4.8.21], compute the HNF of $\mathfrak{a}^{-1}$ on some fixed integral basis.
4. [Pick random element] Using the (α_i) and step 2 of Algorithm 1.3.13, pick a small random element $\alpha \in \mathfrak{a}$.
5. [Check if OK] Form the $m \times 2m$ matrix M whose first m columns give the product of α by the basis elements of $\mathfrak{a}^{-1}$, and the last m columns give a $\mathbb{Z}$-basis of $\mathfrak{b}$ on the fixed integral basis. Compute the HNF of the ideal sum $\alpha\mathfrak{a}^{-1} + \mathfrak{b}$ by computing the HNF of the matrix M. If this HNF is not equal to the identity matrix, go to step 4.
6. [Euclidean step] Using Algorithm 1.3.2, compute $e \in \alpha\mathfrak{a}^{-1}$ and $f \in \mathfrak{b}$ such that $e + f = 1$.
7. [Terminate] Set $a \leftarrow \alpha/d_1$, $b \leftarrow f/d_2$, $c \leftarrow -d_2$, set $d \leftarrow ed_1/\alpha$ if $\alpha \neq 0$, $d \leftarrow d_1$ otherwise, and terminate the algorithm.

Remarks

(1) In step 5, if we keep the unimodular transformation matrix U of the HNF algorithm, the elements e and f necessary for step 6 can be read off immediately as in Algorithm 1.3.2 by looking at an appropriate column of U.

(2) The special case $\alpha = 0$ can occur only if $\mathfrak{b} = \mathbb{Z}_K$ (after step 1), and since in that case $a = 0$ and $bc = -1$, we can choose any value of d belonging to $\mathfrak{a}^{-1}$. Since after step 1, $\mathfrak{a}$ is an integral ideal, $1 \in \mathfrak{a}^{-1}$, and hence we may take $d = d_1$.

1.4 The Hermite Normal Form Algorithm in Dedekind Domains

In this section we will consider only finitely generated, torsion-free R-modules; we refer to Section 1.7 for torsion modules.

1.4.1 Pseudo-Objects

In view of Theorem 1.2.25, it is natural to give the following definition.

Definition 1.4.1. *Let M be a finitely generated, torsion-free R-module, and set $V = KM$.*

(1) *A* pseudo-element *of V is a sub-R-module of V of the form $\mathfrak{a}\omega$ with $\omega \in V$ and $\mathfrak{a}$ a fractional ideal of R or, equivalently, an equivalence class of pairs $(\omega, \mathfrak{a})$ formed by an element of V and a fractional ideal of R under the equivalence relation $(\omega, \mathfrak{a})\ \mathcal{R}\ (\omega', \mathfrak{a}')$ if and only if $\mathfrak{a}\omega = \mathfrak{a}'\omega'$ as sub-R-modules of rank 1 of V.*

(2) *The pseudo-element $\mathfrak{a}\omega$ is said to be* integral *if $\mathfrak{a}\omega \subset M$.*

(3) *If $\mathfrak{a}_i$ are fractional ideals of R and ω_i are elements of V, we say that $(\omega_i, \mathfrak{a}_i)_{1\le i\le k}$ is a* pseudo-generating set *for M if*

$$M = \mathfrak{a}_1\omega_1 + \cdots + \mathfrak{a}_k\omega_k \ .$$

(4) *We say that $(\omega_i, \mathfrak{a}_i)_{1\le i\le k}$ is a* pseudo-basis *of M if*

$$M = \mathfrak{a}_1\omega_1 \oplus \cdots \oplus \mathfrak{a}_k\omega_k \ .$$

Note that, according to Theorem 1.2.25, any finitely generated, torsion-free module has a pseudo-basis.

Let $(\omega_i, \mathfrak{a}_i)_{1\le i\le n}$ be a pseudo-basis of M. Then n is equal to the rank of M. It is clear that, among other transformations, we can multiply $\mathfrak{a}_i$ by a nonzero element of K as long as we divide ω_i by the same element, and we will still have a pseudo-basis. In particular, if so desired, we may assume that the $\mathfrak{a}_i$ are integral ideals, or that the ω_i are elements of M. On the other hand, it is generally not possible to have both properties at once. For example, let $M = \mathfrak{a}$ be a nonprincipal, primitive integral ideal. The general pseudo-basis of M is $(a, \mathfrak{a}/a)$, and so to have both an element of M and an integral ideal, we would need $a \in \mathfrak{a}$ and $\mathfrak{a}/a \subset R$, which is equivalent to $\mathfrak{a} = aR$, contrary to our choice of $\mathfrak{a}$.

Furthermore, restricting either to elements of M or to integral ideals would be too rigid for algorithmic purposes, so it is preferable not to choose a pseudo-basis of a particular type.

We will systematically represent finitely generated, torsion-free R-modules by pseudo-bases. To be able to do this, we need to know how to compute such pseudo-bases and how to perform usual operations on these pseudo-bases. As in the case of $R = \mathbb{Z}$, the basic algorithm for doing this is the Hermite normal form algorithm, and we will see that such an algorithm does indeed exist. Before doing this, however, let us see how one can go from one basis to another.

The following proposition is a generalization of Proposition 1.3.4.

Proposition 1.4.2. *Let $(\omega_i, \mathfrak{a}_i)_i$ and $(\eta_j, \mathfrak{b}_j)_j$ be two pseudo-bases for an R-module M, and let $U = (u_{i,j})$ be the $n \times n$ matrix giving the η_j in terms of the ω_i (so that $(\eta_1, \dots, \eta_n) = (\omega_1, \dots, \omega_n)U$).*

Set $\mathfrak{a} = \mathfrak{a}_1 \cdots \mathfrak{a}_n$ and $\mathfrak{b} = \mathfrak{b}_1 \cdots \mathfrak{b}_n$. Then $u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$ and $\mathfrak{a} = \det(U)\mathfrak{b}$ (note that, by Theorem 1.2.25, we know that $\mathfrak{a}$ and $\mathfrak{b}$ are in the same ideal class). Conversely, if there exist ideals $\mathfrak{b}_j$ such that $\mathfrak{a} = \det(U)\mathfrak{b}$ (with $\mathfrak{b} = \mathfrak{b}_1 \cdots \mathfrak{b}_n$) and $u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$, then $(\eta_j, \mathfrak{b}_j)_j$ is a pseudo-basis of M, where the η_j are given in terms of the ω_i by the columns of U.

Proof. Since

$$\eta_j \in \mathfrak{b}_j^{-1} M = \mathfrak{b}_j^{-1} \bigoplus_{i=1}^{n} \mathfrak{a}_i \omega_i = \bigoplus_{i=1}^{n} \mathfrak{a}_i \mathfrak{b}_j^{-1} \omega_i \ ,$$

it follows that $u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$.

It is easily proven by linearity or by induction on n that $e = \det(U) \in \mathfrak{a}\mathfrak{b}^{-1}$, so $e\mathfrak{b} \subset \mathfrak{a}$. Similarly, the matrix U^{-1} expresses the ω_j in terms of the η_i, so $\det(U^{-1}) \in \mathfrak{b}\mathfrak{a}^{-1}$. But since $\det(U^{-1}) = 1/e$, we have $\mathfrak{a}/e \subset \mathfrak{b}$ or, equivalently, $\mathfrak{a} \subset e\mathfrak{b}$, from which it follows that $\mathfrak{a} = e\mathfrak{b}$.

Conversely, if U has the above properties, by looking at the adjoint matrix of U it is easy to see that U^{-1} is of a similar form with $\mathfrak{a}$ and $\mathfrak{b}$ exchanged (it is of course essential that $\mathfrak{a} = \det(U)\mathfrak{b}$). If $X = (x_1, \dots, x_n)^t$ is the column vector of components of an element m of M in the pseudo-basis $(\omega_i, \mathfrak{a}_i)_i$, then $m = (\omega_1, \dots, \omega_n)X = (\eta_1, \dots, \eta_n)U^{-1}X$, and $U^{-1}X = (y_1, \dots, y_n)^t$ satisfies $y_i \in \mathfrak{b}_i$ for $1 \le i \le n$. Since the y_i are unique, this shows that $(\eta_j, \mathfrak{b}_j)_j$ is a pseudo-basis of M, proving the proposition. □

It is clear that Proposition 1.3.4 is the special case $n = 2$ of this proposition. Since that special case is constantly used, however, we have presented it separately.

Corollary 1.4.3. *Let M be a finitely generated, torsion-free R-module together with a nondegenerate, bilinear pairing $T(x, y)$ from $M \times M$ to R (for example, $M = \mathbb{Z}_L$, where L is a number field containing K, and $T(x, y) = \mathrm{Tr}_{L/K}(x \cdot y)$). For any pseudo-basis $\mathcal{B} = (\omega_j, \mathfrak{a}_j)$ of M, let $\mathrm{disc}_T(\mathcal{B})$ be the ideal defined by $\mathrm{disc}_T(\mathcal{B}) = \det(T(\omega_i, \omega_j))\mathfrak{a}^2$, where as usual $\mathfrak{a} = \mathfrak{a}_1 \cdots \mathfrak{a}_n$. Then if $\mathcal{B}' = (\eta_j, \mathfrak{b}_j)$ is another pseudo-basis of M, we have $\mathrm{disc}_T(\mathcal{B}') = \mathrm{disc}_T(\mathcal{B})$.*

Proof. Note that, since in general $\omega_j \notin M$, in the above definition we extend the bilinear form T to $V \times V$ (where $V = KM$) by bilinearity.

Let U be the matrix expressing the η_j in terms of the ω_i. We know that $\mathfrak{a} = \det(U)\mathfrak{b}$. By bilinearity, it is clear that if G (resp., G') is the matrix of the $T(\omega_i, \omega_j)$ (resp., $T(\eta_i, \eta_j)$), then $G' = U^t G U$. It follows that

$$\mathrm{disc}_T(\mathcal{B}') = \det(G')\mathfrak{b}^2 = \det(G)\det(U)^2\mathfrak{a}^2/\det(U)^2 = \det(G)\mathfrak{a}^2 = \mathrm{disc}_T(\mathcal{B}).$$

□

Since $\mathrm{disc}_T(\mathcal{B})$ does not depend on the chosen pseudo-basis $\mathcal{B}$, we will denote it by $\mathfrak{d}_T(M)$ and call it the *discriminant ideal* of M with respect to the pairing $T(x,y)$.

Remark. We can also define $\det(T(\omega_i,\omega_j))$ as an element $\overline{d_T(M)} \in K^*/K^{*2}$, since, under a change of pseudo-basis, this determinant is multiplied by $\det(U)^2 \in K^{*2}$. The pair $\mathrm{disc}_T(M) = (\mathfrak{d}_T(M), d_T(M))$ will simply be called the *discriminant* of M with respect to T. Note that knowledge of one of the components of the pair does not imply knowledge of the other; hence the pair itself is useful. In the absolute case where $M = \mathbb{Z}_K$ is the ring of integers of a number field K considered as a $\mathbb{Z}$-module and T is the trace, the discriminant ideal $\mathfrak{d}_T(M)$ gives the absolute value of the usual discriminant, and $d_T(M)$ gives its sign (and some other information already contained in $\mathfrak{d}_T(M)$).

Since we represent finitely generated, torsion-free modules by pseudo-bases, we must also explain how to represent linear maps between such modules. This is done using the following proposition, which is, of course, similar in nature to Proposition 1.4.2.

Proposition 1.4.4. *Let $(\omega_i,\mathfrak{a}_i)_i$ be a pseudo-basis for a finitely generated, torsion-free module M, and similarly $(\omega'_j,\mathfrak{a}'_j)_j$ for a module M'. Let f be a K-linear map from M' to M. There exists a matrix $A = (a_{i,j})$ such that $a_{i,j} \in \mathfrak{a}_i\mathfrak{a}_j'^{-1}$ and*

$$f\left(\sum_j a'_j\omega'_j\right) = \sum_i \left(\sum_j a_{i,j}a'_j\right)\omega_i \ .$$

Conversely, if $A = (a_{i,j})$ is such that $a_{i,j} \in \mathfrak{a}_i\mathfrak{a}_j'^{-1}$ for all i, j, the above formula defines a K-linear map f from M' to M.

Proof. The (very easy) proof is left to the reader (Exercise 11). The matrix A will of course be called the matrix of the map f on the chosen pseudo-bases of M' and M. Note that we need only a matrix and not a pseudo-matrix (see Definition 1.4.5) to represent a map. Thus, we will represent maps by such matrices A. □

1.4.2 The Hermite Normal Form in Dedekind Domains

The main theorem of this section is that the notion of Hermite normal form can be extended to Dedekind domains. As is well known, the Hermite normal form algorithm is a direct generalization of the extended Euclidean algorithm. Since we now have such an algorithm available to us (Theorem 1.3.3), it is not surprising that this can be done.

We first introduce a definition.

Definition 1.4.5. (1) *A* pseudo-matrix *is a pair* (A, I)*, where* $A = (a_{i,j})$ *is an* $n \times k$ *matrix with entries in* K*, and* $I = (\mathfrak{a}_i)$ *is a list of* k *fractional ideals.*

(2) *The* map *associated with this pseudo-matrix is the map* f *from* $\mathfrak{a}_1 \times \cdots \times a_k$ *to* K^n *defined by* $f(a_1, \ldots, a_k) = \sum_{1 \le j \le k} a_j A_j$*, where the* A_j *are the columns of* A*.*

(3) *The* module *associated with this pseudo-matrix is the module* $M = \sum_{1 \le j \le k} \mathfrak{a}_j A_j \subset K^n$*, or in other words the image of the map* f*, so that* $(A_j, \mathfrak{a}_j)$ *is a pseudo-generating set for* M*. We will also call this module the* image *of the pseudo-matrix* (A, I)*.*

(4) *The* kernel *of the pseudo-matrix* (A, I) *is the kernel of the associated map* f*.*

Theorem 1.4.6 (Hermite Normal Form in Dedekind Domains). *Let* (A, I) *be a pseudo-matrix, where* $I = (\mathfrak{a}_i)$ *is a list of* k *fractional ideals, and* $A = (a_{i,j})$ *is an* $n \times k$ *matrix. Assume that* A *is of rank* n *(so* $k \ge n$*) with entries in the field of fractions* K *of* R *(we could just as easily consider the case of a matrix of lower rank). Let* $M = \sum_j \mathfrak{a}_j A_j$ *be the* R*-module associated with the pseudo-matrix* (A, I)*. There exist* k *nonzero ideals* $(\mathfrak{b}_j)_{1 \le j \le k}$ *and a* $k \times k$ *matrix* $U = (u_{i,j})$ *satisfying the following conditions, where we set* $\mathfrak{a} = \mathfrak{a}_1 \cdots \mathfrak{a}_k$ *and* $\mathfrak{b} = \mathfrak{b}_1 \cdots \mathfrak{b}_k$*.*

(1) *For all* i *and* j *we have* $u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$*.*

(2) *We have* $\mathfrak{a} = \det(U)\mathfrak{b}$*.*

(3) *The matrix* AU *is of the following form:*

$$AU = \begin{pmatrix} 0 & 0 & \ldots & 0 & 1 & * & \ldots & * \\ 0 & 0 & \ldots & 0 & 0 & 1 & \ldots & * \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 & 0 & \ldots & 0 & 1 \end{pmatrix},$$

where the first $k - n$ *columns are zero (we will write this in abbreviated form as* $AU = (0|H)$*).*

(4) *If we call* ω_j *the elements corresponding to the nonzero columns of* AU *and* $\mathfrak{c}_j = \mathfrak{b}_{k-n+j}$ *for* $1 \le j \le n$*, then*

$$M = \mathfrak{c}_1 \omega_1 \oplus \cdots \oplus \mathfrak{c}_n \omega_n \ ;$$

in other words, $(\omega_j, \mathfrak{c}_j)_{1 \le j \le n}$ *is a pseudo-basis of the image* M *of the pseudo-matrix* (A, I)*.*

(5) *If we denote by* U_j *the columns of* U*, then* $(U_j, \mathfrak{b}_j)_{1 \le j \le k-n}$ *is a pseudo-basis for the kernel of the pseudo-matrix* (A, I)*.*

Proof. We give the proof of the existence of the HNF as an algorithm, very similar to [Coh0, Algorithm 2.4.5], which is the naive HNF algorithm.

Algorithm 1.4.7 (HNF Algorithm in Dedekind Domains). Given an $n \times k$ matrix $A = (a_{i,j})$ of rank n, and k (fractional) ideals $\mathfrak{a}_j$ in a number field K, this algorithm computes k ideals $\mathfrak{b}_j$ and a $k \times k$ matrix U such that these data satisfy the conditions of Theorem 1.4.6. We will make use only of elementary transformations of the type given in Theorem 1.3.3 combined with Corollary 1.3.5. We denote by A_j (resp., U_j) the columns of A (resp., U).

1. [Initialize] Set $i \leftarrow n$, $j \leftarrow k$, and let U be the $k \times k$ identity matrix.
2. [Check zero] Set $m \leftarrow j$, and while $m \geq 1$ and $a_{i,m} = 0$, set $m \leftarrow m - 1$. If $m = 0$, the matrix A is not of rank n, so print an error message and terminate the algorithm. Otherwise, if $m < j$, exchange A_m with A_j, $\mathfrak{a}_m$ with $\mathfrak{a}_j$, U_m with U_j, and set $m \leftarrow j$.
3. [Put 1 on the main diagonal] Set $A_j \leftarrow A_j/a_{i,j}$, $U_j \leftarrow U_j/a_{i,j}$, and $\mathfrak{a}_j \leftarrow a_{i,j}\mathfrak{a}_j$. (We now have $a_{i,j} = 1$.)
4. [Loop] If $m = 1$, go to step 6. Otherwise, set $m \leftarrow m - 1$, and if $a_{i,m} = 0$, go to step 4.
5. [Euclidean step] (Here $a_{i,j} = 1$ and $a_{i,m} \neq 0$.) Using the algorithm contained in the proof of Theorem 1.3.3, set $\mathfrak{d} = a_{i,m}\mathfrak{a}_m + \mathfrak{a}_j$ and find $u \in \mathfrak{a}_m\mathfrak{d}^{-1}$ and $v \in \mathfrak{a}_j\mathfrak{d}^{-1}$ such that $a_{i,m}u + v = 1$. Then set $(A_m, A_j) \leftarrow (A_m - a_{i,m}A_j, uA_m + vA_j)$, $(U_m, U_j) \leftarrow (U_m - a_{i,m}U_j, uU_m + vU_j)$, and $(\mathfrak{a}_m, \mathfrak{a}_j) \leftarrow (\mathfrak{a}_m\mathfrak{a}_j\mathfrak{d}^{-1}, \mathfrak{d})$. Finally, go to step 4.
6. [Final reductions of row i] For $m = j + 1, \ldots, n$, find $q \in \mathfrak{a}_m\mathfrak{a}_j^{-1}$ such that $a_{i,m} - q$ is small (see below), and set $A_m \leftarrow A_m - qA_j$ and $U_m \leftarrow U_m - qU_j$.
7. [Finished?] If $i = 1$, then output the matrix U, the modified matrix A (the matrix AU in the notation of Theorem 1.4.6), and the modified ideals $\mathfrak{a}_j$ (or $\mathfrak{b}_j$ in the notation of Theorem 1.4.6), and terminate the algorithm. Otherwise, set $i \leftarrow i - 1$, $j \leftarrow j - 1$, and go to step 2.

Proof of Theorem 1.4.6 and Algorithm 1.4.7.

Ignoring step 6 for the moment, we clearly see that this algorithm, which is essentially identical to the one for $\mathbb{Z}$, terminates with a new matrix A of the form required by Theorem 1.4.6. Furthermore, the elementary transformations that are used are either exchanges of columns (and the corresponding ideals) or transformations allowed by Corollary 1.3.5; hence the module $\mathfrak{a}_1\omega_1 + \cdots + \mathfrak{a}_k\omega_k$ stays unchanged.

Call $\mathfrak{a}$ the initial ideal product and $\mathfrak{b}$ the current one. All the elementary operations are of determinant ± 1 (in which case $\mathfrak{b}$ is unchanged), except in step 3 where the determinant is $1/a_{i,j}$ and $\mathfrak{b}$ is multiplied by $a_{i,j}$; hence the relation $\mathfrak{a} = \det(U)\mathfrak{b}$ is preserved throughout. We also clearly have $u_{i,j} \in \mathfrak{a}_i\mathfrak{b}_j^{-1}$. This shows (1), (2), and (3) of the theorem.

Upon termination we have a direct sum, and not simply a sum, since the last n columns of A are then linearly independent, showing (4).

Finally, let us prove (5). Since for all i, j we have $u_{i,j} \in \mathfrak{a}_i \mathfrak{b}_j^{-1}$ and $AU_j = 0$ for $1 \leq j \leq k-n$, it is clear that $(U_j, \mathfrak{b}_j)$ belongs to the kernel of (A, I) for $1 \leq j \leq k-n$. Conversely, let $X \in \mathfrak{a}_1 \times \cdots \times \mathfrak{a}_k$ be an element of the kernel of (A, I). Set $Y = U^{-1}X = (y_1, \ldots, y_k)^t$. Since U is invertible, $AX = 0$ if and only if $AUU^{-1}X = AUY = 0$ and, using the special form of the matrix AU, if and only if $y_j = 0$ for $k-n+1 \leq j \leq k$. Hence, $AX = 0$ if and only if $X = UY = \sum_{1 \leq j \leq k-n} y_j U_j$. By symmetry with (1), $U^{-1} = (v_{i,j})$ with $v_{i,j} \in \mathfrak{b}_i \mathfrak{a}_j^{-1}$, hence $y_i \in \mathfrak{b}_i$ so $X \in \sum_{1 \leq j \leq k-n} \mathfrak{b}_j U_j$, as was to be proved. We will come back to step 6 of the algorithm in Section 1.4.3. □

Remark. Note that this proof gives an algorithm to find an HNF of a matrix, but this algorithm is certainly not polynomial-time since the corresponding naive algorithm for HNF over $\mathbb{Z}$ is already not polynomial-time because of coefficient explosion. The existence of a polynomial-time algorithm for HNF reduction (including finding the matrix U) is rather recent (see [Haf-McC]). Note that in practice, n will be the relative degree of number fields extensions, and so in many cases the naive algorithm will be sufficient.

We now consider the problem of uniqueness in Theorem 1.4.6. We first need a definition.

Definition 1.4.8. *Let (A, I) be a pseudo-matrix with $I = (\mathfrak{a}_j)$. If $i_1, \ldots, i_r$ are r distinct rows of A and $j_1, \ldots, j_r$ are r distinct columns, we define the* minor-ideal *corresponding to these indices as follows. Let d be the determinant of the $r \times r$ minor extracted from the given rows and columns of A. Then the minor-ideal is the ideal $d\mathfrak{a}_{j_1} \cdots \mathfrak{a}_{j_r}$.*

With this definition we can state the following result.

Theorem 1.4.9. *With the notation of Theorem 1.4.6, for $1 \leq j \leq n$, set $\mathfrak{c}_j = \mathfrak{b}_{k-n+j}$. Then the ideals $\mathfrak{c}_j$ are unique. More precisely, if we call $\mathfrak{g}_j = \mathfrak{g}_j(A)$ the ideal generated by all the $(n+1-j) \times (n+1-j)$ minor-ideals in the last $n+1-j$ rows of the matrix A, then $\mathfrak{c}_j = \mathfrak{g}_{n+1-j} \mathfrak{g}_{n-j}^{-1}$.*

Proof. One easily checks that the ideals $\mathfrak{g}_m(A)$ are invariant under the elementary transformations of the type used in Algorithm 1.4.7. In particular, $\mathfrak{g}_j(A) = \mathfrak{g}_j(AU)$. But in the last $n+1-j$ rows of AU there is a single nonzero minor whose value is trivially equal to 1; hence we have $\mathfrak{g}_j(A) = \mathfrak{c}_{n+1-j} \cdots \mathfrak{c}_n$, proving the theorem. □

Proposition 1.4.10. *If AU is of the form given by Theorem 1.4.6, a necessary and sufficient condition for AV to be of the same form with the same ideals $\mathfrak{b}_j$ for $j > k-n$ is that $U^{-1}V$ be a block matrix $\left(\begin{smallmatrix} B & C \\ 0 & D \end{smallmatrix}\right)$ with D an $n \times n$ upper-triangular matrix with 1 on the diagonal such that for each i, j the entry in row i and column j belongs to $\mathfrak{c}_i \mathfrak{c}_j^{-1}$.*

Proof. Trivial and left to the reader. □

Corollary 1.4.11. *For each i and j with $1 \le i < j \le n$, let $S_{i,j}$ be a system of representatives of $K/\mathfrak{c}_i\mathfrak{c}_j^{-1}$. Write $AU = (0|H)$ as in Theorem 1.4.6. Then in that theorem, we may assume that for every i and j such that $i < j$ the entry in row i and column j of the matrix H is in $S_{i,j}$, in which case the matrix H is unique.*

Proof. For $i < j$, let $h_{i,j}$ be the entry in row i and column j of the matrix H. There exists a unique $h'_{i,j} \in S_{i,j}$ such that

$$q = h'_{i,j} - h_{i,j} \in \mathfrak{c}_i\mathfrak{c}_j^{-1} \ .$$

If the H_j are the columns of H, then by Proposition 1.4.10 the replacement of H_j by $H_j - qH_i$ is a legal elementary operation that transforms $h_{i,j}$ into $h'_{i,j}$, proving the existence. The uniqueness follows also from this, since there was a unique possible q. □

1.4.3 Reduction Modulo an Ideal

We can now comment on step 6 of Algorithm 1.4.7. By Corollary 1.4.11, the reduction done in step 6 is legal. Ideally, for each i and j, we would like to find a system of representatives of $K/\mathfrak{c}_i\mathfrak{c}_j^{-1}$ as well as an algorithm for finding the representative of a given element of K. There are at least two different methods for doing this, both of which have advantages and disadvantages.

The first method is to compute the (usual) HNF matrix H of $\mathfrak{c}_i\mathfrak{c}_j^{-1}$ on some fixed integral basis of K. If $(d_i)_{1\le i\le m}$ are the diagonal elements of H (with $m = [K : \mathbb{Q}]$), then we can take $S = \prod_{1\le i\le m} \mathbb{Q}/d_i\mathbb{Z}$ (and, for example, the interval $[0, d_i[$ as system of representatives of $\mathbb{Q}/d_i\mathbb{Z}$). If $x \in K$, we express x as a column vector (with rational entries) on the integral basis and then reduce x modulo $\mathfrak{c}_i\mathfrak{c}_j^{-1}$ from bottom up by subtracting from x suitable multiples of the columns of H so that the coordinates of x fall in the interval $[0, d_i[$ for each i.

We write this out explicitly as an algorithm.

Algorithm 1.4.12 (HNF Reduction Modulo an Ideal). Given an ideal $\mathfrak{a}$ by its $m \times m$ upper-triangular HNF matrix $H = (h_{i,j})$ in some basis of K, and an element $x \in K$ given by a column vector $X = (x_i)$ in the same basis, this algorithm computes a "canonical" representative of x modulo $\mathfrak{a}$, more precisely an element $y \in K$ such that $x - y \in \mathfrak{a}$ and the coordinates y_i of y in the basis satisfy $0 \le y_i < h_{i,i}$.

1. [Initialize] Set $i \leftarrow m$, $y \leftarrow x$.
2. [Reduce] Set $q \leftarrow \lfloor y_i/h_{i,i} \rfloor$, $y \leftarrow y - qH_i$ (recall that H_i is the ith column of H).
3. [Finished?] If $i = 1$, output y and terminate the algorithm; otherwise set $i \leftarrow i - 1$ and go to step 2.

This method has the advantage of giving a unique and well-defined representative of x modulo $\mathfrak{c}_i\mathfrak{c}_j^{-1}$ as well as an algorithm to find it. In practice, however, it often happens that the first few rows of the HNF matrix H are very large, and the others much smaller. Hence the resulting "reduced" element will in fact often be quite large.

The second method consists of first finding an LLL-reduced basis L of $\mathfrak{c}_i\mathfrak{c}_j^{-1}$, which will generally have much smaller entries than the HNF matrix H. We must then find an element $q \in \mathfrak{c}_i\mathfrak{c}_j^{-1}$ such that $x-q$ is small (we already mentioned the need for this in the remarks following Proposition 1.3.1). It is well known that this is a difficult problem (probably NP-complete). If, however, we write $x = \sum_{1\le j\le m} x_j L_j$ with $x_j \in \mathbb{Q}$ (where the L_j are the elements of the basis L) and choose

$$q = \sum_{1\le j\le m} \lfloor x_j \rceil L_j$$

(where $\lfloor a \rceil$ denotes one of the nearest integers to a), it is clear that $q \in \mathfrak{c}_i\mathfrak{c}_j^{-1}$ and that $x - q$ is reasonably "small". Note that it is essential that the basis L be LLL-reduced before doing this operation, otherwise $x - q$ would not be small at all in general.

We write this out explicitly as an algorithm.

Algorithm 1.4.13 (LLL-Reduction Modulo an Ideal). Given an ideal $\mathfrak{a}$ by an $m \times m$ matrix $H = (h_{i,j})$ representing a $\mathbb{Z}$-basis of $\mathfrak{a}$ in some basis of K, and an element $x \in K$ given by a column vector $X = (x_i)$ in the same basis, this algorithm computes a noncanonical but "small" representative of x modulo $\mathfrak{a}$, in other words an element $y \in K$ such that $x - y \in \mathfrak{a}$ and the coordinates y_i of y in the basis are reasonably small.

1. [LLL-reduce] Using the LLL algorithm or one of its variants, let L be the matrix of an LLL-reduced basis of $\mathfrak{a}$.
2. [Find coefficients] Using Gaussian elimination, find the solution $Z = (z_i)$ to the linear system $LZ = X$ (we have $Z = L^{-1}X$, but it is faster to compute Z directly than to invert L unless many elements are to be reduced modulo the same ideal).
3. [Reduce] Set $Y \leftarrow X - \sum_{1\le i\le m} \lfloor z_i \rceil L_i$, output the element y corresponding to Y, and terminate the algorithm.

The main advantage of this method is that the reduced vector will have much smaller entries. The reduction is not unique, however, and takes more time since the LLL algorithm is usually slower than the HNF algorithm, although it can of course be performed once and for all for a given ideal. Only practice can tell which method will be preferable. In the modular HNF method explained below, however, it is essential to use this method to avoid coefficient explosion.

The above algorithm can be improved by using an unpublished idea due to Peter Montgomery. Instead of doing an LLL-reduction of the ideal, which is an expensive operation, we can perform a fast *partial* reduction of the matrix (a matrix A with columns A_j will be said to be partially reduced if for any distinct columns we have $\|A_i \pm A_j\| \geq \|A_j\|$).

The resulting basis will usually not be LLL-reduced, but its entries will be of much smaller size than the Hermite-reduced one. Furthermore, this method is particularly well suited to matrices that have a few rows much larger than the others, such as typical HNF matrices for ideals.

1.5 Applications of the HNF Algorithm

1.5.1 Modifications to the HNF Pseudo-Basis

It is first necessary to make a number of remarks concerning the implementation of the HNF algorithm in Dedekind domains (Algorithm 1.4.7).

Usually a torsion-free R-module M will be given by a generating set expressed in a fixed basis $\mathcal{B}$ of KM. Using Algorithm 1.4.7, we can find a pseudo-basis $(\omega_j, \mathfrak{a}_j)_{1 \leq j \leq n}$ that has the special property of being upper-triangular with 1 on the diagonal when expressed on $\mathcal{B}$.

We can now start modifying this pseudo-basis. We can first choose to have only integral (and even primitive) ideals $\mathfrak{a}_j$ by dividing them by suitable elements of $\mathbb{Q}^*$ and multiplying the corresponding ω_j by the same. Alternatively, we can ask that $\omega_j \in M$, and this is done in a similar manner.

Then we can ask for a pseudo-basis such that all the ideals are equal to R except perhaps the last, whose ideal class will then be the Steinitz class of M. That this is possible follows from Proposition 1.2.19 together with Corollary 1.3.6. By induction, using legal elementary transformations on the matrix A, we can replace ideal pairs $(\mathfrak{a}_j, \mathfrak{a}_{j+1})$ by $(R, \mathfrak{a}_j\mathfrak{a}_{j+1})$, and hence at the end of the process all ideals except perhaps the last one will be equal to R, as desired. Note, however, that to apply Proposition 1.3.12 in a deterministic manner, it is necessary to know the prime decompositions of the norms of the $\mathfrak{a}_j$. In practice, this is always the case, but of course in general this is perhaps not a polynomial-time operation. Thus, in practice we use Algorithm 1.3.16, which is probabilistic but much faster.

Finally, note that if we perform the above transformations on the matrix and the ideals, the resulting pseudo-basis will no longer be represented by a triangular matrix.

If we are still not content with this, we could, if desired, obtain an $(n+1)$-element generating set of our module by replacing $\omega_n\mathfrak{a}_n$ with $a\omega_n + b\omega_n$, where $\mathfrak{a}_n = aR + bR$ is found using Algorithm 1.3.15. This will, of course, not be a direct sum. Note that the search for a and b can be done in deterministic polynomial time if the norm of $\mathfrak{a}_n$ is completely factored, since a can be taken equal to the norm of $\mathfrak{a}_n$.

We may also like to know if our module M is free and find a basis. Using the techniques developed in [Coh0, Chapter 6], once we find a relation matrix that is sufficient to compute the class group and regulator of R, it is quite easy to determine whether or not an ideal is principal and, if it is, to find a generator. Note that [Coh0] assumes the GRH, but evidently the same technique applies as long as we have obtained a relation matrix.

So we test whether $\mathfrak{a}_n$ is a principal ideal. If not, nothing more can be done: according to Theorem 1.2.25, M is not free, so use either the pseudo-basis (probably the best) or the $(n+1)$-element generating set. If $\mathfrak{a}_n = aR$, then after we replace ω_n by $a\omega_n$, $(\omega_j)_{1\le j\le n}$ is an R-basis of M.

If we want to know only whether or not M is free, without explicitly finding a basis, then it is not necessary to use Proposition 1.3.12 inductively: we use the initial HNF pseudo-basis and test whether or not $\mathfrak{a}_1 \dots \mathfrak{a}_n$ is a principal ideal.

1.5.2 Operations on Modules and Maps

As in the absolute case, the existence of an HNF algorithm (including an essential uniqueness statement) allows us to perform all of the standard operations on finitely generated, torsion-free modules. Let M and N be two such modules, assumed to be inside a larger module.

(1) To compute $M + N$, we simply compute the HNF of the concatenation of the HNF pseudo-bases of M and N.

(2) To check whether $M = N$, we simply check that the HNF of M and that of N are the same (this is of course the essential place where we need a unique HNF representative).

(3) To check whether $N \subset M$, we check that the HNF of $M + N$ and that of M are the same. Depending on the context, however, there may be faster methods.

(4) To compute the product MN when this makes sense, we form all the possible products of the generators and their corresponding ideals, and compute the HNF of the resulting pseudo-matrix. Usually, however, there are faster methods. For example, if M is an ideal given by a pseudo–two-element representation (see Definition 2.3.6) and if N is given in HNF, we must only multiply the generators and ideals of N by the two pseudo-elements of M.

(5) To compute the image and the kernel of a map f from N to M, we proceed as follows. Let $(\omega'_j, \mathfrak{a}'_j)_j$ and $(\omega_i, \mathfrak{a}_i)_i$ be pseudo-bases of N and M, respectively, and let A be the matrix of f in these pseudo-bases. Since $(f(\omega'_j), \mathfrak{a}'_j)_j$ is a pseudo-generating set for the image of f (note that the ideals $\mathfrak{a}'_j$ have not changed), we compute the HNF of $(A, (\mathfrak{a}'_j))$ using Algorithm 1.4.7 and thus obtain a pseudo-basis of the image of f. According to Theorem 1.4.6, the pseudo-matrix $(U_j, \mathfrak{b}_j)_{1\le j\le k-n}$ gives a

pseudo-basis of the kernel of f, where U is the transformation matrix given by Theorem 1.4.6.

(6) Computing the *intersection* $M \cap N$ of two modules is slightly more difficult. In [Coh0, Exercise 18 of Chapter 4], we have given a possible solution. However, the following algorithm is more elegant and useful also over $\mathbb{Z}$. (I thank D. Ford for having pointed it out to me, together with a reference to [Zim].)

Algorithm 1.5.1 (Intersection of Modules). Let M and N be two modules of the same rank n given by some pseudo-generating sets. This algorithm computes an HNF pseudo-basis for $M \cap N$.

1. [Compute pseudo-bases of M and N] Using Algorithm 1.4.7, compute the HNF (A, I) and (B, J) of the modules M and N, with $I = (\mathfrak{a}_i)$ and $J = (\mathfrak{b}_j)$ (only a pseudo-basis is necessary, not the HNF).
2. [Compute HNF of big matrix] Let C be the block matrix

$$C = \begin{pmatrix} A & 0 \\ A & B \end{pmatrix} .$$

   Using Algorithm 1.4.7, compute the HNF $(H, (\mathfrak{c}_j))$ of the pseudo-matrix $(C, (I|J))$.
3. [Terminate] Let H_1 be the upper-left $n \times n$ submatrix of H, and let J_1 be the first n ideals $(\mathfrak{c}_j)$. Output the pseudo-matrix (H_1, J_1) as representing a pseudo-basis of $M \cap N$ and terminate the algorithm.

Proof. The HNF reduction done in step 2 can be written in block matrix form,

$$\begin{pmatrix} A & 0 \\ A & B \end{pmatrix} \begin{pmatrix} U_1 & U_2 \\ U_3 & U_4 \end{pmatrix} = \begin{pmatrix} H_1 & H_2 \\ 0 & H_4 \end{pmatrix} ,$$

with the additional conditions $u_{i,j} \in \mathfrak{a}_i \mathfrak{c}_j^{-1}$ and $\mathfrak{a} = \det(U)\mathfrak{c}$, with notation similar to that of Theorem 1.4.6.

This implies in particular that $H_1 = AU_1$ and $AU_1 + BU_3 = 0$. Thus, if V_j is the jth column of U_1, we have $(\mathfrak{c}_j V_j)_i \subset \mathfrak{a}_i$, hence $\mathfrak{c}_j A V_j \subset M$ since $(A, (\mathfrak{a}_i))$ is a pseudo-matrix for M. It follows that the module defined by (H_1, J_1) is a submodule of M. But since $BU_3 = -AU_1$, the same reasoning on B and U_3 shows that it is also a submodule of N; hence it is a submodule of $M \cap N$.

Conversely, an element of $M \cap N$ can be represented as $AX = -BY$ for some vectors $X = (x_i)$ and $Y = (y_j)$ such that $x_i \in \mathfrak{a}_i$ and $y_j \in \mathfrak{b}_j$. It follows that the vector $\binom{X}{Y}$ belongs to the kernel of the pseudo-matrix $((A|B), (I|J))$, and Theorem 1.4.6 (5) tells us that this vector will be in the image of $(\binom{U_1}{U_3}, J_1)$. In particular, X will be in the image of (U_1, J_1); hence our initial element will be in the image of $(AU_1, J_1) = (H_1, J_1)$, as was to be proved. □

Remarks

(1) We have given the algorithm for two modules of the same rank n, because this is the only application that we have in mind (ideals in relative extensions), but the algorithm can easily be generalized to the case where the ranks are different (Exercise 12).

(2) In step 3, if we let H_4 be the lower-right $n \times n$ submatrix of H and J_2 the last n ideals $(\mathfrak{b}_j)$, the same proof shows that (H_4, J_2) is a pseudo-matrix giving a pseudo-basis of $M + N$.

1.5.3 Reduction Modulo $\mathfrak{p}$ of a Pseudo-Basis

Another natural question is the following. Assume that M is a finitely generated, torsion-free module, and that, thanks to the above algorithms, we have written $M = \bigoplus_i \mathfrak{a}_i \omega_i$ in terms of a pseudo-basis.

Let $\mathfrak{p}$ be a prime ideal of R. Then $M/\mathfrak{p}M$ is in a natural manner a *vector space* over the field $k = R/\mathfrak{p}$, and we can ask for a *basis* of this vector space over k. This can be done using the following algorithm.

Algorithm 1.5.2 (Reduction Modulo $\mathfrak{p}$ of a Pseudo-Basis). Let $(\omega_i, \mathfrak{a}_i)$ be a pseudo-basis for a finitely generated, torsion-free R-module M, and let $\mathfrak{p}$ be a (nonzero) prime ideal of R. This algorithm outputs a basis $(\overline{\eta_i})$ for $M/\mathfrak{p}M$ over the field $k = R/\mathfrak{p}$.

1. [Find two-element representation] For each i, use Algorithm 1.3.15 to find u_i and v_i in $\mathfrak{a}_i$ such that $\mathfrak{a}_i = u_i R + v_i R$ (one can, for example, choose $u_i = \mathcal{N}(\mathfrak{a}_i)$, but any other choice will do), and set $\alpha_i \leftarrow u_i$.
2. [Find α_i] For each i, using [Coh0, Algorithm 4.8.17], compute $v_{\mathfrak{p}}(u_i)$ and $v_{\mathfrak{p}}(v_i)$, and if $v_{\mathfrak{p}}(v_i) < v_{\mathfrak{p}}(u_i)$ set $\alpha_i \leftarrow v_i$.
3. [Terminate] For each i, let $\eta_i = \alpha_i \omega_i$. Output the $\overline{\eta_i}$ and terminate the algorithm.

Proof. If $\mathfrak{a}_i = u_i R + v_i R$, then $v_{\mathfrak{p}}(\mathfrak{a}_i) = \min(v_{\mathfrak{p}}(u_i), v_{\mathfrak{p}}(v_i))$; hence by our choice of α_i we have $v_{\mathfrak{p}}(\mathfrak{a}_i) = v_{\mathfrak{p}}(\alpha_i)$, so $\alpha_i \in \mathfrak{a}_i \smallsetminus \mathfrak{p}\mathfrak{a}_i$. This is easily seen to imply that the map from $R/\mathfrak{p}$ to $\mathfrak{a}_i/\mathfrak{p}\mathfrak{a}_i$ sending $\overline{x}$ to $\overline{x\alpha_i}$ is a group isomorphism; hence we have

$$M/\mathfrak{p}M = \bigoplus_i (\mathfrak{a}_i/\mathfrak{p}\mathfrak{a}_i)\overline{\omega_i} = \bigoplus_i (R/\mathfrak{p})\overline{\alpha_i \omega_i} \ ,$$

and thus the $(\overline{\alpha_i \omega_i})$ form a k-basis of $M/\mathfrak{p}M$. □

If $x \in M$, we will also want to compute the coefficients of $\overline{x}$ on the k-basis we have just computed. This is done as follows. By definition, we have $x = \sum_i a_i \omega_i$ with $a_i \in \mathfrak{a}_i$. Thus, with the notation of the algorithm, $x = \sum_i (a_i/\alpha_i)\eta_i$. Since $a_i \in \mathfrak{a}_i$,

$$v_{\mathfrak{p}}(a_i) \geq v_{\mathfrak{p}}(\mathfrak{a}_i) = v_{\mathfrak{p}}(\alpha_i) \ ;$$

in other words, $v_{\mathfrak{p}}(a_i/\alpha_i) \geq 0$. By an algorithmic version of the approximation theorem (for example, Proposition 1.3.11), we can find $y_i \in \mathbb{Z}_K$ such that $v_{\mathfrak{p}}(a_i/\alpha_i - y_i) \geq 1$. Hence in $M/\mathfrak{p}M$ we have $\overline{x} = \sum_i \overline{y_i}\overline{\eta_i}$, which is the desired decomposition. We will later see a more efficient algorithm for computing the y_i (Algorithm 4.2.22, see Exercise 13).

1.6 The Modular HNF Algorithm in Dedekind Domains

1.6.1 Introduction

It is well known that the usual HNF over $\mathbb{Z}$ suffers from coefficient explosion, which often makes the algorithm quite impractical, even for matrices of reasonable size. Since our algorithm is a direct generalization of the naive HNF algorithm, the same phenomenon occurs. Hence, it is necessary to improve the basic algorithm.

In the case of the ordinary HNF, there are essentially two ways of doing this, depending on what one wants.

The first method is the "modular" method. If we can compute the determinant of the lattice generated by the columns of our matrix, all computations can then be done modulo this determinant, and the final HNF matrix can be recovered by a simple GCD procedure (see [Coh0, Algorithm 2.4.6]). This method is polynomial-time, but it has the disadvantage of not computing directly the (unimodular) transformation matrix U. In most cases, this is not needed anyway, but in other cases it is essential (see, for example, the proof of Proposition 1.3.1). If we want the matrix U, it can be recovered from the modular method, but its entries will often be large and the method involves larger matrices (see Algorithm 1.6.3).

The other methods, due essentially to Havas (see [Hav-Maj1], [Hav-Maj2], and the references therein), are more heuristic in nature (they are not provably polynomial-time) but have the advantage of giving small transformation matrices U. Since in our applications to relative extensions of number fields we will often not need the matrix U, we will not consider here the generalization of Havas's algorithms to the Dedekind case, although they certainly can be generalized.

Hence, the purpose of this section is to explain how the usual modular HNF algorithm can be modified to work over Dedekind domains. Although quite simple, this generalization is not absolutely straightforward, so we give some details, closely following the exposition of [Coh0] and [Coh1].

1.6.2 The Modular HNF Algorithm

We have defined above the notion of a minor-ideal of a pseudo-matrix (A, I). In particular, $\mathfrak{g}_1(M)$ is the ideal of R generated by all $n \times n$ minor-ideals of

the pseudo-matrix (A, I). We will say that $\mathfrak{g}_1(M)$ is the *determinantal ideal* of the module M. It is clearly a generalization of the notion of determinant of a lattice.

Since there are $\binom{k}{n}$ minors of order n, it could be a lengthy task to compute $\mathfrak{g}_1(M)$ explicitly, except of course when $k = n$ or even $k = n+1$ (note that the computation of each minor is an ordinary determinant computation that can be done with the usual Gauss–Bareiss pivoting strategy, which only involves exact divisions).

We do not, however, really need the determinantal ideal itself but only an integral multiple of it. Furthermore, if we choose $n-1$ fixed independent columns, and consider the $k-n+1$ order-n minors obtained by choosing successively each of the remaining columns, we have a much more reasonable number of minor-ideals to compute, their computation is very fast (since $n-1$ of the pivoting steps are done once and for all), and the ideal sum of all these minor-ideals gives a reasonably sized multiple of the determinantal ideal $\mathfrak{g}_1(M)$.

Hence, we may assume that we have computed an ideal Δ that is an integral multiple (in other words, a subset) of the determinantal ideal $\mathfrak{g}_1(M)$ of M. We now describe what modifications must be made to Algorithm 1.4.7. We will make the computations in this algorithm modulo Δ, and then we will have to recover the correct HNF pseudo-matrix by suitable ideal operations.

First, we must compute modulo Δ. Recall that the individual columns A_j or ideals $\mathfrak{a}_j$ are quite arbitrary and that only the rank 1 submodule $\mathfrak{a}_j A_j$ of M is a reasonable object to consider. Hence, we must reduce modulo Δ not the column A_j itself, but the module $\mathfrak{a}_j A_j$. In other words, we must reduce the column A_j modulo the ideal $\Delta\mathfrak{a}_j^{-1}$.

Hence, we will modify step 5 of Algorithm 1.4.7 as follows. Before returning to step 4, we will set $A_m \leftarrow A_m \pmod{\Delta\mathfrak{a}_m^{-1}}$ and $A_j \leftarrow A_j \pmod{\Delta\mathfrak{a}_j^{-1}}$. Here, the reduction modulo an ideal is understood in the sense of the LLL-reduction Algorithm 1.4.13.

Since in the inner loop of Algorithm 1.4.7, the column index j is fixed and only m varies, it can also be argued that we should perform only the reduction of the column A_m, and perform the reduction of A_j only when the m-loop is finished. Although this avoids almost half of the (expensive) reductions, it may lead to much larger intermediate entries, so it is not clear if this method is preferable.

Once this modified algorithm is finished, we must execute the following supplementary algorithm to recover the true HNF pseudo-basis of M.

Algorithm 1.6.1 (Modular HNF Algorithm in Dedekind Domains). Given an $n \times k$ matrix $A = (a_{i,j})$ of rank n, and k (fractional) ideals $\mathfrak{a}_j$ in a number field K, this algorithm computes an HNF pseudo-basis (W, I) of the module $M = \sum_j \mathfrak{a}_j A_j$, where W is an $n \times n$ upper-triangular matrix with 1 on the

diagonal, and $I = (\mathfrak{b}_1, \ldots, \mathfrak{b}_n)$ is a list of n ideals. We assume that we have computed a multiple Δ of the determinantal ideal of M.

1. [Compute HNF modulo Δ] Using Algorithm 1.4.7, together with the modifications that we have just described for working modulo Δ, let $B = (b_{i,j})$ be the $n \times n$ HNF matrix obtained by discarding the first $k - n$ zero-columns from the resulting matrix AU, and let $\mathfrak{b}_j$ be the corresponding ideals (we discard in Algorithm 1.4.7 all the statements concerning the matrix U). Then set $\mathfrak{B} \leftarrow \Delta$, $i \leftarrow n$.
2. [Euclidean step] Set $\mathfrak{d} = b_{i,i}\mathfrak{b}_i + \mathfrak{B}$, and using Theorem 1.3.3, find $u \in \mathfrak{b}_i\mathfrak{d}^{-1}$ and $v \in \mathfrak{B}\mathfrak{d}^{-1}$ such that $b_{i,i}u + v = 1$. Then set $W_i \leftarrow uB_i \pmod{\mathfrak{B}\mathfrak{d}^{-1}}$ and $\mathfrak{b}_i \leftarrow \mathfrak{d}$ (here again reduction is done using Algorithm 1.4.13). Set $w_{i,i} \leftarrow 1$. (Note that $ub_{i,i} \equiv 1 \pmod{\mathfrak{B}\mathfrak{d}^{-1}}$, but the reduction modulo $\mathfrak{B}\mathfrak{d}^{-1}$ may not reduce it to 1.)
3. [Finished?] If $i > 1$, set $\mathfrak{B} \leftarrow \mathfrak{B}\mathfrak{d}^{-1}$ and go to step 2. Otherwise, for $i = n-1, n-2, \ldots, 1$, and for $j = i+1, \ldots, n$, using Algorithm 1.4.13, find $q \in \mathfrak{b}_i\mathfrak{b}_j^{-1}$ such that $w_{i,j} - q$ is small, and set $W_j \leftarrow W_j - qW_i$. Output the matrix W and the ideal list $I = (\mathfrak{b}_1, \ldots, \mathfrak{b}_n)$, and terminate the algorithm.

Proof. The proof of this algorithm's validity is essentially the same as in the classical case (see [Coh0, Algorithm 2.4.6] and [Coh1]); for brevity's sake we do not repeat it here. The $\mathfrak{g}_i(A)$, which are defined in the classical case as the GCD of all $i \times i$ minors extracted from the last i rows of A, are replaced in our situation by the minor-ideal $\mathfrak{g}_i(M)$, which plays exactly the same role (and reduces to the classical definition in the case where $\mathbb{Z}_K = \mathbb{Z}$). Note that, according to Proposition 1.3.4, for example (see also the remark after Corollary 1.3.6), the elementary column transformations made in step 3 are legal. □

As in the absolute case, it is more efficient in practice to interleave Algorithms 1.4.7 and 1.6.1 into a single algorithm, analogous to [Coh0, Algorithm 2.4.8]. The proof of this algorithm's validity follows from the proofs given above.

Algorithm 1.6.2 (Modular HNF Algorithm in Dedekind Domains). Given an $n \times k$ matrix $A = (a_{i,j})$ of rank n, and k (fractional) ideals $\mathfrak{a}_j$ in a number field K, this algorithm computes an HNF pseudo-basis (W, I) of the module $M = \sum_j \mathfrak{a}_j A_j$, where W is an $n \times n$ upper-triangular matrix with 1 on the diagonal, and $I = (\mathfrak{b}_1, \ldots, \mathfrak{b}_n)$ is a list of n ideals. We assume that we have computed a multiple Δ of the determinantal ideal of M.

1. [Initialize] Set $i \leftarrow n$, $j \leftarrow k$, and $\mathfrak{B} \leftarrow \Delta$.
2. [Check zero] Set $m \leftarrow j$, and while $m \geq 1$ and $a_{i,m} = 0$, set $m \leftarrow m - 1$. (Note that since we know that Δ is a nonzero ideal, it is not necessary to

check that the matrix A is of maximal rank.) If $m < j$, exchange A_m with A_j and $\mathfrak{a}_m$ with $\mathfrak{a}_j$.

3. [Put 1 on the main diagonal] Set $A_j \leftarrow A_j/a_{i,j}$, $\mathfrak{a}_j \leftarrow a_{i,j}\mathfrak{a}_j$, and $m \leftarrow j$. (We now have $a_{i,j} = 1$.)
4. [Loop] If $m = 1$, go to step 6. Otherwise, set $m \leftarrow m - 1$, and if $a_{i,m} = 0$, go to step 4.
5. [Euclidean step] (Here $a_{i,j} = 1$ and $a_{i,m} \neq 0$.) Using the algorithm contained in the proof of Theorem 1.3.3, set $\mathfrak{d} = a_{i,m}\mathfrak{a}_m + \mathfrak{a}_j$ and find $u \in \mathfrak{a}_m\mathfrak{d}^{-1}$ and $v \in \mathfrak{a}_j\mathfrak{d}^{-1}$ such that $a_{i,m}u + v = 1$. Then set in this order $(A_m, A_j) \leftarrow (A_m - a_{i,m}A_j, uA_m + vA_j)$, $(\mathfrak{a}_m, \mathfrak{a}_j) \leftarrow (\mathfrak{a}_m\mathfrak{a}_j\mathfrak{d}^{-1}, \mathfrak{d})$, $A_m \leftarrow A_m \pmod{\mathfrak{B}\mathfrak{a}_m^{-1}}$, and $A_j \leftarrow A_j \pmod{\mathfrak{B}\mathfrak{a}_j^{-1}}$, where the reduction is done using Algorithm 1.4.13. Finally, go to step 4.
6. [Next row] Set $\mathfrak{d} \leftarrow a_{i,j}\mathfrak{a}_j + \mathfrak{B}$ and using Theorem 1.3.3 once again compute $u \in \mathfrak{a}_j\mathfrak{d}^{-1}$ and $v \in \mathfrak{B}\mathfrak{d}^{-1}$ such that $ua_{i,j} + v = 1$. Set $W_i \leftarrow uA_j \pmod{\mathfrak{B}\mathfrak{d}^{-1}}$ (where the reduction is again done using Algorithm 1.4.13), $\mathfrak{a}_i \leftarrow \mathfrak{d}$, and $w_{i,i} \leftarrow 1$. For $m = j+1, \ldots, n$, using Algorithm 1.4.13 once more, find $q \in \mathfrak{a}_m\mathfrak{a}_j^{-1}$ such that $a_{i,m} - q$ is small, and set $A_m \leftarrow A_m - qA_j$.
7. [Finished?] If $i = 1$, output the matrix W and the modified ideals $\mathfrak{a}_j$, and terminate the algorithm. Otherwise, set $\mathfrak{B} \leftarrow \mathfrak{B}\mathfrak{d}^{-1}$, $i \leftarrow i - 1$, $j \leftarrow j - 1$ and go to step 2.

Remark. The above modular version performs well in practice, and it seems quite plausible that, as in the case of $R = \mathbb{Z}$, this algorithm is, in fact, polynomial-time.

1.6.3 Computing the Transformation Matrix

We finish this section by giving an algorithm that shows one method of recovering the unimodular transformation matrix by using the modular HNF algorithm (this algorithm is, of course, applicable also in the absolute case).

Algorithm 1.6.3 (Modular HNF with Transformation Matrix). Given an $n \times k$ matrix $A = (a_{i,j})$ of rank n, and k (fractional) ideals $\mathfrak{a}_j$ in a number field K, this algorithm computes a transformation matrix U, ideals $\mathfrak{b}_j$ for $1 \leq j \leq k$, an HNF pseudo-matrix (H, I) of the module $M = \sum_j \mathfrak{a}_j A_j$, where H is an $n \times n$ upper-triangular matrix with 1 on the diagonal, and $I = (\mathfrak{b}_{k-n+1}, \ldots, \mathfrak{b}_k)$ as in Theorem 1.4.6. We assume that we have computed a multiple Δ of the determinantal ideal of M.

1. [Find column permutation] Using a standard linear algebra algorithm over a field, find indices j_i for $1 \leq i \leq n$ such that the columns of index j_i of the matrix A are linearly independent. Let P be a permutation matrix sending these indices to the integers $[k - n + 1, k]$ so that the last n columns of the matrix AP are linearly independent.

2. [Apply modular HNF] Write $AP = (A_1|A_2)$ in block matrix form, where A_2 is an $n \times n$ matrix (which will be invertible, by step 1). Let C be the block matrix defined by $C = \begin{pmatrix} I_{k-n} & 0 \\ A_1 & A_2 \end{pmatrix}$. Let $(H, (\mathfrak{b}_j))$ with $H = \begin{pmatrix} H_1 & H_2 \\ 0 & H_4 \end{pmatrix}$ be the result of applying the modular HNF algorithm (Algorithm 1.6.2) to the pseudo-matrix $(C, (\mathfrak{a}_{p(j)}))$, where $p(j)$ is the permutation of the indices induced by the permutation matrix P.
3. [Terminate] Set $U_1 \leftarrow H_1$, $U_2 \leftarrow H_2$, $U_3 \leftarrow -A_2^{-1}A_1H_1$, $U_4 \leftarrow A_2^{-1}(H_4 - A_1H_2)$, $H \leftarrow H_4$, $I \leftarrow (\mathfrak{b}_{k-n+1}, \ldots, \mathfrak{b}_k)$, and $U \leftarrow P\begin{pmatrix} U_1 & U_2 \\ U_3 & U_4 \end{pmatrix}$. Output (H, I), the transformation matrix U, and the ideals $\mathfrak{b}_i$, and terminate the algorithm.

Proof. The proof is left to the reader (Exercise 14). □

1.7 The Smith Normal Form Algorithm in Dedekind Domains

Recall the elementary divisor theorem for torsion-free modules (Theorem 1.2.35).

Theorem. *Let P and N be two torsion-free modules of rank p and n, respectively, such that $N \subset P$ (so $n \leq p$). There exist fractional ideals $\mathfrak{b}_1, \ldots, \mathfrak{b}_p$ of R, a basis $(\omega_1, \ldots, \omega_p)$ of $V = PK$, and integral ideals $\mathfrak{d}_1, \ldots, \mathfrak{d}_n$ such that*

$$P = \mathfrak{b}_1\omega_1 \oplus \cdots \oplus \mathfrak{b}_p\omega_p \quad \text{and} \quad N = \mathfrak{d}_1\mathfrak{b}_1\omega_1 \oplus \cdots \oplus \mathfrak{d}_n\mathfrak{b}_n\omega_n$$

and such that $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$ for $2 \leq i \leq n$.

The ideals $\mathfrak{d}_i$ (for $1 \leq i \leq n$) and the ideal classes *of the ideal products $\mathfrak{b}_1 \cdots \mathfrak{b}_n$ and $\mathfrak{b}_{n+1} \cdots \mathfrak{b}_p$ depend only on P and N.*

In other words, this theorem says that we can find pseudo-bases of P and N that differ only in their ideals, in a specific way. Our main goal is to give an algorithm to find these pseudo-bases. This will be the Smith normal form algorithm (SNF).

Before doing this, we must generalize the notion of a pseudo-matrix. If (A, I) is a pseudo-matrix with $A = (a_{i,j})$ an $n \times k$ matrix with entries in K, and $I = (\mathfrak{a}_i)$ a vector of k ideals, it is natural to consider the linear map f from $\mathfrak{a}_1 \times \cdots \times \mathfrak{a}_k$ to K^n, defined by

$$f(a_1, \ldots, a_k) = \sum_{1 \leq j \leq k} a_j A_j \ ,$$

where as usual A_j denotes the jth column of A, considered as an element of K^n. The image of this map f is exactly the module $M = \sum_j \mathfrak{a}_j A_j$ with which we have worked.

We must now consider the more general situation where the map f is a linear map from $N = \mathfrak{a}_1 \times \cdots \times \mathfrak{a}_n$ to $P = \mathfrak{b}_1 \times \cdots \times \mathfrak{b}_p$ for some other ideals $\mathfrak{b}_i$. If we call i_j the jth canonical injection from $\mathfrak{a}_j$ to N (defined by $i_j(a) = (0, \ldots, 0, a, 0, \ldots, 0)$, where a is at the jth component) and p_i the ith canonical projection from P to $\mathfrak{b}_i$ (defined by $p_i(b_1, \ldots, b_n) = b_i$), we will set

$$f_{i,j} = p_i \circ f \circ i_j \ .$$

This is a linear map from $\mathfrak{a}_j$ to $\mathfrak{b}_i$. Conversely, given any family of linear maps $g_{i,j}$ from $\mathfrak{a}_j$ to $\mathfrak{b}_i$, we can define in a unique manner a linear map f from N to P such that $f_{i,j} = g_{i,j}$.

Let $\mathfrak{a}$ and $\mathfrak{b}$ be two ideals and g an R-linear map from $\mathfrak{a}$ to $\mathfrak{b}$. By tensoring with the field K we can extend this to a K-linear map from K to K (which we denote again by g); such a map is of the form $g(x) = \lambda x$ for some $\lambda \in K$. Conversely, such a λ gives a map from $\mathfrak{a}$ to $\mathfrak{b}$ if and only if $\lambda\mathfrak{a} \subset \mathfrak{b}$, hence if and only if $\lambda \in \mathfrak{b}\mathfrak{a}^{-1}$. This leads us to the following definition.

Definition and Proposition 1.7.1. *Let $N = \mathfrak{a}_1\omega_1 \oplus \cdots \oplus \mathfrak{a}_n\omega_n$ and $P = \mathfrak{b}_1\eta_1 \oplus \cdots \oplus \mathfrak{b}_p\eta_p$ be two torsion-free R-modules given by pseudo-bases, and let $A = (a_{i,j})$ be a $p \times n$ matrix. Let $I = (\mathfrak{b}_1, \ldots, \mathfrak{b}_p)$ and $J = (\mathfrak{a}_1, \ldots, \mathfrak{a}_n)$.*

(1) *We will say that (A, I, J) is an* integral *pseudo-matrix if for each i and j we have $a_{i,j} \in \mathfrak{b}_i\mathfrak{a}_j^{-1}$.*

(2) *Given such a pseudo-matrix (A, I, J), the map f from N to P associated with it is the map defined by setting*

$$f\left(\sum_j x_j\omega_j\right) = \sum_j x_j f(\omega_j) = \sum_j x_j \sum_i a_{i,j}\eta_i = \sum_i \eta_i \left(\sum_j a_{i,j}x_j\right) \ ,$$

which makes sense since $a_{i,j}x_j \in \mathfrak{b}_i$.

(3) *The module M associated with (A, I, J) is the quotient module*

$$P/f(N) = (\mathfrak{b}_1\eta_1 \oplus \cdots \oplus \mathfrak{b}_p\eta_p)/f(\mathfrak{a}_1\omega_1 \oplus \cdots \oplus \mathfrak{a}_n\omega_n) \ .$$

Note that the module M associated with a pseudo-matrix (A, I, J) is a torsion module if and only if $p = n$, that is, if A is a square matrix of nonzero determinant.

We can now state the main theorem of this section. For simplicity we state it for square matrices, but it is easily extended to the general case.

Theorem 1.7.2 (Smith Normal Form in Dedekind Domains). *Let (A, I, J) be an integral pseudo-matrix as above, with $A = (a_{i,j})$ an $n \times n$ matrix and $I = (\mathfrak{b}_1, \ldots, \mathfrak{b}_n)$, and $J = (\mathfrak{a}_1, \ldots, \mathfrak{a}_n)$ two vectors of n ideals such that $a_{i,j} \in \mathfrak{b}_i\mathfrak{a}_j^{-1}$.*

There exist vectors of ideals $(\mathfrak{b}'_1, \ldots, \mathfrak{b}'_n)$ and $(\mathfrak{a}'_1, \ldots, \mathfrak{a}'_n)$ and two $n \times n$ matrices $U = (u_{i,j})$ and $V = (v_{i,j})$ satisfying the following conditions, where for all i we set $\mathfrak{d}_i = \mathfrak{a}'_i\mathfrak{b}'^{-1}_i$, and we set $\mathfrak{a} = \mathfrak{a}_1 \cdots \mathfrak{a}_n$, $\mathfrak{b} = \mathfrak{b}_1 \cdots \mathfrak{b}_n$, $\mathfrak{a}' = \mathfrak{a}'_1 \cdots \mathfrak{a}'_n$, and $\mathfrak{b}' = \mathfrak{b}'_1 \cdots \mathfrak{b}'_n$.

(1) $\mathfrak{a} = \det(U)\mathfrak{a}'$ *and* $\mathfrak{b}' = \det(V)\mathfrak{b}$ (*note the reversal*).
(2) *The matrix* VAU *is the* $n \times n$ *identity matrix.*
(3) *The* $\mathfrak{d}_i$ *are integral ideals, and for* $2 \leq i \leq n$ *we have* $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$.
(4) *For all* i, j *we have* $u_{i,j} \in \mathfrak{a}_i \mathfrak{a}_j'^{-1}$ *and* $v_{i,j} \in \mathfrak{b}_i' \mathfrak{b}_j^{-1}$.

Proof. Again we prove this theorem by giving an explicit algorithm for constructing the Smith normal form. We follow closely [Coh0, Algorithm 2.4.14], except that we do not work modulo the determinant (although such a modular version of the Smith normal form algorithm is easily written).

Algorithm 1.7.3 (SNF Algorithm in Dedekind Domains). Given an invertible $n \times n$ matrix $A = (a_{i,j})$, and two lists of n (fractional) ideals $I = (\mathfrak{b}_i)$ and $J = (\mathfrak{a}_j)$ in a number field K, this algorithm computes two other lists of n ideals $\mathfrak{b}_i'$ and $\mathfrak{a}_j'$ and two $n \times n$ matrices U and V such that these data satisfy the conditions of Theorem 1.7.2. We will only make use of elementary transformations of the type given in Theorem 1.3.3 combined with Corollary 1.3.5. We denote by A_j (resp., U_j) the columns of A (resp., U), and by A_j' (resp., V_j') the rows of A (resp., V).

1. [Initialize i] Set $i \leftarrow n$, and let U and V be the $n \times n$ identity matrix. If $n = 1$, output $\mathfrak{b}_1$, $\mathfrak{a}_1$, U, and V, and terminate the algorithm.
2. [Initialize j for row reduction] Set $j \leftarrow i$ and $c \leftarrow 0$.
3. [Check zero] If $j = 1$, go to step 5. Otherwise, set $j \leftarrow j - 1$. If $a_{i,j} = 0$, go to step 3.
4. [Euclidean step] Using the algorithm of Theorem 1.3.3, set $\mathfrak{d} \leftarrow a_{i,i}\mathfrak{a}_i + a_{i,j}\mathfrak{a}_j$ and find $u \in \mathfrak{a}_i\mathfrak{d}^{-1}$ and $v \in \mathfrak{a}_j\mathfrak{d}^{-1}$ such that $a_{i,i}u + a_{i,j}v = 1$. Then set $(A_j, A_i) \leftarrow (a_{i,j}A_j - a_{i,i}A_i, uA_i + vA_j)$, $(U_j, U_i) \leftarrow (a_{i,j}U_j - a_{i,i}U_i, uU_i + vU_j)$, $(\mathfrak{a}_j, \mathfrak{a}_i) \leftarrow (\mathfrak{a}_i\mathfrak{a}_j\mathfrak{d}^{-1}, \mathfrak{d})$. Finally, go to step 3.
5. [Initialize j for column reduction] Set $j \leftarrow i$, and if $a_{i,i} \neq 1$, set $U_i \leftarrow U_i/a_{i,i}$, $\mathfrak{a}_i \leftarrow a_{i,i}\mathfrak{a}_i$, $a_{i,i} \leftarrow 1$.
6. [Check zero] If $j = 1$, go to step 8. Otherwise, set $j \leftarrow j - 1$. If $a_{j,i} = 0$, go to step 6.
7. [Euclidean step] Using the algorithm of Theorem 1.3.3, set $\mathfrak{d} \leftarrow \mathfrak{b}_i^{-1} + a_{j,i}\mathfrak{b}_j^{-1}$ and find $u \in \mathfrak{b}_i^{-1}\mathfrak{d}^{-1}$ and $v \in \mathfrak{b}_j^{-1}\mathfrak{d}^{-1}$ such that $u + a_{j,i}v = 1$. Then set $(A_j', A_i') \leftarrow (a_{j,i}A_j' - A_i', uA_i' + vA_j')$, $(V_j', V_i') \leftarrow (a_{j,i}V_j' - V_i', uV_i' + vV_j')$, $(\mathfrak{b}_j, \mathfrak{b}_i) \leftarrow (\mathfrak{b}_i\mathfrak{b}_j\mathfrak{d}, \mathfrak{d}^{-1})$. Finally, set $c \leftarrow c + 1$ and go to step 6.
8. [Repeat stage i?] If $c > 0$, go to step 2.
9. [Check the rest of the matrix] Set $\mathfrak{b} \leftarrow \mathfrak{a}_i\mathfrak{b}_i^{-1}$. For $1 \leq k, l < i$, check whether $a_{k,l}\mathfrak{a}_l\mathfrak{b}_k^{-1} \subset \mathfrak{b}$. As soon as this is not the case, set $\mathfrak{d} \leftarrow \mathfrak{b}_i\mathfrak{b}_k^{-1}$. Let d be an element of $\mathfrak{d}$ such that $a_{k,l}d \notin \mathfrak{a}_i\mathfrak{a}_l^{-1}$ (such an element must exist and is easy to find — for example, by looking at the $\mathbb{Z}$-basis of $\mathfrak{d}$ given by the ordinary HNF). Set $A_i' \leftarrow A_i' + dA_k'$ and $V_i' \leftarrow V_i' + dV_k'$, and go to step 2.
10. [Next stage] (Here $a_{k,l}\mathfrak{a}_l\mathfrak{b}_k^{-1} \subset \mathfrak{b}$ for all $k, l < i$.) If $i \geq 3$, set $i \leftarrow i - 1$ and go to step 2. Otherwise, set $U_1 \leftarrow U_1/a_{1,1}$, $\mathfrak{a}_1 \leftarrow a_{1,1}\mathfrak{a}_1$, and $a_{1,1} \leftarrow 1$,

output the matrices U and V, the two ideal lists $(\mathfrak{b}_i)$ and $(\mathfrak{a}_j)$, and terminate the algorithm.

Proof. Contrary to the HNF algorithm whose proof was immediate, there are several things to be checked. First we must check that this algorithm is valid. It is easily verified that all the elementary operations used are legal ones and that the identities $\mathfrak{a} = \det(U)\mathfrak{a}'$ and $\mathfrak{b}' = \det(V)\mathfrak{b}$ are preserved throughout. Furthermore, upon termination the matrix A will be the identity matrix and we will have $a_{j,j}\mathfrak{b}_j'^{-1}\mathfrak{a}_j' \subset a_{i,i}\mathfrak{b}_i'^{-1}\mathfrak{a}_i'$ for all $j < i$. Hence, since $a_{i,i} = a_{j,j} = 1$, we obtain from the definition of the $\mathfrak{d}_i$ that $\mathfrak{d}_j \subset \mathfrak{d}_i$ for all $j < i$. In addition, it is easily checked that the ideal $\mathfrak{c} = \sum_{i,j} a_{i,j}\mathfrak{a}_j\mathfrak{b}_i^{-1}$ is preserved by all the elementary transformations of rows and columns that we perform. Since we have assumed that $a_{i,j} \in \mathfrak{b}_i\mathfrak{a}_j^{-1}$, it follows that $\mathfrak{c}$ is an integral ideal. But in the final pseudo-matrix we have $\mathfrak{c} = \sum_i \mathfrak{a}_i'\mathfrak{b}_i' = \sum_i \mathfrak{d}_i = \mathfrak{d}_n$, since $\mathfrak{d}_j \subset \mathfrak{d}_i$ for $j < i$. Thus $\mathfrak{d}_n$ is an integral ideal; hence all the $\mathfrak{d}_i$ are integral ideals.

Note that we could interpret all the $\mathfrak{d}_i$ in the same way by taking the sum of *all* $(n-i) \times (n-i)$ minor-ideals of the pseudo-matrix, where the minor-ideals of an integral pseudo-matrix (A, I, J) are defined as for a pseudo-matrix (A, I) (Definition 1.4.8), except that we must multiply the determinant d by the product of the ideals $\mathfrak{a}_j\mathfrak{b}_i^{-1}$ and not only by the product of the ideals $\mathfrak{a}_j$.

We must now show that the algorithm terminates. First note that the effect of steps 2 to 8 on the triplet $(a_{i,i}, \mathfrak{a}_i, \mathfrak{b}_i^{-1})$ is to transform it into

$$\left(1, \sum_{j \le i} a_{i,j}\mathfrak{a}_j, \mathfrak{b}_i^{-1} + \sum_{j<i} a'_{j,i}\mathfrak{b}_j^{-1}\right) ,$$

where the $a'_{j,i}$ are the entries of the matrix after step 4. Hence, the product $a_{i,i}\mathfrak{a}_i\mathfrak{b}_i^{-1}$, which is an integral ideal throughout the algorithm (since it is included in the ideal $\mathfrak{c} = \mathfrak{d}_n$), can only get larger. Since all the ideals are nonzero, steps 2 to 8 can leave this product unchanged only if $a_{i,j} = 0$ and $a'_{j,i} = 0$ for all $j < i$, and this implies that $c = 0$, which is the termination condition of the loop from steps 2 to 8. Thus, we have a strictly increasing sequence of integral ideals, which is therefore finite. Hence we reach step 9 after a finite number of steps.

One loop from step 9 back to step 5 again transforms the triplet $(1, \mathfrak{a}_i, \mathfrak{b}_i^{-1})$ into

$$\left(1, \mathfrak{a}_i + \sum_{j<i} da_{k,j}\mathfrak{a}_j, \mathfrak{b}_i^{-1}\right) .$$

Hence, since $da_{k,l} \notin \mathfrak{a}_i\mathfrak{a}_l^{-1}$, it follows that the new ideal $\mathfrak{a}_i$ is strictly larger, and hence the new $a_{i,i}\mathfrak{a}_i\mathfrak{b}_i^{-1}$ also. We again have a strictly increasing sequence of integral ideals, which is therefore finite; hence we execute step 9 only a finite number of times, and so the algorithm terminates. □

Remarks

(1) Considering step 7 of the algorithm, in practice it will probably be better to keep the ideals $\mathfrak{b}_i^{-1}$ and not the ideals $\mathfrak{b}_i$ themselves, so as to diminish the number of ideal inversions.

(2) As mentioned earlier, it is very easy to introduce a modular version of the SNF algorithm, as in [Coh0, Algorithm 2.4.14]. Such a variant is necessary in certain cases to avoid coefficient explosion. In addition, the algorithm is easily modified to deal with singular or nonsquare matrices.

(3) Note that the module M associated with the pseudo-matrix (A, I, J) will be isomorphic to

$$R/\mathfrak{d}_1 \oplus \cdots \oplus R/\mathfrak{d}_n \ ,$$

and thus this gives the complete structure of M as an R-module.

1.8 Exercises for Chapter 1

1. Let K be a number field, and let $x \in K^*$. With the notation of Proposition 1.2.7, show that

$$\prod_{1 \le i \le r_1+r_2} |x|_{\sigma_i}^{n_i} = 1/|\mathcal{N}(x)| \quad \text{and} \quad \prod_{\mathfrak{p}} |x|_{\mathfrak{p}} = |\mathcal{N}(x)| \ ,$$

thus proving Proposition 1.2.7.

2. Give a counterexample to Proposition 1.2.8 if one asks that $|x|_i \le 1$ for all places $|\ |_i \notin S$.

3. With the notation of Definition 1.2.14, show that the relation

$$(a_1/a_2)\alpha \mathcal{R} (b_1/b_2)\beta \iff b_2 a_1 \alpha - a_2 b_1 \beta = 0$$

is not in general an equivalence relation.

4. Someone remembered the definition of a projective module P as a module such that for every surjective map f from a module F to P and any map g from a module G to P, there exists a map h from G to F such that $g = f \circ h$ (see the diagrams below, where the first diagram illustrates Definition 1.2.16 and the second the present definition). Is this definition correct? If yes, prove it; otherwise, give a counterexample.

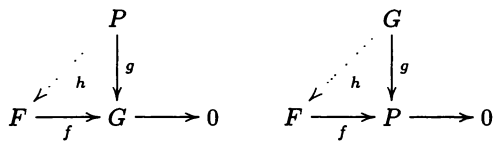

5. Let R be a Dedekind domain, let $\mathfrak{a}$ and $\mathfrak{b}$ be two fractional ideals of R, and let a, b, c, and d be elements of R.

   a) Show that a necessary and sufficient condition for the existence of ideals $\mathfrak{c}$ and $\mathfrak{d}$ satisfying Proposition 1.3.4 is that $ad-bc \neq 0$ and $(a\mathfrak{b}+b\mathfrak{a})(c\mathfrak{b}+d\mathfrak{a}) = (ad-bc)\mathfrak{a}\mathfrak{b}$.

b) If this condition is satisfied, show that

$$\mathfrak{c} = \frac{cb + da}{ad - bc} = \frac{\mathfrak{a}}{a} \cap \frac{\mathfrak{b}}{b} = (ab + ba)^{-1}\mathfrak{a}\mathfrak{b} \quad \text{and}$$
$$\mathfrak{d} = \frac{ab + ba}{ad - bc} = \frac{\mathfrak{a}}{c} \cap \frac{\mathfrak{b}}{d} = (cb + da)^{-1}\mathfrak{a}\mathfrak{b} \ .$$

6. Let $\mathfrak{a}$ be a nonzero fractional ideal of a Dedekind domain R. Show that there is a canonical isomorphism between $\mathfrak{a}^{-1}$ and the R-module of R-linear maps from $\mathfrak{a}$ to R.
7. Using the Hermite and Smith normal form algorithms, give another proof of the structure theorem for finitely generated modules over Dedekind domains.
8. Let R be a Dedekind domain with field of fractions K.
   a) For each α, β in K, show that there exist u and v in R such that
   $$u\alpha^2 + v\beta^2 = \alpha\beta \ .$$
   b) Using Algorithm 1.3.2, give an algorithm to compute u and v when $R = \mathbb{Z}_K$ is the ring of integers of a number field.
   c) Show that the result of a) can be false if R is not a Dedekind domain (take, for example, $R = \mathbb{Z}[X]$ or $R = \mathbb{Z}[\sqrt{8}]$).
   d) Show that the result of a) may be true even if R is not a Dedekind domain (take, for example, $R = \mathbb{Z}[\sqrt{5}]$).
   e) More generally, if $R = \mathbb{Z}[\sqrt{D}]$ with D a nonsquare integer, show that the result of a) is valid if and only if D is squarefree (which is *not* the same condition as saying that R is a Dedekind domain).
9. Let R be a Dedekind domain and $\mathfrak{a}$ a fractional ideal of R. Set $\mathfrak{n} = \mathfrak{a} \cap R$ and $\mathfrak{d} = \mathfrak{a}^{-1} \cap R$. Show that $\mathfrak{n}$ and $\mathfrak{d}$ are coprime integral ideals such that $\mathfrak{a} = \mathfrak{n}\mathfrak{d}^{-1}$.
10. Modify Algorithm 1.3.14 so that it instead computes a $\beta \in K$ such that $\beta\mathfrak{a}^{-1}$ is an integral ideal coprime to $\mathfrak{b}$.
11. Prove Proposition 1.4.4.
12. Write a generalization of Algorithm 1.5.1 to the case where the modules have different ranks.
13. Using the ideas of Algorithm 4.2.22, write an efficient algorithm to compute the coefficients y_i such that $\overline{x} = \sum_i \overline{y_i}\overline{\eta_i}$, with the notation of Section 1.5.3.
14. Prove the validity of Algorithm 1.6.3.

2. Basic Relative Number Field Algorithms

Having the necessary tools for dealing theoretically and algorithmically with modules over Dedekind domains, we are now going to study in detail relative extensions of number fields. In the first section, we emphasize the *field-theoretic* properties, while in the rest of this chapter we study the *ring-theoretic* properties.

2.1 Compositum of Number Fields and Relative and Absolute Equations

2.1.1 Introduction

A number field L can be represented in many different ways, all having their relative advantages and disadvantages. In [Coh0] we consider number fields L to be given as $L = \mathbb{Q}(\theta)$, where θ is a root of some polynomial $T \in \mathbb{Q}[X]$. A number field L is thus explicitly considered as a finite extension of $\mathbb{Q}$, in other words as an *absolute extension*, with $T(X)$ being an *absolute defining polynomial* or an *absolute equation* for the number field L.

In most problems of algebraic number theory, it is often more natural to consider *relative* extensions L/K. In other words, we have a *base field* K, and a number field L, which is given as $L = K(\theta)$, where θ is a root of some polynomial $T \in K[X]$, called a *relative defining polynomial* or a *relative equation* for L over K. Of course, L is still a number field in its own right and as such can also be given by an absolute defining polynomial if so desired (we will see in Section 2.1.5 how to achieve this), but it is usually preferable to keep the field L given by its relative defining polynomial. There are at least two reasons for this. First, the relative defining polynomial will be considerably simpler than the absolute one (for instance, it will be of much lower degree). Second, the K-structure on L gives considerably more arithmetical information than considering L on its own.

If $K = \mathbb{Q}(\theta_1)$ and $L = K(\theta_2)$, we can write $L = \mathbb{Q}(\theta_1, \theta_2)$, in this case considered as a tower of extensions $L/K/\mathbb{Q}$. In the special case where the minimal monic polynomial of θ_2 in $K[X]$ has in fact coefficients in $\mathbb{Q}$, we can consider the number field $K_2 = \mathbb{Q}(\theta_2)$, and in this case L is a *compositum* of the number fields K and K_2.

There are still other ways to represent a number field L. For example, we can choose a $\mathbb{Q}$-basis of L (or a K-basis, if we view L as a relative extension), and represent L by the multiplication table of this basis. If $n = [L : \mathbb{Q}]$ (or $n = [L : K]$ in the relative case), this is an $n \times n \times n$ array with entries in $\mathbb{Q}$ (or K). Although computationally more cumbersome and expensive, this representation can sometimes also be useful, but we will not study it here.

A completely different way is to represent a number field by a *system* of polynomials in *several* indeterminates, assuming that the solution is zero-dimensional. This touches upon difficult problems of computational algebraic geometry, in particular Gröbner basis computations, and we refer to the abundant literature on the subject (see, for example, [Co-Li-Os] and [Bec-Wei]).

In this chapter, we fix a number field K and identify it with a subfield of $\mathbb{C}$. We will only consider number fields L that are extensions of K and contained in the algebraic closure $\overline{K}$ of K in $\mathbb{C}$. In particular, L will be a K-vector space, and one of the embeddings of K (and L) in $\mathbb{C}$ is the identity. We will say that K is the *base field*, and L is a *relative extension*, or a *number field over* K. If L is a relative extension of a number field K, the dimension $n = \dim_K(L)$ will be denoted $[L : K]$ and called the *relative degree* of L over K. If M is a relative extension of L, we clearly have the transitivity relation $[M : K] = [M : L][L : K]$.

2.1.2 Étale Algebras

Let $L = K(\theta)$ be a relative extension of number fields, where θ is the root of an irreducible polynomial $T(X) \in K[X]$. As already remarked in [Coh0], we have an isomorphism of L with $K[X]/T(X)K[X]$ obtained by sending θ to the class of X. A natural question is to consider the case where T is not irreducible. If the factorization of T in $K[X]$ is given by $T(X) = \prod_i T_i(X)^{e_i}$, where the T_i are nonassociate, irreducible polynomials in $K[X]$, the Chinese remainder theorem tells us that

$$K[X]/T(X)K[X] \simeq \prod_i K[X]/T_i^{e_i}(X)K[X]$$

(see Exercise 1). If $e_i = 1$, $K[X]/T_i(X)K[X]$ is a number field; while if $e_i > 1$, $K[X]/T_i^{e_i}(X)K[X]$ has nonzero nilpotent elements and hence is an inseparable algebra over K. These algebras have nasty properties, and we want to exclude them. We have the following proposition.

Proposition 2.1.1. *Let K be a number field, and let A be a finite-dimensional commutative K-algebra (in other words, a finite-dimensional K-vector space with an additional commutative ring structure with unit, compatible with the vector space structure). The following three properties are equivalent.*

(1) *A has no nonzero nilpotent elements;*
(2) *The equation $x^2 = 0$ in A implies $x = 0$;*

(3) *The minimal polynomial in $K[X]$ of any element $a \in A$ is squarefree.*

Proof. That (1) implies (2) is trivial. Let us prove that (2) implies (3), so assume (2), and let $a \in A$. The set I_a of polynomials $P \in K[X]$ such that $P(a) = 0$ is clearly an ideal of $K[X]$. Furthermore, since A is of finite dimension n, say, the elements $1, a, \ldots, a^n$ are K-linearly dependent; hence I_a is nonzero. Therefore, I_a is generated by a monic polynomial $P_a \in K[X]$, which will be called as usual the *minimal polynomial* of a in A. Assume that $P_a(X) = Q^2(X)R(X)$ in $K[X]$, and let $b = Q(a)R(a)$. We have $b^2 = P_a(a)R(a) = 0$; hence by (2), $b = 0$. But this means that $Q(X)R(X)$ is a multiple of the minimal polynomial $Q^2(X)R(X)$. If follows that $Q(X)$ is constant, so P_a is squarefree, as claimed.

Finally, if $a^n = 0$, the minimal polynomial of a must be a divisor of X^n, and it must be squarefree by (3), so it must be equal to X. Hence $a = 0$ and so (3) implies (1). □

Definition 2.1.2. *Let K be a number field, and let A be a commutative K-algebra. We say that A is an* étale algebra (*or a* separable algebra) *over K if A is finite-dimensional over K and satisfies the equivalent properties of Proposition 2.1.1.*

In particular, a number field L that is an extension of K is trivially an étale algebra. We will see in Corollary 2.1.6 that in fact every étale algebra is isomorphic to a product of number fields.

Note that the minimal polynomial of an element of an étale algebra is squarefree but not necessarily irreducible. In fact, this is another way of saying that the difference between étale algebras and number fields is the existence of zero divisors.

Before stating and proving the primitive element theorem, which is one of the main results about étale algebras, we prove the following apparently unrelated proposition (see [Coh0, Exercise 4 of Chapter 3]).

Proposition 2.1.3. *Let B be a commutative ring with unit, and let T_1 and T_2 be polynomials in $B[X]$. There exist polynomials $U_1(X)$ and $U_2(X)$ in $B[X]$ such that*

$$U_1(X)T_1(X) + U_2(X)T_2(X) = \mathcal{R}(T_1(X), T_2(X)) \in B \ ,$$

where as usual $\mathcal{R}(T_1(X), T_2(X))$ denotes the resultant of the polynomials $T_1(X)$ and $T_2(X)$.

Proof. Let M be the Sylvester matrix associated to the polynomials T_1 and T_2 (see [Coh0, Lemma 3.3.4]). If $\deg(T_i) = n_i$, let $U_1(X) = \sum_{0 \le i < n_2} x_i X^i$ and $U_2(X) = \sum_{0 \le i < n_1} y_i X^i$ be arbitrary polynomials of degree less than or equal to $n_2 - 1$ and $n_1 - 1$, respectively, and let

$$Z = (x_{n_2-1}, \ldots, x_0, y_{n_1-1}, \ldots, y_0) \quad \text{and} \quad \mathcal{X} = (X^{n_1+n_2-1}, \ldots, X, 1)^t \ .$$

Then the matrix M can be defined by the equation

$$ZM\mathcal{X} = U_1(X)T_1(X) + U_2(X)T_2(X) \ .$$

Let M^{adj} be the adjoint matrix of M. By definition, we have

$$M^{\mathrm{adj}}M = \det(M)I_{n_1+n_2} = \mathcal{R}(T_1(X), T_2(X))I_{n_1+n_2} \ .$$

Hence, if Z is the last row of the matrix M^{adj}, we have

$$ZM\mathcal{X} = \mathcal{R}(T_1(X), T_2(X)) = U_1(X)T_1(X) + U_2(X)T_2(X)$$

if $U_1(X)$ and $U_2(X)$ are related to Z as above. Since the entries of the adjoint matrix, hence of Z, are in the ring B, this proves the existence of the polynomials $U_1(X)$ and $U_2(X)$ in $B[X]$ and also gives an algorithm to find them. The algorithm mentioned in [Coh0, Exercise 5 of Chapter 3], based on the subresultant algorithm is, however, much better in practice. □

Remark. The point of this proposition is that the polynomials U_1 and U_2 can be chosen in $B[X]$, and not in a larger ring.

We can now prove a special case of the primitive element theorem.

Lemma 2.1.4. *Let $A = K[\theta_1, \theta_2]$ be an étale algebra generated by two elements. There exists $\theta \in A$ such that $A = K[\theta]$, and θ can be taken of the form $\theta = \theta_2 + k\theta_1$ for some $k \in \mathbb{Z}$.*

Remark. Note the use of square brackets in the expression $A = K[\theta]$. Since A is not a field in general, we consider only polynomials in θ, and not rational functions. On the other hand, when A is a field, it is easy to see that $K(\theta) = K[\theta]$, and this is the notation most commonly used in this case, to emphasize that A is a field.

Proof. Let T_1 and T_2 be the minimal monic polynomials of θ_1 and θ_2, respectively. Since A is an étale algebra, these polynomials are squarefree, hence have distinct roots in some fixed algebraic closure $\overline{K}$ of K.

Let $(\theta_1^{(i)})$ (resp., $(\theta_2^{(j)})$) be the roots of T_1 (resp., T_2) in $\overline{K}$. Since in general A is not a field, we cannot assume that $\theta_1 = \theta_1^{(i)}$ for some i. For any fixed i, we can send θ_1 to $\theta_1^{(i)}$ and extend by linearity and multiplicativity, and since $\theta_1^{(i)}$ is a root of $T_1(X)$, this gives a K-algebra homomorphism from A to $\overline{K}$, which is injective if and only if A is a field. A similar statement is true for T_2.

Choose $k \in \mathbb{Z}$ different from the finite set of values

$$\frac{\theta_2^{(j)} - \theta_2^{(j')}}{\theta_1^{(i')} - \theta_1^{(i)}} \quad \text{for } i \neq i' \ ,$$

and set $\theta = \theta_2 + k\theta_1$. I claim that $A = K[\theta]$. Indeed, since θ_1 and θ_2 are elements of A, we have $K[\theta] \subset A$. Conversely, let

$$R(X) = \prod_{i,j}(X - (\theta_2^{(j)} + k\theta_1^{(i)})) \ .$$

By Galois theory or, equivalently, by the theorem on symmetric functions, or by the explicit formula $R(X) = \mathcal{R}_Y(T_1(Y), T_2(X - kY))$, which we will use below (where $\mathcal{R}_Y$ denotes the resultant with respect to the variable Y), we have $R \in K[X]$. Furthermore, our choice of k, and the fact that T_1 and T_2 are squarefree, ensure that

$$(i,j) \neq (i',j') \implies \theta_2^{(j)} + k\theta_1^{(i)} \neq \theta_2^{(j')} + k\theta_1^{(i')} \ .$$

Indeed, if we had equality in the right-hand side, we would have

$$k(\theta_1^{(i')} - \theta_1^{(i)}) = (\theta_2^{(j)} - \theta_2^{(j')}) \ .$$

Hence by our choice of k, we would have $\theta_1^{(i')} - \theta_1^{(i)} = \theta_2^{(j)} - \theta_2^{(j')} = 0$.Therefore $i' = i$ and $j' = j$ since the polynomials T_1 and T_2 are squarefree.

It follows that, with our choice of k, $R(X)$ is a squarefree polynomial in $K[X]$.

Let us come back to the elements of our algebra. By abuse of notation, set $R(X,Z) = \mathcal{R}_Y(T_1(Y), T_2(X - ZY))$ and $\theta(Z) = \theta_2 + Z\theta_1$, so that $R(X) = R(X,k)$ and $\theta = \theta(k)$. I claim that

$$R(\theta(Z), Z) = \mathcal{R}_Y(T_1(Y), T_2(\theta(Z) - ZY)) = 0 \ .$$

Note that this is not completely trivial: indeed, one cannot say that this follows from the fact that the polynomials $T_1(Y)$ and $T_2(\theta(Z) - ZY)$ have the common root $Y = \theta_1$, since this implies vanishing of the resultant only over a field.

Applying Proposition 2.1.3 to the ring $B = K[X,Z]$, we see that there exist *polynomials* $U_1(X,Y,Z)$ and $U_2(X,Y,Z)$ in $K[X,Y,Z]$ such that

$$U_1(X,Y,Z)T_1(Y) + U_2(X,Y,Z)T_2(X - ZY) = R(X,Z) \ .$$

Replacing Y by θ_1 and X by $\theta(Z) = \theta_2 + Z\theta_1$ gives $R(\theta(Z), Z) = 0$, as claimed. In particular, replacing Z by k, we obtain $R(\theta, k) = R(\theta) = 0$.

If we take the derivative of the equality $R(\theta(Z), Z) = 0$ with respect to the formal variable Z, and denote by R'_X and R'_Z the partial derivatives of $R(X,Z)$ with respect to X and Z, we obtain $\theta_1 R'_X(\theta(Z), Z) + R'_Z(\theta(Z), Z) = 0$, hence in particular for $Z = k$ the equation

$$\theta_1 R'_X(\theta, k) + R'_Z(\theta, k) = 0 \ . \tag{1}$$

Note that $R'_X(\theta, k) = R'(\theta)$. Since the polynomial $R(X) = R(X,k)$ is squarefree, there exist polynomials U and V in $K[X]$ such that $U(X)R'(X) + V(X)R(X) = 1$, and since $R(\theta) = 0$, we obtain $U(\theta)R'(\theta) = 1$. Hence, multiplying the identity (1) by $U(\theta)$, we obtain $\theta_1 = -U(\theta)R'_Z(\theta, k) \in K[\theta]$, and

evidently $\theta_2 = \theta - k\theta_1 \in K[\theta]$ also, so $A \subset K[\theta]$, thus finishing the proof of the lemma.

Note that we have proved much more, since we have also found a squarefree polynomial of which θ is a root, and an expression for θ_1 and θ_2 in terms of θ. We will use this explicitly in the algorithms given below. The expression for θ_1 comes from the implicit function theorem in an algebraic setting. □

We can now easily prove the main classical theorem about étale algebras.

Theorem 2.1.5 (Primitive Element Theorem). *Let K be a number field and A be an étale algebra of dimension n over K. There exists $\theta \in A$ (called a* primitive element*) such that $A = K[\theta]$, in other words such that $1, \theta, \ldots, \theta^{n-1}$ is a K-basis of A.*

Proof. Since A is finite-dimensional over K, there exist elements $\theta_1, \ldots, \theta_m$ such that $A = K[\theta_1, \ldots, \theta_m]$. For example, we can take for the θ_i a K-basis of A. We prove the theorem by induction on m. It is trivial for $m = 1$, and for $m = 2$ it is nothing else than Lemma 2.1.4. Let $m \geq 3$. By induction, we assume that we have proved it for all $i \leq m - 1$. Since the theorem is true for $m = 2$, we can find $\alpha \in A$ such that $K[\theta_{m-1}, \theta_m] = K[\alpha]$. But then $A = K[\theta_1, \ldots, \theta_{m-2}, \alpha]$ is generated by $m - 1$ elements, and we conclude by our induction hypothesis. □

As an immediate consequence we obtain the following corollary.

Corollary 2.1.6. *Let A be an étale algebra over K.*

(1) *There exists a squarefree monic polynomial $T(X) \in K[X]$ (called as above a* defining polynomial *for A/K) such that A is isomorphic to $K[X]/T(X)K[X]$.*

(2) *If $T(X) = \prod_{1 \leq i \leq g} T_i(X)$ is a decomposition of $T(X)$ into irreducible polynomials in $K[X]$, then A is isomorphic to the product $K_1 \times \cdots \times K_g$, where K_i is the number field defined by $K_i = K[X]/T_i(X)K[X]$.*

(3) *Conversely, any finite product of number fields over K (with componentwise multiplication and addition) is an étale algebra over K.*

Proof. (1). By Theorem 2.1.5, we know that $A = K[\theta]$ for some $\theta \in A$. If T is the minimal monic polynomial of θ in $K[X]$, then by definition of an étale algebra the polynomial $T(X)$ is squarefree, and the map sending θ to the class of X clearly gives an isomorphism from A to $K[X]/T(X)K[X]$.

(2). This is simply a restatement of the Chinese remainder theorem. More explicitly, since the polynomials $U_i(X) = T(X)/T_i(X)$ are (globally) coprime, there exist polynomials $V_i(X)$ such that $\sum_{1 \leq i \leq g} U_i(X)V_i(X) = 1$. We set $e_i = U_i(\theta)V_i(\theta)$. Then the map that sends the class of P in $K[X]/T(X)K[X]$ to $(\overline{e_1 P}, \ldots, \overline{e_g P})$ in $K_1 \times \cdots \times K_g$ is easily seen to be an algebra isomorphism (see Exercise 1).

(3). A product of fields cannot have nilpotent elements, so (3) is clear. □

As we shall see below, the introduction of étale algebras in addition to number fields is not simply a desire to generalize. First, they occur naturally in many contexts, as we shall see immediately, when we want to compute the compositum of number fields. Second, many, if not most, of the algorithms that we have given for number fields do not really use the nonexistence of zero divisors and hence are directly applicable to étale algebras. So from an algorithmic point of view, étale algebras are almost as simple as number fields.

Let us look at the important example of the notions of discriminant, integral basis, and the corresponding algorithms such as the round 2 algorithms (see [Coh0, Section 6.1]).

Let A be an étale algebra over K, which by Corollary 2.1.6 can be assumed to be equal to $K_1 \times \cdots \times K_g$ for some number fields K_j of degree n_j over K. Thus, $n = [A : K] = \prod_{1 \le j \le g} n_j$. We will identify each K_j with the subalgebra of A formed by elements whose components on K_l for $l \neq j$ are zero. Equivalently, if $e_j = (0, \ldots, 1, \ldots, 0) \in A$, where 1 is at the jth position, the e_j form a complete family of mutually orthogonal idempotents, and we identify K_j with $e_j A$.

If σ is a K-linear field homomorphism from A into $\mathbb{C}$, then on each K_j, σ restricts to some K-linear embedding $\sigma_{i,j}$ of K_j into $\mathbb{C}$. Conversely, if for each j we choose K-linear embeddings $\sigma_{i,j}$ of K_j into $\mathbb{C}$, it is clear that there exists a unique K-linear field homomorphism σ from A to $\mathbb{C}$ which restricts to $\sigma_{i,j}$ for each j given by

$$\sigma\left(\sum_j x_j e_j\right) = \sum_j \sigma_{i,j}(x_j) e_j \ .$$

Hence, as in the number field case, there exist $n = \prod n_j$ K-linear ring homomorphisms of A into $\mathbb{C}$. These are no longer embeddings, however; in other words, they are not injective since A is not a field in general. As in the number field case, we will denote them by σ_i.

The notions of trace and norm and more generally of characteristic polynomial now generalize without difficulty and can be expressed either directly, for example, via resultants, or through the embeddings σ_i.

The definitions of discriminant and integral basis (or pseudo-basis in the relative case; see Section 2.2.3) then go through without change. The reader who wants to explore further is warned, however, that the *separability* condition mentioned in [Coh0, remark after the proof of Proposition 4.4.1] is essential, although we can ignore it as long as we are in characteristic zero.

Finally, the reader can easily check that the description of the absolute round 2 algorithm given in [Coh0, Section 6.1] is valid without modification for étale algebras. In fact, it is applicable in even more general contexts.

Similarly, the relative round 2 algorithm that will be given in Section 2.4 is also valid without change.

We can also define the Galois group of an étale algebra and give algorithms to compute it, analogous to the algorithms given in [Coh0, Section 6.3]. Note, however, that the Galois groups are no longer *transitive* subgroups of S_n; hence their classification is usually more complex (see Exercises 5 to 7).

2.1.3 Compositum of Two Number Fields

We consider the following problem. Let L_1 and L_2 be two number fields defined over the base field K by their relative defining polynomials $T_1(X)$ and $T_2(X)$, respectively. We would like to compute the compositum L of L_1 and L_2, which by definition is the smallest number field containing both L_1 and L_2.

As such, the above problem does not make any sense, for two reasons. First, we must embed L_1 and L_2 into a fixed algebraic closure $\overline{K}$ of K. Once this is done, the compositum does make sense, but on the other hand the polynomial T_1 alone does not in general determine the number field L_1 since any root of T_1 can be chosen, so it is impossible to distinguish L_1 from its conjugate fields over K. Thus, to determine L_1 we must give not only a polynomial T_1 but also some way to distinguish the root θ_1 of T_1 such that $L_1 = K(\theta_1)$ from the other roots of T_1.

Let $L = L_1L_2$ be the compositum of L_1 and L_2. To find a polynomial defining L over K, we prove the following proposition.

Proposition 2.1.7. *Let $L_1 = K(\theta_1)$ and $L_2 = K(\theta_2)$ be two number fields defined over K, and let T_1 (resp., T_2) be the minimal monic polynomial of θ_1 (resp., θ_2) over K. Set*

$$R(X,Z) \leftarrow \mathcal{R}_Y(T_1(Y), T_2(X - ZY)) \ ,$$

where $\mathcal{R}_Y$ denotes the resultant with respect to Y. Then we have the following.

(1) *There exists an integer $k \in \mathbb{Z}$ such that the polynomial $R(X,k)$ is squarefree.*
(2) *If k is chosen as in* (1)*, then the compositum $L = L_1L_2$ is given by $L = K(\theta)$ with $\theta = \theta_2 + k\theta_1$, and the minimal polynomial $T(X)$ of θ is one of the irreducible factors of $R(X,k)$ in $K[X]$.*
(3) *If k and θ are as in* (1) *and* (2)*, we have*

$$\theta_1 = -\frac{R'_Z}{R'_X}(\theta, k), \quad \theta_2 = \theta - k\theta_1, \quad \text{with} \quad R'_Z = \frac{\partial R}{\partial Z}, \ R'_X = \frac{\partial R}{\partial X} \ .$$

Proof. By definition, $L = K(\theta_1, \theta_2)$. By the proof of Lemma 2.1.4, there exists $k \in \mathbb{Z}$ such that $R(X,k)$ is squarefree, and if $\theta = \theta_2 + k\theta_1$, then $L = K(\theta)$ with θ a root of $R(X,k) = 0$. Since L is a field, the minimal

polynomial of θ is an irreducible factor of $R(X,k)$ in $K[X]$. Finally, in that proof, we have also seen that $\theta_1 R'_X(\theta,k) + R'_Z(\theta,k) = 0$, and since we are in a field and $R(X,k)$ is squarefree, we obtain the formula for θ_1, hence for θ_2, given in the proposition. □

In practice, a number field is often given only up to isomorphism, and thus it is not possible to specify a specific root θ_1, but only the polynomial $T_1(X)$ of which it is a root. In that case, all irreducible factors of the resultant $R(X,k)$ give a possible compositum, and it makes perfectly good sense to consider them all. In fact, since $R(X,k)$ is squarefree, $K[X]/R(X,k)K[X]$ is an étale algebra that is isomorphic to the product of all possible compositums of the number fields L_1 and L_2 defined by T_1 and T_2. It is reasonable to call this algebra *the* compositum of the number fields defined by T_1 and T_2. Thus, we are led to the following algorithm.

Algorithm 2.1.8 (Compositum of Two Number Fields). Given two irreducible polynomials T_1 and T_2 in $K[X]$, this algorithm computes the relative defining polynomials for all possible compositums $L = K(\theta)$ of the number fields determined by T_1 and T_2, respectively, and expresses the generic roots θ_1 and θ_2 of T_1 and T_2 in terms of θ.

1. [Compute resultant] Using, for example, the subresultant algorithm ([Coh0, Algorithm 3.3.7]) over the ring $K[X,Z]$, compute

$$R(X,Z) \leftarrow \mathcal{R}_Y(T_1(Y), T_2(X - ZY)) \ ,$$

   where X and Z are formal variables. We denote as above by R'_X (resp., R'_Z) the partial derivative of $R(X,Z)$ with respect to X (resp., Z).

2. [Find integer k] For $k = \pm 1, \pm 2, \ldots$, compute $s \leftarrow \gcd(R(X,k), R'_X(X,k))$ until $s = 1$.

3. [Compute θ_1] (Here $R(X,k)$ is squarefree.) Using the extended Euclidean algorithm, compute polynomials U and V in $K[X]$ such that $U(X)R'_X(X,k) + V(X)R(X,k) = 1$, and set $A_1(X) \leftarrow -U(X)R'_Z(X,k) \bmod R(X,k)$.

4. [Factor $R(X,k)$] Using, for example, [Coh0, Algorithm 3.5.7] if $K = \mathbb{Q}$ and [Coh0, Algorithm 3.6.4] otherwise, factor $R(X,k)$ in $K[X]$ as $R(X,k) = \prod_{1 \le i \le g} R_i(X)$ (we already know that $R(X,k)$ is squarefree).

5. [Terminate] For $i = 1$ to $i = g$, output $R_i(X)$ as the irreducible relative defining polynomial of a compositum of number fields determined by T_1 and T_2, output $\theta_1 \leftarrow A_1(X) \bmod R_i(X)$, $\theta_2 \leftarrow X - kA_1(X) \bmod R_i(X)$ as roots of T_1 and T_2, respectively, and terminate the algorithm.

Remark. The polynomials $R_i(X)$ output by the above algorithm often have large coefficients, hence it is almost always necessary to modify them before doing further work. For this, if the base field K is equal to $\mathbb{Q}$, we use a polynomial reduction algorithm such as the one in [Coh0, Algorithm 4.4.12].

In the general case, we have to use relative polynomial reduction algorithms which will be described in Section 2.4.2. In any case, we obtain a polynomial $B_i(X)$ such that the minimal monic polynomial over K of $\eta = B_i(\theta)$ is a polynomial $S_i(X)$ that is hopefully simpler than $R_i(X)$ (up to a multiplicative constant, we have $S_i(X) = \mathcal{R}_Y(R_i(Y), X - B_i(Y))$). Since we now work with $S_i(X)$ and its root η, we must express θ_1 and θ_2 in terms of η. Using Algorithm 2.1.12 below, we can compute a polynomial B_i^{-1} such that $\theta = B_i^{-1}(\eta)$; hence $\theta_1 = C_1(\eta)$ and $\theta_2 = C_2(\eta)$ with

$$C_1(X) = A_1(B_i^{-1}(X)) \bmod S_i(X) \quad \text{and}$$
$$C_2(X) = B_i^{-1}(X) - kA_1(B_i^{-1}(X)) \bmod S_i(X) \ .$$

Another way to obtain smaller polynomials is to modify the type of elements chosen in Algorithm 2.1.8. Instead of trying elements of the form $\theta = \theta_2 + k\theta_1$, we may try more generally any polynomial in θ_1 and θ_2 with rational coefficients. This usually leads to rather complicated computations unless the polynomials are very simple, such as the linear polynomials we have just chosen. But it is also reasonable to look at $\theta = \theta_1\theta_2 + k_1\theta_1 + k_2\theta_2$ for small integers k_1 and k_2. This means that instead of computing $R(X, k) = \mathcal{R}_Y(T_1(Y), T_2(X - kY))$, we compute

$$R(X, k_1, k_2) = \mathcal{R}_Y(T_1(Y - k_2), T_2((X - k_1Y + k_1k_2)/Y)Y^{n_2})$$

with $n_2 = \deg(T_2)$. As usual, if this polynomial is squarefree it defines the compositum of the number fields defined by T_1 and T_2 as an étale algebra. We can also recover θ_1 and θ_2 by proving a proposition analogous to Proposition 2.1.7 (see Exercise 8). It is, however, much better in this case to use the direct method explained in Section 2.1.4.

Very often we can simply take $k_1 = k_2 = 0$, hence $\theta = \theta_1\theta_2$, and in this case it frequently happens that the polynomial $R(X, 0, 0)$ is simpler than the polynomial output by Algorithm 2.1.8.

We give an example. Assume that $K = \mathbb{Q}$ and that $T_1(X) = T_2(X) = X^3 - 2$. We apply Algorithm 2.1.8. After step 1, we find that

$$R(X, Z) = X^9 - 6(Z^3 + 1)X^6 + 12(Z^6 - 7Z^3 + 1)X^3 - 8(Z^3 + 1)^3 \ .$$

In step 2, $k = 1$ and $k = -1$ do not work, but both $k = \pm 2$ work, so we choose $k = 2$, for example, so $R(X, 2) = X^9 - 54X^6 + 108X^3 - 5832$. In step 3, we obtain $A_1(X) = (X^7 - 63X^4 + 1242X)/2268$. In step 4, we get the factorization into irreducibles in $\mathbb{Q}[X]$ as $R(X, 2) = (X^3 - 54)(X^6 + 108)$. In step 5, we first output $R_1(X) = X^3 - 54$, which is the trivial compositum of the field $\mathbb{Q}(2^{1/3})$ with itself, and $\theta_1 \leftarrow A_1(X) \bmod R_1(X)$ gives $X/3 \bmod R_1(X)$, which is indeed the change of variable necessary to transform $R_1(X)$ into the initial polynomial $X^3 - 2$. Note that $\theta_2 = X - 2\theta_1$ is also equal to $X/3$, as it should be.

We then output $R_2(X) = X^6 + 108$, which is the nontrivial compositum of $\mathbb{Q}(2^{1/3})$ with itself, hence its Galois closure, as well as $\theta_1 = -X^4/36 + X/2 \bmod R_2(X)$ and $\theta_2 = X^4/18 \bmod R_2(X)$. Note that these formulas for θ_1 and θ_2 come immediately from the algorithm but would not have been so simple to obtain directly. The reader can check that $\theta_1^3 \equiv \theta_2^3 \equiv 2 \pmod{X^6 + 108}$.

A polynomial reduction algorithm such as [Coh0, Algorithm 4.4.11] gives the new polynomial $S_2(X) = X^6 - 3X^5 + 5X^3 - 3X + 1$ (whose discriminant is more than 10^{10} times smaller than that of $R_2(X)$) and the polynomial $B_2(X) = -X^5/54 - X^3/36 + X/3 + 1/2$. Algorithm 2.1.12 gives us $B_2^{-1}(X) = 4X^5 - 10X^4 - 6X^3 + 19X^2 + 11X - 9$, and so we obtain finally $C_1(X) = 2X^5 - 5X^4 - 3X^3 + 10X^2 + 5X - 5$ and $C_2(X) = -X^2 + X + 1$.

On the other hand, if we want to use $\theta = \theta_1\theta_2 + k_1\theta_1 + k_2\theta_2$, we find that $(k_1, k_2) = (0, 0)$ does not work (it will never work when $T_1 = T_2$; see Exercise 10), but $(k_1, k_2) = (-1, 0)$ works and gives

$$\begin{aligned} R(X, -1, 0) &= X^9 - 6X^6 + 228X^3 - 8 \\ &= (X^3 + 6X - 2)(X^6 - 6X^4 - 4X^3 + 36X^2 + 12X + 4) \ . \end{aligned}$$

The third degree factor defines as usual the same number field defined by T_1 and T_2, and the sixth degree factor defines its Galois closure. Although more coefficients are nonzero, it is a slightly simpler polynomial than the polynomial $X^6 + 108$ (for example, its discriminant is 144 times smaller), and of course polynomial reduction leads to the same polynomial as the one found above.

2.1.4 Computing θ_1 and θ_2

The formula $\theta_1 = -R'_Z(\theta, k)/R'_X(\theta, k)$ implicitly used in Algorithm 2.1.8 has the advantage of simplicity but is usually not the most efficient. Indeed, although $R'_X(X, k)$ can be obtained as the derivative of the single-variable resultant $R(X)$, there does not seem to be any direct way of computing $R'_Z(\theta, k)$ without computing a two-variable resultant $R(X, Z)$. This, however, is a rather expensive operation if we directly use the subresultant algorithm. It is generally better to use modular variants, which amounts to computing $R(X, k)$ for several values of k, which is exactly what is needed in step 2. There are at least two ways to obtain θ_1 without knowing $R(X, Z)$ as a two-variable polynomial.

The first way is by looking more closely at the structure of the subresultant algorithm ([Coh0, Algorithm 3.3.7]). This algorithm follows the steps of an ordinary Euclidean algorithm, except that pseudo-divisions are used instead of divisions. In our case, this means that we start with the polynomials $T_1(Y)$ and $T_2(X - kY)$ considered as polynomials in Y only and essentially perform successive Euclidean steps until we reach a constant polynomial in

Y, which will be the desired resultant $R(X) = R(X,k)$ if we follow the normalizations of the algorithm properly. At each stage, the polynomial in Y is a linear combination with coefficients in $K[X,Y]$ of the polynomials $T_1(Y)$ and $T_2(X - kY)$. In particular, they will vanish when we set simultaneously $Y = \theta_1$ and $X = \theta_2 + k\theta_1$.

It can be shown that, when the final remainder $R(X)$ is squarefree, the degree of the *preceding* polynomial in the sequence will be exactly equal to 1 in the variable Y, and up to a constant, it will be equal to $R'(X)Y + R'_Z(X,k)$ (see Exercise 12). Since this will vanish when we set $Y = \theta_1$ and $X = \theta_2 + k\theta_1$, we obtain the formula $\theta_1 = -R'(X)/R'_Z(X,k)$ as before, but without the need for computing $R(X,Z)$ explicitly.

The second way to compute θ_1 is direct. Let N be the $n_1n_2 \times n_1n_2$ matrix whose rows are indexed by pairs (i_1,i_2) with $0 \le i_1 < n_1$ and $0 \le i_2 < n_2$, and whose columns, indexed by j for $0 \le j < n_1n_2$, contain the coefficients of $\theta^j = (\theta_2 + k\theta_1)^j$ on $\theta_1^{i_1}\theta_2^{i_2}$, which can easily be computed by induction using the polynomials T_1 and T_2. Since we know that θ_1 belongs to $K(\theta)$, the column vector V representing θ_1 (whose entries are equal to 0 except for $(i_1,i_2) = (1,0)$ for which $V_{(1,0)} = 1$) belongs to the image of N in $K^{n_1n_2}$. By Gaussian elimination, we can thus find a column vector W such that $V = NW$, and θ_1 is thus equal to $(1,\theta,\dots,\theta^{n_1n_2-1})W$.

Note that if we add an (n_1n_2+1)st column to the matrix N representing $\theta^{n_1n_2}$, the ordinary *kernel* of this matrix gives the polynomial $R(X,k)$, thus giving a way other than the subresultant to compute it. Practice shows that, suitably implemented, these ideas lead to much better performance than implementations based on the subresultant algorithm, even with the improvement mentioned above (see Exercise 13).

Both of the methods just described can of course also be applied to the case where one chooses $\theta = \theta_1\theta_2$ or more generally $\theta = \theta_1\theta_2 + k\theta_2$. Since this case often gives simpler results, we isolate it as a formal algorithm.

Algorithm 2.1.9 (Compositum of Two Number Fields Using $\theta_1\theta_2$). Let $T_1(X)$ and $T_2(X)$ be two monic irreducible polynomials in $K[X]$ of degree n_1 and n_2. This algorithm computes a relative defining polynomial for all the possible compositums $L = K(\theta)$ of the number fields determined by T_1 and T_2, respectively, and expresses the generic roots θ_1 and θ_2 of T_1 and T_2 in terms of θ.

1. [Modify θ_1] By trying $k = 0, \pm1$, etc., find k such that the characteristic polynomial of $\theta_2(\theta_1 + k)$ is squarefree. Set $T_1(X) \leftarrow T_1(X - k)$, then write $T_1(X) = \sum_{0 \le i_1 \le n_1} t_{1,i_1}X^{i_1}$ and $T_2(X) = \sum_{0 \le i_2 \le n_2} t_{2,i_2}X^{i_2}$.
2. [Set up big matrix] Set $n \leftarrow n_1n_2$, and construct the $n \times (n+1)$ matrix $N = (N_{(i_1,i_2),j})$ whose rows are indexed by pairs (i_1,i_2) with $0 \le i_1 < n_1$ and $0 \le i_2 < n_2$, and whose columns are indexed by integers j such that $0 \le j \le n$ as follows. Set $N_{(0,0),0} \leftarrow 1$, $N_{(i_1,i_2),0} \leftarrow 0$ for all $(i_1,i_2) \ne (0,0)$. Then for $j = 0,\dots,j = n-1$, and for all (i_1,i_2), set

$$N_{(i_1,i_2),j+1} \leftarrow N_{(i_1-1,i_2-1),j} - t_{1,i_1} N_{(n_1-1,i_2-1),j} - t_{2,i_2} N_{(i_1-1,n_2-1),j} + t_{1,i_1} t_{2,i_2} N_{(n_1-1,n_2-1),j} \ ,$$

where $N_{(i_1,i_2),j}$ is taken equal to 0 if either $i_1 < 0$ or $i_2 < 0$.

3. [Compute kernel] Using [Coh0, Algorithm 2.3.1], compute the kernel of N, which will be a one-dimensional space generated by some element $(r_0, \ldots, r_n)^t$ with $r_n \neq 0$ which we normalize so that $r_n = 1$. Set $R(X) \leftarrow \sum_{0 \le j \le n} r_j X^j$.
4. [Compute inverse image] Change the nth column of the matrix N by setting $N_{(1,0),n} \leftarrow 1$ and $N_{(i_1,i_2),n} \leftarrow 0$ for $(i_1, i_2) \neq (1, 0)$. Then once again using [Coh0, Algorithm 2.3.1], compute the kernel of this new matrix N, which must again be a one-dimensional space generated by some element $(a_0, \ldots, a_n)^t$ with $a_n \neq 0$ which we normalize so that $a_n = -1$.
5. [Compute θ_1 and θ_2] Set $A_1(X) \leftarrow \sum_{0 \le j < n} a_j X^j$. Using the extended Euclidean algorithm, compute polynomials U and V in $K[X]$ such that $U(X)A_1(X) + V(X)R(X) = 1$ and set $A_2(X) \leftarrow XU(X) \bmod R(X)$.
6. [Factor $R(X)$] Using, for example, [Coh0, Algorithm 3.5.7] if $K = \mathbb{Q}$ and [Coh0, Algorithm 3.6.4] otherwise, factor $R(X)$ in $K[X]$ as $R(X) = \prod_{1 \le i \le g} R_i(X)$ (we already know that $R(X)$ is squarefree).
7. [Terminate] For $i = 1$ to $i = g$, output $R_i(X)$ as the irreducible defining polynomial of a compositum of number fields determined by T_1 and T_2, output $\theta_1 \leftarrow A_1(X) - k \bmod R_i(X)$, $\theta_2 \leftarrow A_2(X) \bmod R_i(X)$ as roots of the initial T_1 and T_2, respectively, where k has been computed in step 1, and terminate the algorithm.

Remarks

(1) If $k = 0$ does not work in step 1, we change θ_1 into $\theta_1 + k$ and hence $T_1(X)$ into $T_1(X - k)$ for some small k. The proof of the existence of such a k is essentially identical to the proof of the primitive element theorem (Exercise 2). To avoid any risk of confusion, it is preferable to do this change *before* using this algorithm, and forget about the initial polynomial $T_1(X)$ entirely, rather than handling $\theta_2(\theta_1 + k)$. In other words, we could reasonably ask the algorithm also to output the new polynomial $T_1(X)$, and set $\theta_1 \leftarrow A_1(X) \bmod R_i(X)$ in the last step.

(2) After computing $\theta_1 + k$, to recover θ_2 we use in step 5 the trivial formula $\theta_2 = \theta/(\theta_1 + k)$. We can also obtain θ_2 directly by still another kernel computation where we set $N_{(0,1),n} \leftarrow 1$ and $N_{(i_1,i_2),n} \leftarrow 0$ for $(i_1, i_2) \neq (0, 1)$.

(3) Since we have two (or three if we use the preceding remark) matrix kernels to compute of $n \times (n + 1)$ matrices whose first n columns are the same, we can considerably speed up the algorithm by solving the two (or three) linear systems at once. This is a simple modification of the algorithm for computing a matrix inverse [Coh0, Algorithm 2.2.2], where the work is done only on two (or three) columns instead of n. The details are left

to the reader (see Exercise 14), but an actual implementation must use this.

One last problem can be asked in the context of the compositum of number fields. Let K_1 and K_2 be two extensions of K determined by the polynomials T_1 and T_2 as above, and let L be the étale algebra compositum of K_1 and K_2, so that a defining polynomial for L/K is the resultant $R(X, k)$, and $\theta = \theta_2 + k\theta_1$ (if desired, instead of the full compositum L, we could also consider one of the number fields obtained by factoring $R(X, k)$). We would like to compute a relative defining polynomial for L/K_1 and for L/K_2. The answer to this problem is trivial but deserves to be mentioned. Since $\theta = \theta_2 + k\theta_1$, we clearly have $T_2(\theta - k\theta_1) = T_2(\theta_2) = 0$. Hence, if we set $U_1(X) = T_2(X - k\theta_1)$, we have $U_1 \in K_1[X]$ and $U_1(\theta) = 0$. In addition, since L is the compositum considered as an étale algebra, we have $[L : K_1] = [K_2 : K] = \deg(T_2)$, so U_1 is the minimal polynomial of θ over K_1; hence it is a relative defining polynomial for L/K_1 (if we had taken L to be a number field associated to an irreducible factor of $R(X, k)$, we would have had to consider a suitable factor of U_1).

Similarly, since $k \neq 0$, $U_2(X) = k^{\deg(T_1)} T_1((X - \theta_2)/k)$ is the minimal polynomial of θ over K_2; hence it is a relative defining polynomial for L/K_2.

2.1.5 Relative and Absolute Defining Polynomials

Let L_1 be a number field over K, defined as $L_1 = K(\theta_1)$, where θ_1 is a root of the irreducible polynomial $T_1 \in K[X]$ of degree n_1. Let $L_2 = L_1(\theta_2)$ be a relative extension, defined by a root θ_2 of the polynomial $T_2 \in L_1[X]$, irreducible over L_1 of degree $n_2 = [L_2 : L_1]$. In this section, we give an algorithm that allows us to go back and forth from the representation of L_2 as an L_1-extension to the representation of L_2 as a K-extension. (In the case where $K = \mathbb{Q}$, this of course allows us to go back and forth from relative to absolute defining polynomials.) We will see that this is a natural generalization of the algorithm for computing the compositum of two number fields (Algorithms 2.1.8 or 2.1.9).

The following theorem is the analog (in fact, a generalization) of Proposition 2.1.7.

Theorem 2.1.10. *Let $L_1 = K(\theta_1)$ and $L_2 = L_1(\theta_2)$ be two number fields, where θ_1 is a root of the irreducible polynomial $T_1(X) \in K[X]$ of degree n_1, and θ_2 is a root of the polynomial $T_2(X) \in L_1(X)$ of degree n_2, assumed to be irreducible in $L_1(X)$. If $T_2(X) = \sum_{m=0}^{n_2} A_m(\theta_1) X^m$, we set $W(X, Y) = \sum_{m=0}^{n_2} A_m(Y) X^m$, which makes sense only modulo $T_1(Y)K[X]$. Set*

$$R(X, Z) = \mathcal{R}_Y(T_1(Y), W(X - ZY, Y)) \ .$$

Then we have the following.

(1) *There exists an integer* $k \in \mathbb{Z}$ *such that the polynomial* $R(X,k)$ *is squarefree.*
(2) *If* k *is chosen as in* (1)*, then* $R(X,k)$ *is irreducible in* $K[X]$*, and* $L_2 = K(\theta)$*, where* $\theta = \theta_2 + k\theta_1$ *is a root of* $R(X,k)$*.*
(3) *If* k *and* θ *are as in* (1) *and* (2)*, we have*

$$\theta_1 = -\frac{R'_Z}{R'_X}(\theta,k), \quad \theta_2 = \theta - k\theta_1 \ .$$

Proof. The proof is very close to that of Lemma 2.1.4 and Proposition 2.1.7 (see also [Coh0, Lemma 3.6.2]).

(1). Let $\Omega = \overline{L_2}$ be some algebraic closure of L_2. Then Ω is also an algebraic closure of K and of L_1. We denote by $\theta_1^{(i)}$ (resp., $\theta_2^{(j)}$) the roots of T_1 (resp., T_2) in Ω, chosen so that $\theta_1 = \theta_1^{(1)}$ and $\theta_2 = \theta_2^{(1)}$. Note that the $\theta_1^{(i)}$ (resp., the $\theta_2^{(j)}$) are distinct since T_1 and T_2 are irreducible and in particular squarefree. Let $k \in \mathbb{Z}$. The roots of $R(X,k)$ in Ω are the numbers X such that there exists a common root of $T_1(Y)$ and $W(X - kY, Y)$, so that $Y = \theta_1^{(i)}$ and $W(X - k\theta_1^{(i)}, \theta_1^{(i)}) = 0$.

Set

$$T_2^{(i)} = \sum_{m=0}^{n_2} A_m(\theta_1^{(i)})X^m = W(X, \theta_1^{(i)}) \ ,$$

and let $\theta_2^{(i,j)}$ be the roots of $T_2^{(i)}$ in Ω, ordered so that $\theta_2^{(1,j)} = \theta_2^{(j)}$. Thus the roots of $R(X,k)$ are the numbers $\gamma^{(i,j)} = \theta_2^{(i,j)} + k\theta_1^{(i)}$. Furthermore, using as before Sylvester's determinant, it is easy to show that $R(X,k)$ is a polynomial in X of degree at most equal to $n_1 n_2$. If we choose $k \in \mathbb{Z}$ different from the finite set of values

$$\frac{\theta_2^{(i,j)} - \theta_2^{(i',j')}}{\theta_1^{(i')} - \theta_1^{(i)}} \quad \text{for } i \neq i' \ ,$$

the $n_1 n_2$ values $\gamma^{(i,j)}$ are distinct, and hence the polynomial $R(X,k)$ is squarefree of degree exactly equal to $n_1 n_2$, proving (1).

We prove (2) and (3) simultaneously. Let k be chosen as in (1). Keeping the notation of the proof of Lemma 2.1.4, we obtain without change that for $\theta = \theta_2 + k\theta_1$, we have $\theta_1 = -(R'_Z/R'_X)(\theta,k)$ and $\theta_2 = \theta - k\theta_1$. Since θ is a root of $R(X,k) \in K[X]$, it is a root of some irreducible factor $R_1(X)$ of $R(X,k)$ in $K[X]$. But then the number field $K(\theta)$ defined over K by the polynomial R_1 contains L_1 (since θ_1 is a rational function of θ) and hence also contains L_2 since $\theta_2 = \theta - k\theta_1$. Since $[L_2 : K] = [L_2 : L_1][L_1 : K] = n_1 n_2 = \deg(R(X,k))$, it follows that $\deg(R_1) = \deg(R(X,k))$; in other words, $R(X,k)$ is irreducible in $K[X]$. □

The corresponding algorithm is also essentially identical to Algorithm 2.1.8. Considering its importance it is useful to give it separately.

Algorithm 2.1.11 (Relative to Absolute Defining Polynomial). Given an irreducible polynomial $T_1 \in K[X]$ defining a number field $L_1 = K(\theta_1)$ and a polynomial $T_2 \in L_1[X]$, irreducible in $L_1[X]$, hence defining a number field $L_2 = L_1(\theta_2)$, this algorithm computes a defining polynomial for L_2 over K, in other words, an irreducible polynomial $R \in K[X]$ such that $L_2 = K(\theta)$ with θ a root of R. Furthermore, it also computes θ_1 and θ_2 as polynomials in θ, and the small integer k such that $\theta = \theta_2 + k\theta_1$.

1. [Compute resultant] If $T_2(X) = \sum_m A_m(\theta_1)X^m$ for some polynomials A_m, set $W(X,Y) \leftarrow \sum_m A_m(Y)X^m$. Using, for example, the subresultant algorithm ([Coh0, Algorithm 3.3.7]) over the ring $K[X,Z]$, compute

$$R(X,Z) \leftarrow \mathcal{R}_Y(T_1(Y), W(X - ZY, Y)) \ ,$$

   where X and Z are formal variables. We will denote by R'_X (resp., R'_Z) the partial derivative of $R(X,Z)$ with respect to X (resp., Z).

2. [Find integer k] For $k = 0, \pm 1, \pm 2, \ldots$, compute

$$s \leftarrow \gcd(R(X,k), R'_X(X,k))$$

   until $s = 1$.

3. [Terminate] (Here $R(X,k)$ is irreducible.) Using the extended Euclidean algorithm, compute polynomials U and V in $K[X]$ such that $U(X)R'_X(X,k) + V(X)R(X,k) = 1$, and set $\theta_1 \leftarrow -U(X)R'_Z(X,k) \bmod R(X,k)$. Output the defining polynomial $R(X,k)$ for L_2/K, output k, θ_1, $\theta_2 = X - k\theta_1 \bmod R(X)$, and terminate the algorithm.

Again we give an example. Let $K = \mathbb{Q}$, $L_1 = K(\theta_1)$, where θ_1 is a root of the polynomial $T_1(X) = X^3 - 2$, and $L_2 = L_1(\theta_2)$, where θ_2 is a root of the polynomial $T_2(X) = X^2 - \theta_1 X + 1$. We compute that

$$\begin{aligned} R(X,Z) = X^6 &+ 3X^4 - 2(2Z^3 + 3Z^2 + 3Z + 1)X^3 \\ &+ 3X^2 + 6(2Z^3 + 3Z^2 + Z)X \\ &+ (4Z^6 + 12Z^5 + 12Z^4 + 4Z^3 + 1) \ . \end{aligned}$$

Here we can take $k = 0$ (this was never possible when computing a compositum; see Exercise 16), and hence $R(X) = X^6 + 3X^4 - 2X^3 + 3X^2 + 1$ is an absolute defining polynomial for L_2. Furthermore, the computations of step 3 show that $\theta_1 = -(X^5 + 3X^3 - 2X^2 - 2X) \bmod R(X)$.

The above algorithm is used when it is necessary to have a defining polynomial for L_2 over K and not only over L_1 (although it is usually better *not* to work over K, but sometimes it is impossible to do otherwise). The results are then used as follows. If $\alpha \in L_1$ is given as a polynomial in θ_1, thanks to the formula expressing θ_1 in terms of θ we can immediately compute α in terms of θ and hence consider it as an element of L_2. Conversely,

if $\alpha = \sum_{0 \le i < n_1 n_2} a_i \theta^i$ is an element of L_2 which for some reason is known to belong to L_1 (for example, α may be the relative trace or norm of some other element of L_2), we want to express α as a polynomial in θ_1 and not only as a polynomial in θ. Thus, we want to find $b_j \in K$ such that $\alpha = \sum_{0 \le j < n_1} b_j \theta_1^j$. Since θ_1 is known as a polynomial in θ, we can compute (once and for all if this has to be done for several α) coefficients $c_{i,j}$ for $0 \le i < n_1 n_2$ and $0 \le j < n_1$, such that $\theta_1^j = \sum_{0 \le i < n_1 n_2} c_{i,j} \theta^i$. Finding the b_j is thus equivalent to solving the linear system of $n_1 n_2$ equations in n_1 unknowns $\sum_{0 \le j < n_1} c_{i,j} a_j = b_i$ for $0 \le i < n_1 n_2$. This is done via a straightforward pivoting method (see, for example, [Coh0, Algorithm 2.3.4]). Since there are many more equations than unknowns in the system, this gives an excellent check of the correctness of preceding computations. In fact, the system has a solution if and only if α does belong to L_1, so it provides a new verification of this fact.

Remark. It is important to note the similarities and differences between the two problems studied above. The computation of the compositum of two number fields corresponds clearly to the special case of the computation of a relative extension L_2/L_1 in which the defining polynomial $T_2(X)$ not only belongs to $L_1[X]$, but in fact to $K[X]$, or in other words, that the corresponding two-variable polynomial $W(X,Y)$ defined above does not depend on Y. However, it is not quite a special case, since in the computation of absolute defining polynomials, we assume that the polynomial T_2 is irreducible in $L_1(X)$, while for the compositum we have the weaker assumption that T_2 should be irreducible in $K[X]$.

We could, however, write a common theorem and a common algorithm by considering finite étale algebras over K instead of field extensions of K, and in that case the polynomial T_2 need not be irreducible. We leave this as an easy exercise for the reader (Exercise 17).

Finally, note that most of the remarks made after Algorithm 2.1.8 — in particular those about other ways of computing θ_1 — still apply here and must be used in a serious implementation.

Another interesting special case of the problem of computing absolute defining polynomials is the *reversion* of an algebraic number. Assume that $L_1 = K(\theta_1)$ is a number field of degree n over K defined by an irreducible polynomial T_1, and let $\theta_2 = A(\theta_1)$, which is also known to be of degree n. Since $K(\theta_2) = K(\theta_1)$, we must be able to express θ_1 in terms of θ_2. Using Algorithm 2.1.11, we let $L_2 = L_1(\theta_2)$, where θ_2 is a root of the polynomial T_2 of degree 1 over $L_1(X)$ defined by $T_2(X) = X - A(\theta_1)$. Clearly, $L_2 = L_1$. We have $W(X,Y) = X - A(Y)$, and hence $R(X,Z) = \mathcal{R}_Y(T_1(Y), X - ZY - A(Y))$. The polynomial $R(X,0) = \mathcal{R}_Y(T_1(Y), X - A(Y))$ is the characteristic polynomial of α in L_1 (see [Coh0, Proposition 4.3.4]). Hence, since α is of degree n, $R(X,0)$ is irreducible and is thus equal to the minimal polynomial of α. Theorem 2.1.10 then tells us that $\theta_1 = -R'_Z(\theta_2,0)/R'_X(\theta_2,0)$.

We can also solve the problem directly using the following algorithm.

Algorithm 2.1.12 (Reversion of an Algebraic Number). Let $L = K(\theta_1)$ be a number field of degree n over K defined by an irreducible polynomial T_1, and let $\theta_2 = A(\theta_1)$ be an element of degree n. This algorithm computes a polynomial $B(X) \in K[X]$ of degree less than n such that $\theta_1 = B(\theta_2)$.

1. [Compute powers of θ_2] For $0 \leq j < n$, compute $A_j(X) \leftarrow A^j(X) \bmod T_1(X)$, and let $A_j(X) = \sum_{0 \leq i < n} a_{i,j} X^i$.
2. [Solve linear system] Let $M = (a_{i,j})$ be the matrix of the $a_{i,j}$. Using ordinary Gaussian pivoting in K, find a solution $B = (b_0, \ldots, b_{n-1})^t$ to the linear system $MB = (1, 0, \ldots, 0)^t$. If the system has no solution, output an error message saying that θ_2 is of degree strictly less than n and terminate. If the system has more than one solution, output an error message saying that θ_1 is of degree strictly less than n and terminate.
3. [Terminate] Set $B(X) \leftarrow \sum_{0 \leq j < n} b_j X^j$, output $B(X)$, and terminate the algorithm.

Proof. The (easy) proof of this algorithm's validity is left to the reader (Exercise 18). □

Finally, consider the following problem. Assume that we have a field extension L/K and that, in addition to the data for the base field K, we know only an absolute defining polynomial for L (over $\mathbb{Q}$ or some other subfield k of K). We want to find a defining polynomial for L/K. This is simply the *subfield* problem considered in [Coh0, Section 4.5]. There are thus several methods to do this, which are equivalent to factoring the absolute polynomial defining L in the number field K. Any such factor of the correct degree gives a relative defining polynomial. Note, however, that if the number fields are specified together with embeddings (in $\mathbb{C}$, for example), then one must choose among the factors of the correct degree, selecting the one that corresponds to the given extension. The details are left to the reader (Exercise 19).

2.1.6 Compositum with Normal Extensions

We keep the situation of the preceding section, but we will specialize to $K = \mathbb{Q}$ (although most of what will be said applies with essentially no change to the general case), so we change notation.

Let $K = \mathbb{Q}(\theta_1)$ be defined by a root of an irreducible polynomial $T_1(X) \in \mathbb{Q}[X]$ of degree n_1, and let $L = K(\theta_2)$ be a relative extension, defined by a root θ_2 of the polynomial $T_2 \in K[X]$, irreducible over K of degree $n_2 = [L : K]$. Recall that a special case of this situation is the compositum of two number fields.

In the preceding section, we gave an algorithm for computing an absolute defining polynomial $T(X)$ for $L/\mathbb{Q}$ and for expressing θ_1 and θ_2 in terms of the generic root θ of T.

We will want to compute arithmetical invariants of L such as its discriminant and integral basis, or to perform polynomial reduction on its defining polynomial, either relative or absolute. Later we will see how to use the relative defining polynomial (which, as usual, gives simpler results). If instead we want to use the absolute defining polynomial alone, the first major problem is to *factor* its discriminant. This is generally quite a hard task (with the technology available in 1999, factoring 100-digit numbers is already quite hard; the records, harnessing huge amounts of computing power, are in the 170-digit range). In the special case where L is obtained as a compositum with a *normal* extension of K (and also in more general cases), we have algebraic methods for obtaining a partial factorization, which we now explain.

Thus, let L be the compositum of K with a number field K_2 assumed to be *normal* over $\mathbb{Q}$ with Galois group G. This is an important situation that occurs, for example, when we deal with Kummer extensions (see Chapter 5), where we first need to adjoin an nth root of unity, so that $K_2 = \mathbb{Q}(\zeta_n)$ and $G \simeq (\mathbb{Z}/n\mathbb{Z})^*$.

Let $T_2(X)$ be a polynomial defining $K_2/\mathbb{Q}$, and let θ_2 be a root of T_2 in K_2. As we have seen in Section 2.1.3, if we set

$$R(X,Z) = \mathcal{R}_Y(T_1(Y), T_2(X - ZY)) \ ,$$

we can find $k \in \mathbb{Z}$ such that $R(X) = R(X,k)$ is squarefree and defines our étale algebra $L = \mathbb{Q}[\theta]$ with $\theta = \theta_2 + k\theta_1$, and we also have

$$\theta_1 = -\frac{R'_Z}{R'_X}(\theta, k), \quad \theta_2 = \theta - k\theta_1 \ .$$

If we want to work directly with the absolute polynomial $R(X)$, we must begin by factoring its discriminant. Since we have introduced parasitic factors, its discriminant is generally large and hence difficult to factor. Using the algebraic structure present in the construction of $R(X)$, however, we can considerably simplify the factoring process.

To simplify the computations, we will assume that T_1 and T_2 are monic and with integer coefficients, since it is easy to reduce to this case. The reader can easily modify the computations given below to the case where no such preliminary reduction is made (see Exercise 20). Since K_2 is normal over $\mathbb{Q}$, we can index the roots of T_2 by G, so that if $K_2 = \mathbb{Q}(\beta)$, we can set $\beta_\sigma = \sigma(\beta)$ for any $\sigma \in G = \mathrm{Gal}(K_2/\mathbb{Q})$, and these will be all the roots of T_2. Thus, we can write

$$T_1(X) = \prod_{0 \le i < n_1} (X - \alpha_i) \quad \text{and} \quad T_2(X) = \prod_{\sigma \in G} (X - \beta_\sigma) \ ,$$

and we let $n = n_1 n_2$ be the absolute degree of L.

The definition of the resultant (see [Coh0, Definition 3.3.2]) shows that

$$R(X,Z) = \prod_{0\le i<n_1,\sigma\in G} (X - (\alpha_i Z + \beta_\sigma)) \ ,$$

hence the formula for the discriminant ([Coh0, Proposition 3.3.5]) gives

$$\begin{aligned} D(Z) &= \mathrm{disc}_X(R(X,Z)) \\ &= (-1)^{n(n-1)/2} \prod_{(i_2,\sigma_2)\neq(i_1,\sigma_1)} ((\alpha_{i_2} - \alpha_{i_1})Z + (\beta_{\sigma_2} - \beta_{\sigma_1})) \\ &= (-1)^{n(n-1)/2} P_1(Z)\cdot P_2 \ , \end{aligned}$$

with

$$P_1(Z) = \prod_{i_2\neq i_1} \prod_{\sigma_1,\sigma_2} ((\alpha_{i_2} - \alpha_{i_1})Z + (\beta_{\sigma_2} - \beta_{\sigma_1}))$$

and

$$P_2 = \prod_{i_2=i_1} \prod_{\sigma_2\neq\sigma_1} (\beta_{\sigma_2} - \beta_{\sigma_1}) = (-1)^{n_1 n_2(n_2-1)/2}\,\mathrm{disc}(T_2)^{n_1} \ ,$$

where, of course, i_1 and i_2 vary implicitly between 0 and $n_1 - 1$, while σ_1 and σ_2 are in G. Hence,

$$D(Z) = (-1)^{n(n+n_2-2)/2}\,\mathrm{disc}(T_2)^{n_1} P_1(Z) \ .$$

The discriminant of the absolute polynomial $R(X) = R(X,k)$ defining L is thus equal to $D(k)$. We will see below that up to sign $P_1(k)$ is a *square*, but it is simpler to keep it in the present form for now.

Let us compute this value in the simplest possible rational terms. Here we will use in an essential way the fact that the number field K_2 is a normal extension of $\mathbb{Q}$. Grouping terms with a given $s = \sigma_1^{-1}\sigma_2$, we have

$$D(k) = (-1)^{n(n+n_2-2)/2}\,\mathrm{disc}(T_2)^{n_1} \prod_{s\in G} D_s(k) \ ,$$

with

$$D_s(k) = \prod_{\sigma\in G} \prod_{0\le i_2\neq i_1<n_1} ((\alpha_{i_2} - \alpha_{i_1})k + (\beta_{\sigma s} - \beta_\sigma)) \ .$$

By Galois theory, since $D_s(k)$ is invariant both by the Galois group of the Galois closure of $K/\mathbb{Q}$ and by G, $D_s(k)$ is a rational integer. In addition, if we denote by 1_G the unit element of G, then

$$\begin{aligned} D_{1_G}(k) &= k^{n_1 n_2(n_1-1)} \Big(\prod_{0\le i_2\neq i_1<n_1} (\alpha_{i_2} - \alpha_{i_1}) \Big)^{n_2} \\ &= k^{n_1 n_2(n_1-1)} (-1)^{n_1 n_2(n_1-1)/2}\,\mathrm{disc}(T_1)^{n_2} \ ; \end{aligned}$$

hence

$$D(k) = (-1)^a k^{n_1 n_2 (n_1-1)} \operatorname{disc}(T_1)^{n_2} \operatorname{disc}(T_2)^{n_1} \prod_{s \in G, s \neq 1_G} D_s(k) \ ,$$

with

$$a = \frac{n(n+n_2-2)+n(n_1-1)}{2} \equiv \frac{n_1(n_1+1)}{2} n_2(n_2+1) \equiv 0 \pmod 2 \ .$$

Furthermore, it is easy to see that

$$\begin{aligned} D_{s^{-1}}(k) &= \prod_{\sigma \in G} \prod_{0 \leq i_2 \neq i_1 < n_1} ((\alpha_{i_2} - \alpha_{i_1})k + (\beta_{\sigma s^{-1}} - \beta_\sigma)) \\ &= \prod_{\sigma \in G} \prod_{0 \leq i_2 \neq i_1 < n_1} ((\alpha_{i_2} - \alpha_{i_1})k + (\beta_\sigma - \beta_{\sigma s})) \\ &= (-1)^{n_1 n_2 (n_1-1)} D_s(k) = D_s(k) \ . \end{aligned}$$

In addition, if $s^2 = 1_G$ and $s \neq 1_G$, let H be a system of right coset representatives of G modulo $< s >$, so that $G = H \cup Hs$ (disjoint union). Then

$$\begin{aligned} D_s(k) &= \prod_{\sigma \in H} \prod_{i_2 \neq i_1} ((\alpha_{i_2} - \alpha_{i_1})k + (\beta_{\sigma s} - \beta_\sigma))((\alpha_{i_2} - \alpha_{i_1})k + (\beta_{\sigma s^2} - \beta_{\sigma s})) \\ &= E_s(k)^2 \ , \end{aligned}$$

with

$$E_s(k) = \prod_{\sigma \in H} \prod_{i_2 \neq i_1} ((\alpha_{i_2} - \alpha_{i_1})k + (\beta_{\sigma s} - \beta_\sigma)) \ .$$

Since $E_s(k)$ is still invariant by the Galois groups, it is rational. Hence it follows that $D_s(k)$ is a square when $s^2 = 1_G$ and $s \neq 1_G$.

Thus we have finally obtained the following result.

Lemma 2.1.13. *With the above notation, we have*

$$D(k) = k^{n_1 n_2 (n_1-1)} \operatorname{disc}(T_1)^{n_2} \operatorname{disc}(T_2)^{n_1} \prod_{s \in G, s \neq 1_G} D_s(k) \ .$$

Furthermore, for all s, we have $D_{s^{-1}} = D_s$, and if $s^2 = 1_G$ and $s \neq 1_G$, then D_s is the square of a rational integer.

Hence, we have split our large discriminant $D(k)$ as a product of smaller pieces $D_s(k)$. This already shows that $D(k)$ must factor relatively easily. This is still theoretical, however, since we must also give a purely algebraic way of computing $D_s(k)$.

To do this, we make the following observation. Let

$$U(X) = \mathcal{R}_Y(T_1(Y), T_1(Y+X))/X^{n_1}$$

be the resultant in Y of T_1 with a shifted version of the *same* polynomial T_1 divided by X^{n_1}. Then

$$U(X) = \prod_{0 \le i_2 \ne i_1 < n_1} (\alpha_{i_2} - \alpha_{i_1} + X) \ .$$

Hence for $s \ne 1_G$,

$$\begin{aligned} D_s(k) &= \prod_{\sigma \in G} \left(k^{n_1(n_1-1)} U((\beta_{\sigma s} - \beta_\sigma)/k) \right) \\ &= k^{n_1 n_2 (n_1-1)} \prod_{\sigma \in G} U((\beta_{\sigma s} - \beta_\sigma)/k) \ . \end{aligned}$$

If we set

$$V_s(X) = \prod_{\sigma \in G} (X - (\beta_{\sigma s} - \beta_\sigma)) \ ,$$

we have

$$V_s(X) = \prod_{\sigma \in G} (X - \sigma(s(\beta) - \beta)) = C_{s(\beta)-\beta}(X) \ ,$$

where $C_\alpha(X)$ denotes the characteristic polynomial of α in the number field K_2 (see [Coh0, Definition 4.3.1]). Since $K_2/\mathbb{Q}$ is a normal extension, $s(\beta)$ is a polynomial in β with rational coefficients, and hence we can set $s(\beta) = A_s(\beta)$ with $A_s \in \mathbb{Q}[X]$. Note that A_s can be computed algorithmically using one of the algorithms for the field isomorphism problem ([Coh0, Section 4.5]). Thus, using [Coh0, Proposition 4.3.4], we have

$$V_s(X) = \mathcal{R}_Y(T_2(Y), X + Y - A_s(Y)) \ .$$

Finally, coming back to $D_s(k)$, we see that

$$\mathcal{R}_X(U(X), V_s(kX)) = k^{n_1 n_2 (n_1-1)} \prod_{\sigma \in G} U((\beta_{\sigma s} - \beta_\sigma)/k) = D_s(k) \ .$$

We summarize what we have obtained in the following theorem.

Theorem 2.1.14. *Let $K_1 = \mathbb{Q}(\theta_1)$ and $K_2 = \mathbb{Q}(\theta_2)$ be number fields of respective degrees n_1 and n_2, and let $T_1(X)$ and $T_2(X)$ be the minimal monic polynomials of θ_1 and θ_2, respectively. Assume that K_2 is a normal extension of $\mathbb{Q}$ with Galois group G. Let $R(X) = R(X,k)$ be an absolute defining polynomial for the compositum L of K_1 and K_2 as computed by Algorithm 2.1.8 (R is squarefree but not necessarily irreducible).*

For $s \in G$, $s \ne 1_G$, define $A_s(X)$ to be the polynomial expressing $s(\theta_2)$ in terms of θ_2, and set $V_s(X) = \mathcal{R}_Y(T_2(Y), X + Y - A_s(Y))$ (this depends only on the number field K_2 and on s).

Let $U(X) = \mathcal{R}_Y(T_1(Y), T_1(Y+X))/X^{n_1}$ (this depends only on the number field K_1), and for $s \ne 1_G$, set

$$D_s(k) = \mathcal{R}_X(U(X), V_s(kX)) \ .$$

Then

(1) *for all* $s \in G$, $s \neq 1_G$, *we have* $D_s(k) \in \mathbb{Z}$;
(2) *we have the decomposition*

$$\operatorname{disc}(R(X)) = k^{n_1 n_2 (n_1 - 1)} \operatorname{disc}(T_1)^{n_2} \operatorname{disc}(T_2)^{n_1} \prod_{s \in G, s \neq 1_G} D_s(k) \ ;$$

(3) *for all* $s \in G$, *we have* $D_{s^{-1}}(k) = D_s(k)$;
(4) *if* $s^2 = 1_G$ *and* $s \neq 1_G$, *then* $D_s(k)$ *is the square of a rational integer.*

Remarks

(1) To use this theorem in practice, we let I be the set of elements s of G such that $s^2 = 1_G$ and $s \neq 1_G$, and we let G_1 be a complete set of representatives for the equivalence relation on $G - I - \{1_G\}$ whose equivalence classes are the pairs $\{s, s^{-1}\}$. Then

$$\operatorname{disc}(R(X)) = k^{n_1 n_2 (n_1 - 1)} \operatorname{disc}(T_1)^{n_2} \operatorname{disc}(T_2)^{n_1} \Big(\prod_{s \in G_1} D_s(k) \prod_{s \in I} E_s(k) \Big)^2 \ ,$$

with $E_s(k) = D_s(k)^{1/2} \in \mathbb{Z}$ as above.
(2) If instead of choosing $\theta = \theta_2 + k\theta_1$ we choose $\theta = \theta_1\theta_2$ (or, more generally, $\theta = \theta_1\theta_2 + k_1\theta_1 + k_2\theta_2$) as in Algorithm 2.1.9, a completely similar theorem holds; see Exercise 9.

Let us look at an example. Let $T_1(X) = X^4 - X^3 + 2X + 1$ and let $T_2(X) = (X^{11} - 1)/(X - 1)$ be the 11th cyclotomic polynomial, which defines a cyclic extension of $\mathbb{Q}$. An absolute defining polynomial for the compositum, which is of degree 40, obtained by choosing $\theta = \theta_1 + \theta_2$, has a discriminant of several hundred digits. Even after casting out small prime factors less than 500,000, say, and noting that the unfactored part is a square, the number that remains to be factored still has 110 digits. Factoring such a number is in general a feasible but formidable task. We know, however, that it must factor in relatively small parts $D_s(k)$, and indeed all the $D_s(k)$ have around 34 digits, which are considerably easier numbers to factor.

If, on the other hand, we use $\theta = \theta_1\theta_2$, we are in a very favorable case. First, the discriminant of the compositum obtained in this way has less than half the digits of the preceding one. Second, it is divisible only by very small primes (the largest being 1319), so factoring becomes trivial. Even if it were not so, we could have used the analog of Theorem 2.1.14 given in Exercise 9.

The method explained above can be generalized to the case where there exist only some nontrivial $\mathbb{Q}$-automorphisms of K_2, and also to the case of relative normal extensions of K_1 not necessarily defined by a compositum. The methods are completely similar, and the details are left to the reader.

2.2 Arithmetic of Relative Extensions

The preceding section was essentially field-theoretical. In the present section, which is mainly theoretical, we study the arithmetic of relative extensions, in particular, the properties of the rings of algebraic integers.

We first briefly explain how the usual notions for absolute extensions extend to the relative case. We follow closely [Coh0, Section 4.1.2 and following]. Let K be a base field, $L = K(\theta)$ a relative extension, and $T(X) \in K[X]$ the minimal monic polynomial of θ which is irreducible in $K[X]$. More generally, we could assume that L is only an *étale algebra* over K, in other words, that $T(X)$ is only squarefree.

2.2.1 Relative Signatures

We start with the following simple, but important, theorem.

Theorem 2.2.1. *Let $L = K(\theta)$ be an extension of number fields with $T(\theta) = 0$ as above, and let $n = \deg(T) = [L : K]$. Let σ be an embedding (an* injective *field homomorphism) of K into an arbitrary field Ω (not necessarily a number field). Assume that the polynomial $T^\sigma(X)$ has n roots in Ω, where T^σ denotes the polynomial obtained from T by applying σ to all the coefficients. Then σ can be extended to exactly n embeddings of L into Ω.*

Proof. Indeed, let $\alpha = A(\theta) \in L = K(\theta)$. If ϕ is an extension of σ to L, we must have $\phi(\alpha) = \phi(A(\theta)) = A^\sigma(\phi(\theta))$. Since $T(\theta) = 0$, we must have $T^\sigma(\phi(\theta)) = 0$, hence $\phi(\theta)$ must be one of the n roots β_i of T^σ in Ω, so there are at most n embeddings. Conversely, if we set $\phi(\alpha) = A^\sigma(\beta_i)$ and if $\alpha = A_1(\theta)$ for some other polynomial A_1, we have $A_1(X) - A(X) = T(X)U(X)$, hence

$$A_1^\sigma(\beta_i) = A^\sigma(\beta_i) + T^\sigma(\beta_i)U^\sigma(\beta_i) = A^\sigma(\beta_i) \ ,$$

so ϕ is a well-defined embedding of L extending σ. □

Corollary 2.2.2. *Let k be a number field (for example, $k = \mathbb{Q}$), K and K' two extensions of k, and L/K an extension of K of relative degree n. We assume that all our number fields are subfields of $\mathbb{C}$. Any k-isomorphism of σ from K to K' extends to exactly n k-embeddings from L into $\mathbb{C}$.*

The following proposition gives a more precise way of stating these results.

Proposition 2.2.3. *Let $L = K(\theta)$ be a relative extension of number fields, and let T be the minimal monic polynomial of θ, as above. Let $m = [K : \mathbb{Q}]$ and $n = [L : K]$, so that $[L : \mathbb{Q}] = nm$. For each of the m embeddings τ_i of K into $\mathbb{C}$, denote by T^{τ_i} the polynomial obtained from T by applying τ_i on the coefficients. Then we have the following.*

(1) *Each of the m embeddings τ_i of K into $\mathbb{C}$ extends to exactly n embeddings of L into $\mathbb{C}$, given by $\theta \mapsto \theta_{i,j}$, where the $\theta_{i,j}$ are the roots of T^{τ_i} for $1 \leq j \leq n$.*
(2) *There exist exactly n embeddings $\sigma_{j,K}$ of L into $\mathbb{C}$ which are K-linear, given by $\theta \mapsto \theta_j$, where the θ_j are the roots of the polynomial T in $\mathbb{C}$.*

Proof. If σ is an embedding of L into $\mathbb{C}$, then $\sigma|_K$ is an embedding of K into $\mathbb{C}$, so $\sigma|_K = \tau_i$ for a certain i. If $\sigma(\theta) = \theta'$, applying σ to the equality $T(\theta) = 0$ we obtain $T^{\tau_i}(\theta') = 0$, and hence $\theta' = \theta_{i,j}$ for a certain index j. Conversely, it is clear that

$$\sigma\left(\sum_k a_k \theta^k\right) = \sum_k \tau_i(a_k)\theta_{i,j}^k$$

defines an embedding of L into $\mathbb{C}$, which extends τ_i.

This embedding will be K-linear if and only if $\tau_i(a) = a$ for all $a \in K$ — in other words, if τ_i is the identity map (recall that we have explicitly embedded K into $\overline{\mathbb{Q}} \subset \mathbb{C}$) — hence there are exactly n K-linear embeddings, namely those that extend the identity of K. □

Definition 2.2.4. *Let L/K be a relative extension of relative degree n, and let τ be an embedding of K into $\mathbb{C}$. We say that τ is* ramified *in L if τ is a real embedding (that is, if $\tau(K) \subset \mathbb{R}$) and if at least one of the extensions of τ to L is not a real embedding. It is* unramified *otherwise (in particular, a nonreal embedding is unramified).*

In terms of defining polynomials, if $T(X) \in K[X]$ is a polynomial defining the field L over K, then a real embedding τ is unramified if and only if T^τ has only real roots in $\mathbb{C}$.

The following is a simple consequence of this definition.

Proposition 2.2.5. *Let L/K be a relative extension of relative degree n. Denote by (r_1, r_2) (resp., (R_1, R_2)) the signature of the number field K (resp., L). If all the embeddings τ of K are unramified in L, we have $(R_1, R_2) = (nr_1, nr_2)$.*

More generally, if $R_{1,i}$ (resp., $2R_{2,i}$) is the number of real (resp., nonreal) roots of T^{τ_i} for $1 \leq i \leq r_1$, we have the formula

$$(R_1, R_2) = \left(\sum_{1 \leq i \leq r_1} R_{1,i}, \quad nr_2 + \sum_{1 \leq i \leq r_1} R_{2,i}\right).$$

Proof. If τ is a nonreal embedding of K, any extension of τ to L must also be nonreal since L is an extension of K. On the other hand, if $\tau = \tau_i$ is a real embedding, the polynomial T^{τ_i} has $R_{1,i}$ real and $2R_{2,i}$ nonreal roots for some nonnegative integers $R_{1,i}$ and $R_{2,i}$ such that $R_{1,i} + 2R_{2,i} = n$. Hence the signature of L is equal to (R_1, R_2) with

$$R_1 = \sum_{1 \le i \le r_1} R_{1,i} \quad \text{and} \quad R_2 = \sum_{1 \le i \le r_1} R_{2,i} + nr_2 \ ,$$

as claimed. In the special case where all the τ_i are unramified, we have $R_{1,i} = n$ and $R_{2,i} = 0$, proving the formulas of the proposition. □

Note that when τ is a nonreal embedding, the polynomial T^τ does not necessarily have an even number of nonreal roots, since it is not invariant by complex conjugation.

Recall that an extension L/K of number fields is Galois (or normal) if L is globally invariant by the $[L : K]$ K-linear embeddings of L into $\mathbb{C}$. If this is the case, the set of such embeddings is a group, called the *Galois group* of L/K and denoted by $\mathrm{Gal}(L/K)$ (remember that all our number fields are assumed to be subfields of $\mathbb{C}$). If an element $x \in L$ is such that $g(x) = x$ for all $g \in \mathrm{Gal}(L/K)$, then Galois theory tells us that $x \in K$.

If L/K is a Galois extension, then Proposition 2.2.5 simplifies considerably, as follows.

Corollary 2.2.6. *Keep the notation of Proposition 2.2.5, and assume in addition that L/K is a Galois extension.*

(1) *If k is the number of ramified real places of K in L/K, we have $(R_1, R_2) = (n(r_1 - k), n(r_2 + k/2))$.*
(2) *If n is odd, we have $k = 0$, so $(R_1, R_2) = (nr_1, nr_2)$.*

Proof. Let τ be a real embedding of K. If τ has a real extension to L, then since L/K is Galois, the roots of the defining polynomial T^τ can be expressed as polynomials with coefficients in $\tau(K)$ of any one of them. Hence if one root is real, all of them are, and if one is nonreal, all of them are. Thus, either τ is unramified, or *all* the extensions of τ to L are nonreal. In the case where n is odd, T^τ is an odd-degree polynomial with real coefficients — hence has at least one real root — so all real places τ are unramified, thus proving the corollary. □

Corollary 2.2.7. *Keep the notation of Proposition 2.2.5. Then*

(1) *we have $R_1 + 2R_2 = n(r_1 + 2r_2)$ and $R_1 \le nr_1$ (or, equivalently, $R_2 \ge nr_2$), and if n is odd, $R_1 \ge r_1$;*
(2) *if L/K is a Galois extension, then in addition $n \mid R_1$ (or, equivalently, $n \mid 2R_2$), and if n is odd, $R_1 = nr_1$ and $R_2 = nr_2$;*
(3) *conversely, if (1) is satisfied, there exists a relative extension L/K of signature (R_1, R_2), and if (1) and (2) are satisfied, there exists a Galois (even a cyclic) extension L/K of signature (R_1, R_2).*

Proof. Statement (1) immediately follows from Proposition 2.2.5, and (2) follows from Corollary 2.2.6. The cases n odd follow from the fact that a real polynomial of odd degree has at least one real root.

Conversely, assume (1). Since $R_1 \leq nr_1$ and $R_1 \equiv nr_1 \pmod 2$, we can find r_1 integers $n_1, \ldots, n_{r_1}$ such that $n_i \leq n$, $n_i \equiv n \pmod 2$, and $\sum_{1 \leq i \leq r_1} n_i = R_1$. For each $i \leq r_1$, choose a monic squarefree polynomial $P_i(X) \in \mathbb{R}[X]$ of degree n having exactly n_i real roots; for example,

$$P_i(X) = \prod_{1 \leq j \leq n_i} (X - j) \prod_{1 \leq j \leq (n-n_i)/2} (X^2 + j^2) \ .$$

Write $P_i(X) = \sum_{0 \leq j \leq n} a_{i,j} X^j$, and let ε be a sufficiently small, positive real number. By the approximation theorem (Proposition 1.2.8), we can find $\alpha_j \in K$ such that $\left|\alpha_j^{(i)} - a_{i,j}\right| \leq \varepsilon$ for all $i \leq r_1$, where as usual $\alpha_j^{(i)}$ denotes the ith conjugate of α_j. Since P_i is squarefree, if ε is small enough, by continuity the modified polynomials $Q_i(X) = \sum_{0 \leq j \leq n} \alpha_j^{(i)} X^j$ will have the same number of real roots as P_i, in other words n_i. By the approximation theorem once again, we may also modify α_j so that $Q_i(X)$ is irreducible in $K[X]$. Once this is done, we take $Q(X) = \sum_{0 \leq j \leq n} \alpha_j X^j$ and it is clear that a root of $Q(X)$ defines an extension L/K having the required signature.

Assume now that (1) and (2) are satisfied. Choose a large prime $p > 2$ such that $p \equiv 1 \pmod{2n}$ and $p \nmid d(K)$. Since $\mathbb{Q}(\zeta_p)/\mathbb{Q}$ is ramified only at p and $K/\mathbb{Q}$ is not, it follows that $K \cap \mathbb{Q}(\zeta_p) = \mathbb{Q}$. Let $\eta = \zeta_p + \zeta_p^{-1}$, so that $\mathbb{Q}(\eta)$ is the totally real subfield of degree $(p-1)/2$ of $\mathbb{Q}$. Denote by $k = \mathbb{Q}(\theta)$ the unique totally real subfield of degree n of $\mathbb{Q}(\eta)$, which exists since $\mathrm{Gal}(\mathbb{Q}(\eta)/\mathbb{Q}) \simeq (\mathbb{Z}/((p-1)/2)\mathbb{Z})$ and $p \equiv 1 \pmod{2n}$. Since $K \cap \mathbb{Q}(\zeta_p) = \mathbb{Q}$, we also have $K \cap k = \mathbb{Q}$, so $K(\theta)/K$ is a cyclic extension of degree exactly equal to n, and by construction all the conjugates of θ over any real place of K are real. Note, in addition, that the discriminant of $K(\theta)/K$ will be divisible only by prime ideals above p, since this is the case for $K(\zeta_p)$.

If n is odd, then by (2) the signature of $K(\theta)$ is equal to (nr_1, nr_2), so $L = K(\theta)$ is a field with suitable signature. Thus assume n even. Let $G = \mathrm{Gal}(K(\theta)/K)$, let σ be a generator of G, and set $s = \sigma^{n/2}$, which is thus an element of order 2 in G.

Let α be any element of $K^* \smallsetminus K^{*2}$ having zero $\mathfrak{p}$-adic valuation for all prime ideals $\mathfrak{p}$ above p. It follows in particular that $K\left(\sqrt{\alpha}\right)$ is linearly disjoint from $K(\theta)$ over K. Consider the field $L = K\left((\theta - s(\theta))\sqrt{\alpha}\right)$. I claim that for a suitable choice of α the field L will have the desired properties.

First, since $K\left(\sqrt{\alpha}\right)$ is linearly disjoint from $K(\theta)$ over K, L is the compositum of $K(\theta - s(\theta))$ with $K\left(\sqrt{\alpha}\right)$. In addition, the Galois conjugates of $u = (\theta - s(\theta))\sqrt{\alpha}$ are $u_i = (\sigma^i(\theta) - s(\sigma^i(\theta)))\sqrt{\alpha}$ for $0 \leq i < n$, since changing $\sqrt{\alpha}$ into $-\sqrt{\alpha}$ is equivalent to changing i into $i + n/2$ modulo n. If we choose θ to be a normal basis of $K(\theta)/K$, the u_i are distinct and hence L/K is a cyclic extension of degree n.

Next, let σ_i be a real embedding of K. Since all the embeddings of θ above a real embedding of K are real by assumption, it follows that σ_i is ramified in L/K if and only if $\sigma_i(\alpha) < 0$. Thus, if we choose α so that $\sigma_i(\alpha) < 0$ for R_1/n

of the real embeddings of K and $\sigma_i(\alpha) > 0$ for the others, the number of real embeddings of L will be equal to R_1, as desired. (I thank Bjorn Poonen for the idea leading to this last proof.) □

2.2.2 Relative Norm, Trace, and Characteristic Polynomial

If we denote by $\sigma_{i,K}$ the n K-linear embeddings of L into $\mathbb{C}$, then for $\alpha \in L$ we define the (relative) characteristic polynomial $C_\alpha(X)$ of α by

$$C_\alpha(X) = \prod_{1 \le i \le n} (X - \sigma_{i,K}(\alpha)) \ ,$$

which belongs to $K[X]$ by Galois theory. If

$$C_\alpha(X) = \sum_{0 \le i \le n} (-1)^{n-i} s_{n-i}(\alpha) X^i \ ,$$

then $s_1(\alpha)$ is called the *relative trace* of α and denoted $\mathrm{Tr}_{L/K}(\alpha)$, and $s_n(\alpha)$ is called the *relative norm* of α and denoted $\mathcal{N}_{L/K}(\alpha)$. Evidently the trace is additive and the norm is multiplicative. Note that similar statements are not true for the other coefficients of the characteristic polynomial, which explains the importance of these functions.

Furthermore, Proposition 2.2.3 implies immediately that the relative trace and norm are *transitive*, in other words, that they satisfy

$$\mathrm{Tr}_{L/\mathbb{Q}}(\alpha) = \mathrm{Tr}_{K/\mathbb{Q}}(\mathrm{Tr}_{L/K}(\alpha)) \quad \text{and} \quad \mathcal{N}_{L/\mathbb{Q}}(\alpha) = \mathcal{N}_{K/\mathbb{Q}}(\mathcal{N}_{L/K}(\alpha)) \ .$$

More generally, the characteristic polynomial itself (and hence all of its coefficients) satisfies the transitivity property (see Exercise 21). This allows an absolute characteristic polynomial to be computed from a relative one.

As in the absolute case, a characteristic polynomial can be computed using resultants. Let $L = K(\theta)$. Then, if T is the minimal monic polynomial of θ, and if $\alpha = A(\theta)$ for some polynomial $A \in K[X]$, then $C_\alpha(X) = \mathcal{R}_Y(T(Y), X - A(Y))$, where $\mathcal{R}_Y$ denotes the resultant with respect to Y. In particular, we have $\mathcal{N}_{L/K}(\alpha) = \mathcal{R}(T(X), A(X))$.

2.2.3 Integral Pseudo-Bases

We now explain in more detail how to generalize the notions of integral basis and discriminant. This is a little less straightforward than for the preceding notions, and it uses most of the ideas of Chapter 1.

As usual, let L/K be a relative extension of degree n. The ring of integers $\mathbb{Z}_L$ of L is not only a finitely generated free $\mathbb{Z}$-module but is clearly also a finitely generated $\mathbb{Z}_K$-module. The ring $\mathbb{Z}_K$ is, however, in general not a

principal ideal domain but only a Dedekind domain, and hence $\mathbb{Z}_L$ is not necessarily free (see Exercise 22 for an example). The theory developed in Chapter 1 tells us that $\mathbb{Z}_L$ has a pseudo-basis over $\mathbb{Z}_K$, and any such basis will be called a *relative integral pseudo-basis*, or simply an *integral pseudo-basis*.

If we assume that L is given by $L = K(\theta)$ with θ an algebraic *integer*, then $\mathbb{Z}_K[\theta] \subset \mathbb{Z}_L$. Hence, as in the absolute case this implies that the relative HNF pseudo-basis of $\mathbb{Z}_L$ in the K-basis $(1, \theta, \ldots, \theta^{n-1})$ must satisfy some conditions, as follows.

Proposition 2.2.8. *Let M be a $\mathbb{Z}_K[\theta]$-module that is projective of rank n as a $\mathbb{Z}_K$-module, and let $(\omega_i, \mathfrak{a}_i)_{1 \le i \le n}$ be a pseudo-basis of M in relative HNF on the basis $(1, \theta, \ldots, \theta^{n-1})$, where θ is assumed to be an algebraic integer.*

(1) *The ideals $\mathfrak{q}_i = \mathfrak{a}_i^{-1}$ are divisible by $\mathfrak{q}_1$, and we have*

$$\mathfrak{a}_1 \subset \mathfrak{a}_2 \subset \cdots \subset \mathfrak{a}_n \ ;$$

in other words,

$$\mathfrak{q}_1 \mid \mathfrak{q}_2 \mid \cdots \mid \mathfrak{q}_n \ .$$

(2) *For all $i \le n$ we have $\omega_i \in \mathbb{Z}_K[\theta]$; in other words, if $(H, (\mathfrak{a}_i))$ is the pseudo-matrix representing the pseudo-basis (ω_i, α_i), then the entries of H are in $\mathbb{Z}_K$.*

Proof. We will prove (1) and (2) simultaneously by showing by induction on j that $\omega_j \in \mathbb{Z}_K[\theta]$ and $\mathfrak{a}_{j-1} \subset \mathfrak{a}_j$ for $j > 1$. Since $\omega_1 = 1$, this is trivially true for $j = 1$. Assume that it is true up to $j - 1$, and let a be any element of $\mathfrak{a}_{j-1}$. Since M is a $\mathbb{Z}_K[\theta]$-module, we have $\mathfrak{a}_{j-1}\theta\omega_{j-1} \subset M$; hence in particular,

$$a\theta\omega_{j-1} = \sum_{1 \le i \le n} x_i\omega_i \quad \text{with} \quad x_i \in \mathfrak{a}_i \ .$$

Since the matrix of the ω_i is upper-triangular with 1 on the diagonal, we obtain $x_i = 0$ for $i > j$ and $x_j = a$. Since this is true for any $a \in \mathfrak{a}_{j-1}$, we therefore have $\mathfrak{a}_{j-1} \subset \mathfrak{a}_j$. In addition,

$$a\omega_j = a\theta\omega_{j-1} - \sum_{1 \le i < j} x_i\omega_i \ .$$

Since $x_i \in \mathfrak{a}_i$, and by induction we have $\mathfrak{a}_1 \subset \cdots \subset \mathfrak{a}_{j-1}$ and $\omega_i \in \mathbb{Z}_K[\theta]$ for $i < j$, we have for all $a \in \mathfrak{a}_{j-1}$

$$a(\omega_j - \theta\omega_{j-1}) \in \mathfrak{a}_{j-1}\mathbb{Z}_K[\theta] \ ,$$

Hence

$$\mathfrak{a}_{j-1}(\omega_j - \theta\omega_{j-1}) \subset \mathfrak{a}_{j-1}\mathbb{Z}_K[\theta] \ ,$$

from which we deduce that $\omega_j - \theta\omega_{j-1} \in \mathbb{Z}_K[\theta]$, hence that $\omega_j \in \mathbb{Z}_K[\theta]$, proving our induction hypothesis. Thus $\mathfrak{a}_1 \subset \mathfrak{a}_2 \subset \cdots \subset \mathfrak{a}_n$; hence by taking inverses, $\mathfrak{q}_1 \supset \mathfrak{q}_2 \supset \cdots \supset \mathfrak{q}_n$, showing that all the $\mathfrak{q}_i$ are divisible by $\mathfrak{q}_1$ and proving the proposition. □

Corollary 2.2.9. *Let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis of $\mathbb{Z}_L$ in HNF on $(1, \theta, \ldots, \theta^{n-1})$, where θ is assumed to be an algebraic integer.*

(1) *The ideals $\mathfrak{q}_i = \mathfrak{a}_i^{-1}$ are integral ideals, $\mathfrak{a}_1 = \mathfrak{q}_1 = \mathbb{Z}_K$, and*

$$\mathbb{Z}_K = \mathfrak{q}_1 \mid \mathfrak{q}_2 \mid \cdots \mid \mathfrak{q}_n \ .$$

(2) *For all $i \leq n$ we have $\omega_i \in \mathbb{Z}_K[\theta]$.*

(3) *For every $i \leq j$, we have*

$$\mathfrak{q}_i\mathfrak{q}_{j+1-i} \mid \mathfrak{q}_j \ .$$

Proof. Since $\mathfrak{a}_1 = \mathbb{Z}_L \cap K = \mathbb{Z}_K$, we have $\mathfrak{a}_1 = \mathfrak{q}_1 = \mathbb{Z}_K$ and (1) and (2) are restatements of Proposition 2.2.8. The proof of (3) is very similar to the proof of the proposition: since the leading term of $\omega_i\omega_{j+1-i}$ is θ^{j-1}, we must have $\mathfrak{a}_i\mathfrak{a}_{j+1-i} \subset \mathfrak{a}_j$, so $\mathfrak{q}_i\mathfrak{q}_{j+1-i} \mid \mathfrak{q}_j$. Note that (3) combined with the fact that the $\mathfrak{q}_i$ are integral implies (1). □

Remark. The above proposition and its corollary are generalizations to the relative case of [Coh0, Theorem 4.7.5 and Corollary 4.7.6]. Using the notation of [Coh0, Corollary 4.7.6], property (3) translates into $d_i d_{j+1-i} \mid d_j$ and is not given in [Coh0].

2.2.4 Discriminants

For $1 \leq i \leq n$, let $\sigma_{i,K}$ be the K-linear embeddings of L into $\mathbb{C}$, and let $\alpha_1, \ldots, \alpha_n$ be n elements of L. Since $\mathrm{Tr}_{L/K}(\alpha) = \sum_{1 \leq i \leq n} \sigma_i(\alpha)$, as in the absolute case we find that

$$\det(\sigma_i(\alpha_j))^2 = \det(\mathrm{Tr}_{L/K}(\alpha_i\alpha_j)) \ .$$

This common quantity belongs to K and will be called the *discriminant* of the α_j and denoted $d(\alpha_1, \ldots, \alpha_n)$. We have $d(\alpha_1, \ldots, \alpha_n) = 0$ if and only if the α_j are K-linearly dependent (see Exercise 23).

Let $\mathcal{B} = (\omega_j, \mathfrak{a}_j)$ be a relative integral pseudo-basis. Then, according to Corollary 1.4.3 and the remarks that follow, we can give two different invariants that together generalize the discriminant in the absolute case. First, the ideal

$$\mathfrak{d}(L/K) = d(\omega_1, \ldots, \omega_n)(\mathfrak{a}_1 \cdots \mathfrak{a}_n)^2$$

which we will call the *relative discriminant ideal* (or simply the *discriminant ideal*) of L/K. Second, the quantity

$$\overline{d(L/K)} = \overline{d(\omega_1, \ldots, \omega_n)} \in K^*/K^{*2}$$

considered as an element of K^*/K^{*2}, in other words, modulo nonzero squares. The pair $\mathrm{disc}(L/K) = (\mathfrak{d}(L/K), \overline{d(L/K)})$ will simply be called the *relative discriminant* of L over K.

As mentioned in Chapter 1, each component of the pair gives (related) information. For example, in the absolute case the discriminant ideal $\mathfrak{d}(L/K)$ gives the absolute value of the discriminant, while $\overline{d(L/K)}$ gives its sign, along with other information.

If $L = K(\theta)$ and θ is chosen to be an algebraic integer, the minimal monic polynomial T of θ will have coefficients in $\mathbb{Z}_K$. Thus $\mathbb{Z}_K[\theta] \subset \mathbb{Z}_L$, and we can consider the module $M = \mathbb{Z}_L/\mathbb{Z}_K[\theta]$. Since $\mathbb{Z}_L$ and $\mathbb{Z}_K[\theta]$ are both of $\mathbb{Z}_K$-rank equal to $n = [L : K]$, it follows that M is a torsion $\mathbb{Z}_K$-module. The order-ideal of M (in the sense of Definition 1.2.33) will be called the *index-ideal* of $\mathbb{Z}_K[\theta]$ (or, by abuse of language, of θ) in $\mathbb{Z}_L$ and denoted by $\mathfrak{f}$.

As in the absolute case, we have $d(1, \theta, \ldots, \theta^{n-1}) = \mathrm{disc}(T)$, where $\mathrm{disc}(T)$ is the discriminant of the polynomial T, and we have the formula

$$\mathrm{disc}(T)\mathbb{Z}_K = \mathfrak{d}(L/K)\mathfrak{f}^2 \ ,$$

where $\mathfrak{d}(L/K)$ is the relative discriminant ideal of L as defined above. It is clear that in K^*/K^{*2} we have $\overline{d(L/K)} = \overline{\mathrm{disc}(T)}$.

Note also that if $(\omega_i, \mathfrak{a}_i)$ is an integral pseudo-basis in HNF, then the matrix of the ω_j on the θ^{i-1} has determinant 1; hence

$$d(\omega_1, \ldots, \omega_n) = d(1, \theta, \ldots, \theta^{n-1}) = \mathrm{disc}(T) \ ,$$

so that

$$\mathfrak{d}(L/K) = \mathrm{disc}(T)(\mathfrak{a}_1 \cdots \mathfrak{a}_n)^2 \ .$$

Using the notation of Corollary 2.2.9, it follows that the index-ideal $\mathfrak{f}$ is given by

$$\mathfrak{f} = (\mathfrak{a}_1 \cdots \mathfrak{a}_n)^{-1} = \mathfrak{q}_1 \cdots \mathfrak{q}_n \ .$$

In Section 2.4, we give an algorithm for computing relative integral pseudo-bases and relative discriminants.

One of the main reasons for introducing the discriminant, in both the absolute case and the relative case, is that it is an *invariant* of the number field, more precisely of its ring of integers. It should be noted, however, that we can define *finer* invariants, although it seems that they have not been used in the literature. The invariance of the discriminant (or the discriminant ideal in the relative case) comes from the invariance of the determinant of a bilinear form by change of basis. The determinant is equal to the product of the elementary divisors of the Smith normal form however, and each of these is also an invariant. More precisely, we have the following proposition.

Proposition 2.2.10. *Let $\mathcal{B} = (\omega_j, \mathfrak{a}_j)$ be a relative pseudo-basis. Let $T = (\mathrm{Tr}_{L/K}(\omega_i\omega_j))$, $I = (\mathfrak{a}_1^{-1}, \ldots, \mathfrak{a}_n^{-1})$, and $J = (\mathfrak{a}_1, \ldots, \mathfrak{a}_n)$. Then (T, I, J) is an integral pseudo-matrix. For $1 \leq i \leq n$, let $\mathfrak{d}_i$ be the elementary divisors of this pseudo-matrix in the sense of Theorem 1.7.2. The $\mathfrak{d}_i$ are independent of the chosen pseudo-basis $\mathcal{B}$ (hence are invariants of the field extension), and their product is equal to the relative discriminant ideal.*

Proof. The proof is straightforward and left to the reader (see Exercises 24 and 25). □

It is natural to call these ideals $\mathfrak{d}_i$ the *elementary discriminantal divisors* of the field extension. We have stated the above proposition in the relative case, but evidently it also gives nontrivial invariants in the absolute case.

2.2.5 Norms of Ideals in Relative Extensions

As usual, let L/K be a relative extension of number fields, and let I be a nonzero integral ideal of $\mathbb{Z}_L$. The absolute norm of I is the order of $\mathbb{Z}_L/I$, but, as above in the case of $\mathbb{Z}_L$, we have a richer structure since $\mathbb{Z}_L/I$ is a $\mathbb{Z}_K$-torsion module.

Definition 2.2.11. *Let L/K be a relative extension, and let I be an integral ideal of $\mathbb{Z}_L$. The* relative norm *of I is the order-ideal of the $\mathbb{Z}_K$-torsion module $\mathbb{Z}_L/I$, or the index-ideal $[\mathbb{Z}_L : I]$ in the sense of Definition 1.2.33. It is an ideal of $\mathbb{Z}_K$ denoted $\mathcal{N}_{L/K}(I)$. In other words, if $\mathbb{Z}_L/I = \bigoplus_i (\mathbb{Z}_K/\mathfrak{d}_i)\overline{\alpha_i}$ as a torsion $\mathbb{Z}_K$-module, then $\mathcal{N}_{L/K}(I) = \prod_i \mathfrak{d}_i$.*

Since we can identify integral ideals of $\mathbb{Z}$ with positive integers, the above definition generalizes the usual definition of the norm of an ideal, and thus we can use the same notation. We will later give other equivalent definitions of the norm of an ideal.

Proposition 2.2.12. (1) *If I and J are two integral ideals of $\mathbb{Z}_L$, we have*

$$\mathcal{N}_{L/K}(IJ) = \mathcal{N}_{L/K}(I)\mathcal{N}_{L/K}(J) \ .$$

(2) *We have*

$$\mathcal{N}_{K/\mathbb{Q}}(\mathcal{N}_{L/K}(I)) = \mathcal{N}_{L/\mathbb{Q}}(I) \ .$$

(3) *If $\alpha \in \mathbb{Z}_L$, we have*

$$\mathcal{N}_{L/K}(\alpha\mathbb{Z}_L) = \mathcal{N}_{L/K}(\alpha)\mathbb{Z}_K \ .$$

Proof. The proof of (1) is exactly as in the absolute case (see, for example, [Coh0, Proposition 4.6.8]), replacing the index $[M : N]$ by the index-ideal $[M : N]$ (that is, by the order-ideal of M/N; see Definition 1.2.33).

For (2), write $\mathbb{Z}_L/I = \bigoplus_i (\mathbb{Z}_K/\mathfrak{d}_i)\overline{\alpha_i}$. Then $\mathcal{N}_{L/K}(I) = \prod_i \mathfrak{d}_i$. For each i, write $\mathbb{Z}_K/\mathfrak{d}_i \simeq \bigoplus_j \mathbb{Z}/d_{i,j}\mathbb{Z}$, so that $N_{K/\mathbb{Q}}(\mathfrak{d}_i) = \prod_j d_{i,j}$, and hence by

(1), $N_{K/\mathbb{Q}}(\mathcal{N}_{L/K}(I)) = \prod_i \prod_j d_{i,j}$. On the other hand, we have $\mathbb{Z}_L/I \simeq \bigoplus_{i,j} \mathbb{Z}/d_{i,j}\mathbb{Z}$, so $N_{L/\mathbb{Q}} = \prod_i \prod_j d_{i,j}$ by Proposition 1.2.34, proving (2).

(3). Let $(\omega_j, \mathfrak{a}_j)$ be a relative integral pseudo-basis of $\mathbb{Z}_L$ over $\mathbb{Z}_K$, let $\sigma_{1,K}, \ldots, \sigma_{n,K}$ be the K-linear embeddings of L into $\mathbb{C}$, and let Ω be the matrix defined by

$$\Omega = \begin{pmatrix} \sigma_1(\omega_1) & \ldots & \sigma_1(\omega_n) \\ \vdots & \ddots & \vdots \\ \sigma_n(\omega_1) & \ldots & \sigma_n(\omega_n) \end{pmatrix} .$$

Since $\alpha \in \mathbb{Z}_L$, multiplication by α induces a $\mathbb{Z}_K$-linear map from $\mathbb{Z}_L$ to itself, which can be represented by the matrix M_α expressing the $\alpha\omega_j$ in terms of the ω_i. In other words, we have

$$(\omega_1, \ldots, \omega_n) M_\alpha = \alpha(\omega_1, \ldots, \omega_n) .$$

Applying the σ_i to the above equality, we deduce that

$$\Omega \cdot M_\alpha = \begin{pmatrix} \sigma_1(\alpha) & \ldots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \ldots & \sigma_n(\alpha) \end{pmatrix} \cdot \Omega .$$

Since the matrix of the $\sigma_i(\omega_j)$ is invertible (its square multiplied by the square of the product of the $\mathfrak{a}_i$ is the relative discriminant ideal $\mathfrak{d}(L/K)$), it follows that the characteristic polynomial of M_α is the same as the characteristic polynomial $C_\alpha(X)$ of α; in particular, the relative norm $\mathcal{N}_{L/K}(\alpha)$ is equal to the determinant of M_α.

Let $I = (\mathfrak{a}_i)$ be the list of ideals in the integral pseudo-basis. Since $\alpha \in \mathbb{Z}_L$, the pseudo-matrix (M_α, I, I) is an integral pseudo-matrix in the sense of Definition 1.7.1, and the map f associated to it is multiplication by α. Using Theorem 1.7.2 and the subsequent remarks, we see that

$$\mathbb{Z}_L/\alpha\mathbb{Z}_L \simeq \bigoplus_{1 \le i \le n} \mathbb{Z}_K/\mathfrak{d}_i$$

for some (unique) integral ideals $\mathfrak{d}_i$ satisfying $\mathfrak{d}_{i-1} \subset \mathfrak{d}_i$ for $2 \le i \le n$, which can be computed using the SNF algorithm in Dedekind domains (Algorithm 1.7.3). Set $\mathfrak{d} = \prod_{1 \le i \le n} \mathfrak{d}_i$. Using the formulas and notation of Theorem 1.7.2, we have, since $\mathfrak{a} = \mathfrak{b}$,

$$\mathbb{Z}_K = \det(V) \det(M_\alpha) \det(U) \mathbb{Z}_K = \mathfrak{b}'\mathfrak{b}^{-1} \mathcal{N}_{L/K}(\alpha) \mathfrak{a}\mathfrak{a}'^{-1} = \mathcal{N}_{L/K}(\alpha)\mathfrak{d}^{-1} ;$$

hence

$$\mathcal{N}_{L/K}(\alpha)\mathbb{Z}_K = \mathfrak{d} = \mathcal{N}_{L/K}(\alpha\mathbb{Z}_L) ,$$

proving the proposition. □

Remarks

(1) As in the absolute case, the result of (1) (multiplicativity of the norm on ideals) is valid only for the maximal order $\mathbb{Z}_L$ and not for a suborder.
(2) The result of (2) above (transitivity of the norm on ideals) is the same if we replace $\mathbb{Q}$ by any other number field k.
(3) We remark in [Coh0] that for an absolute extension,

$$\mathcal{N}_{K/\mathbb{Q}}(\alpha\mathbb{Z}_K) = \left|\mathcal{N}_{K/\mathbb{Q}}(\alpha)\right|$$

with an absolute-value sign. This is equivalent to

$$\mathcal{N}_{K/\mathbb{Q}}(\alpha\mathbb{Z}_K)\mathbb{Z} = \mathcal{N}_{K/\mathbb{Q}}(\alpha)\mathbb{Z} \ ,$$

as claimed above.
(4) Thanks to the multiplicativity of the norm, as in the absolute case we can define the norm of a fractional ideal. It is equal to the fractional index-ideal in the sense of Definition 1.2.33.

An equivalent definition of the norm of an ideal results from the following proposition.

Proposition 2.2.13. *Let I be a fractional ideal of L. Then $\mathcal{N}_{L/K}(I)$ is the ideal of K generated by all the $\mathcal{N}_{L/K}(\alpha)$ for $\alpha \in I$. More precisely, there exist α and β in I such that $\mathcal{N}_{L/K}(I) = \mathcal{N}_{L/K}(\alpha)\mathbb{Z}_K + \mathcal{N}_{L/K}(\beta)\mathbb{Z}_K$.*

Proof. Clearly, if $\alpha \in I$ then $\mathcal{N}_{L/K}(\alpha) \in \mathcal{N}_{L/K}(I)$. To prove the converse, we proceed in two steps. First, by the approximation theorem in Dedekind domains, we can find $\alpha \in L$ such that $v_{\mathfrak{P}}(\alpha) = v_{\mathfrak{P}}(I)$ for all $\mathfrak{P}$ above the prime ideals $\mathfrak{p}$ of K such that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(I)) \neq 0$, and $v_{\mathfrak{P}}(\alpha) \geq 0$ for all other $\mathfrak{P}$. With this choice of α, it is clear that $\alpha \in I$ and that $\mathcal{N}_{L/K}(\alpha) = \mathcal{N}_{L/K}(I)\mathfrak{a}$ with $\mathfrak{a}$ an integral ideal of $\mathbb{Z}_K$ coprime to $\mathcal{N}_{L/K}(I)$. Applying once again the approximation theorem, we can find $\beta \in \mathbb{Z}_L$ such that $v_{\mathfrak{P}}(\beta) = v_{\mathfrak{P}}(I)$ for all $\mathfrak{P}$ above the prime ideals $\mathfrak{p}$ of K such that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha)) \neq 0$, and $v_{\mathfrak{P}}(\beta) \geq 0$ for all other $\mathfrak{P}$. It is clear that $\beta \in I$ and that $\mathcal{N}_{L/K}(\beta) = \mathcal{N}_{L/K}(I)\mathfrak{b}$, where $\mathfrak{b}$ is an integral ideal coprime to $\mathcal{N}_{L/K}(\alpha)$, hence in particular to $\mathfrak{a}$. Thus, the ideal generated by $\mathcal{N}_{L/K}(\alpha)$ and $\mathcal{N}_{L/K}(\beta)$ is equal to $\mathcal{N}_{L/K}(I)$, proving the proposition. □

We consider now the special case of prime ideals. If $\mathfrak{P}$ is a prime ideal of $\mathbb{Z}_L$, the ideal $\mathfrak{p} = \mathfrak{P} \cap \mathbb{Z}_K$ is clearly a prime ideal of $\mathbb{Z}_K$, and we say that $\mathfrak{P}$ is *above* $\mathfrak{p}$, or that $\mathfrak{p}$ is *below* $\mathfrak{P}$. We have the usual formulas

$$\mathfrak{p}\mathbb{Z}_L = \prod_{1 \leq i \leq g} \mathfrak{P}_i^{e_i} \ ,$$

where the $\mathfrak{P}_i$ are all the ideals above $\mathfrak{p}$ and the $e_i = e(\mathfrak{P}_i/\mathfrak{p})$ are the *ramification indices*. If

$$f_i = f(\mathfrak{P}_i/\mathfrak{p}) = \dim_{\mathbb{Z}_K/\mathfrak{p}}(\mathbb{Z}_L/\mathfrak{P}_i) \ ,$$

we call f_i the *residual degree* of $\mathfrak{P}_i$, and we have the formula

$$\sum_{1 \le i \le g} e_i f_i = n = [L : K] \ .$$

In Section 2.4.3 we give an algorithm for computing the $\mathfrak{P}_i$, e_i, and f_i, generalizing the Buchmann–Lenstra algorithm [Coh0, Algorithm 6.2.9]. For now, we note the following.

Lemma 2.2.14. *Let $\mathfrak{P}$ be a prime ideal of L above $\mathfrak{p}$, and let $f = f(\mathfrak{P}/\mathfrak{p})$ be its residual degree. Then*

$$\mathcal{N}_{L/K}(\mathfrak{P}) = \mathfrak{p}^f \ .$$

Proof. We have $\mathbb{Z}_L/\mathfrak{P} \simeq (\mathbb{Z}_K/\mathfrak{p})^f$ as $\mathbb{Z}_K/\mathfrak{p}$-modules, hence also as $\mathbb{Z}_K$-modules, thus we conclude by Proposition 1.2.34. Note that the elementary divisors $\mathfrak{d}_i$ of $\mathbb{Z}_L/\mathfrak{P}$ are given by $\mathfrak{d}_i = \mathfrak{p}$ for $1 \le i \le f$ and $\mathfrak{d}_i = \mathbb{Z}_K$ for $f < i \le n$. □

Remarks

(1) We can use this lemma to give still another definition of the norm of an ideal: we *define* $\mathcal{N}_{L/K}(\mathfrak{P})$ as $\mathfrak{p}^f$, and extend to all fractional ideals by multiplicativity. Thanks to Proposition 2.2.12, this definition agrees with the preceding one.

(2) We thus have seen three definitions of the relative norm of an ideal I: first as the order-ideal of the torsion module $\mathbb{Z}_L/I$; second as the ideal generated by the norms of the elements of I; third as the power product of the norms of the prime ideals dividing I with the definition given above.

(3) A fourth definition is to set

$$\mathcal{N}_{L/K}(I) = \Big(\prod_i \sigma_i(I)\Big) \cap K \ ,$$

where the σ_i are all the embeddings of L into $\mathbb{C}$ and the product is considered in the Galois closure of L/K in $\mathbb{C}$ (see Exercise 26).

2.3 Representation and Operations on Ideals

2.3.1 Representation of Ideals

Let $L = K(\theta)$ be a relative extension of degree n, where we assume θ to be an algebraic integer, and let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis, which we

may assume to be in relative HNF. This pseudo-basis is represented by a pseudo-matrix $(H_Z, \mathfrak{a}_i)$ in HNF on $(1, \theta, \ldots, \theta^{n-1})$.

Now let I be a (nonzero) ideal of L. Since I is a torsion-free $\mathbb{Z}_K$-module of rank n, it also has a $\mathbb{Z}_K$-pseudo-basis $(\gamma_i, \mathfrak{c}_i)_{1 \le i \le n}$ such that

$$I = \bigoplus_{1 \le i \le n} \mathfrak{c}_i \gamma_i \ ,$$

and we may also assume that this pseudo-basis is given on $(1, \theta, \ldots, \theta^{n-1})$ by a pseudo-matrix $(H_I, \mathfrak{c}_i)$ in HNF.

The following proposition gives most of the information that we need on the ideal I. In particular, it allows us to determine whether I is an integral ideal and to compute $\mathcal{N}_{L/K}(I)$.

Proposition 2.3.1. *Keep the above notation, let $H = H_Z^{-1} H_I$ (which is the matrix giving the γ_j on the basis of the ω_i), and write $H = (h_{i,j})$.*

(1) *Let $\mathfrak{q}_i = \mathfrak{a}_i^{-1}$ (which are integral ideals by Corollary 2.2.9). Then for all $i \le j$ we have*

$$\mathfrak{c}_j \mid \mathfrak{c}_j \mathfrak{q}_{j+1-i} \mid \mathfrak{c}_i \ ,$$

and, in particular,

$$\mathfrak{c}_n \mid \mathfrak{c}_{n-1} \mid \cdots \mid \mathfrak{c}_1 \ .$$

(2) *For all j we have $\gamma_j \in \mathbb{Z}_K[\theta]$, and for all i and j we have $h_{i,j} \in \mathbb{Z}_K$.*

(3) *I is an integral ideal if and only if for all i and j with $i \le j$ we have $h_{i,j} \in \mathfrak{a}_i \mathfrak{c}_j^{-1}$, which implies in particular $\mathfrak{c}_j \subset \mathfrak{a}_j$ for all j (since $h_{j,j} = 1$).*

(4) *For all I (not necessarily integral) we have*

$$\mathcal{N}_{L/K}(I) = \prod_{1 \le j \le n} \mathfrak{c}_j \mathfrak{a}_j^{-1} = \prod_{1 \le j \le n} \mathfrak{c}_j \mathfrak{q}_j \ .$$

Proof. The proof of (1) is essentially identical to the proof of Corollary 2.2.9 (3): since the leading term of $\gamma_i \omega_{j+1-i}$ is θ^{j-1} and I is an ideal, we must have $\mathfrak{c}_i \mathfrak{a}_{j+1-i} \subset \mathfrak{c}_j$, in other words $\mathfrak{c}_j \mathfrak{q}_{j+1-i} \mid \mathfrak{c}_i$, and (1) follows since $\mathfrak{q}_{j+1-i}$ is an integral ideal.

Since I is a $\mathbb{Z}_L$-module, it is in particular a $\mathbb{Z}_K[\theta]$-module and a projective $\mathbb{Z}_K$-module of rank n, so Proposition 2.2.8 (2) implies that $\gamma_j \in \mathbb{Z}_K[\theta]$. This means that the matrix H_I giving the γ_j on the basis $(1, \theta, \ldots, \theta^{n-1})$ has entries in $\mathbb{Z}_K$. This is also the case for the matrix H_Z, and since this matrix has determinant 1, it follows that $H = H_Z^{-1} H_I$ also has entries in $\mathbb{Z}_K$, proving (2).

The ideal I is integral if and only if for all j we have $\mathfrak{c}_j \gamma_j \subset \mathbb{Z}_L$, if and only if for all $c \in \mathfrak{c}_j$

$$c \sum_{i \le j} h_{i,j} \omega_i = \sum_{i \ge 1} x_i \omega_i \quad \text{with} \quad x_i \in \mathfrak{a}_i \ ,$$

hence if and only if for all $c \in \mathfrak{c}_j$, $ch_{i,j} \in \mathfrak{a}_i$, hence $h_{i,j} \in \mathfrak{c}_j^{-1}\mathfrak{a}_i$, proving (3).

Let us prove (4). By the elementary divisor theorem for torsion-free modules (Theorem 1.2.35), there exists a K-basis (e_i), ideals $\mathfrak{b}_i$ and $\mathfrak{d}_i$ such that $\mathfrak{d}_i \mid \mathfrak{d}_{i-1}$ for $i \geq 2$, and

$$\mathbb{Z}_L = \bigoplus_i \mathfrak{b}_i e_i, \quad I = \bigoplus \mathfrak{d}_i \mathfrak{b}_i e_i \ .$$

By definition, we have $\mathcal{N}_{L/K}(I) = \prod_i \mathfrak{d}_i$.

Let Ω be the matrix giving the ω_j in terms of the e_i. Then the matrix giving the α_j in terms of the e_i is equal to ΩH. By Proposition 1.4.2, since $(e_j, \mathfrak{b}_j)$ and $(\omega_j, \mathfrak{a}_j)$ are both pseudo-bases of $\mathbb{Z}_L$, we have

$$\prod_j \mathfrak{b}_j = \det(\Omega) \prod_j \mathfrak{a}_j \ .$$

Similarly,

$$\prod_j \mathfrak{d}_j \mathfrak{b}_j = \det(\Omega H) \prod_j \mathfrak{c}_j \ .$$

Since $\det(H) = 1$, it follows by dividing that

$$\prod_j \mathfrak{d}_j = \prod_j \mathfrak{c}_j \mathfrak{a}_j^{-1} \ ,$$

proving (4) and the proposition. □

In view of this proposition, we will always assume that an ideal is represented by a pseudo-matrix $(H, \mathfrak{c}_i)$ in HNF on a basis $(\omega_1, \ldots, \omega_n)$ (although the pseudo-matrix giving the integral pseudo-basis on $(1, \theta, \ldots, \theta^{n-1})$ must evidently also be kept). This has the usual advantages of the HNF, in particular the uniqueness property (see Proposition 2.3.2 below). The main disadvantage of the HNF representation is that it is costly, particularly for ideal operations.

In fact, considering the above lemma, it would even be more natural to represent the ideal by the pseudo-matrix $(H, \mathfrak{c}_i \mathfrak{a}_i^{-1})$. For this to make sense, we would have had to define the notion of pseudo-matrix with respect to a pseudo-basis and not only with respect to a basis as we have done up to now, so as to take into account not only the ω_i but also the ideals $\mathfrak{a}_i$. We leave the (trivial) definitions and modifications to the reader.

To test ideals for equality (hence also for inclusion using $I \subset J$ if and only if $I + J = J$), we need to have uniqueness of the representation of an ideal. The HNF representation does give uniqueness if one is careful about the choice of the off-diagonal entries (see Corollary 1.4.11). More precisely:

Proposition 2.3.2. *Keep the above notation. Assume that I is an* integral *ideal, so that in particular by Proposition 2.3.1, we have for all i, $\mathfrak{c}_i \subset \mathfrak{a}_i$ (or, equivalently, $\mathfrak{a}_i \mid \mathfrak{c}_i$). Let S_i be a system of representatives of $\mathbb{Z}_K/(\mathfrak{c}_i\mathfrak{a}_i^{-1})$. For all i and j such that $i < j$, choose $c_{i,j} \in \mathfrak{a}_i\mathfrak{c}_j^{-1}$ such that $v_\mathfrak{p}(c_{i,j}) = v_\mathfrak{p}(\mathfrak{a}_i\mathfrak{c}_j^{-1})$ for all prime ideals $\mathfrak{p}$ of $\mathbb{Z}_K$ such that $v_\mathfrak{p}(\mathfrak{a}_i) < v_\mathfrak{p}(\mathfrak{c}_i)$. Then for $i < j$ we may choose $h_{i,j} \in c_{i,j}S_i$, and the pseudo-matrix $(H, \mathfrak{c}_j)$ is then uniquely determined by the ideal I.*

Proof. According to Corollary 1.4.11, to obtain a unique pseudo-matrix $(H, \mathfrak{c}_j)$ we must choose $h_{i,j} \in S_{i,j}$, where $S_{i,j}$ is a system of representatives of $K/\mathfrak{c}_i\mathfrak{c}_j^{-1}$. Since I is an integral ideal, $h_{i,j} \in \mathfrak{a}_i\mathfrak{c}_j^{-1}$, so it is enough to define a system of representatives of $\mathfrak{a}_i\mathfrak{c}_j^{-1}/\mathfrak{c}_i\mathfrak{c}_j^{-1}$. If $c_{i,j}$ satisfies the conditions of the proposition, it is easy to check that the map $x \mapsto c_{i,j}x$ induces an isomorphism from $\mathbb{Z}_K/(\mathfrak{c}_i\mathfrak{a}_i^{-1})$ to $\mathfrak{a}_i\mathfrak{c}_j^{-1}/\mathfrak{c}_i\mathfrak{c}_j^{-1}$ (see Exercise 27), proving the proposition. Note that by Proposition 2.3.1, we have for $i \leq j$, $\mathfrak{a}_i\mathfrak{c}_j^{-1} \subset \mathfrak{a}_i\mathfrak{c}_i^{-1} \subset \mathbb{Z}_K$; hence $\mathfrak{a}_i\mathfrak{c}_j^{-1}$ is an integral ideal. □

Following this proposition, we can give an algorithm that gives a small HNF pseudo-matrix for an integral ideal.

Algorithm 2.3.3 (Small HNF Pseudo-Matrix of an Integral Ideal). Let L/K be a relative extension of degree n. Given an integral ideal I by a pseudo-matrix in HNF $(H, (\mathfrak{c}_j))$ with $H = (h_{i,j})$ expressed on a relative integral pseudo-basis $(\omega_i, \mathfrak{a}_i)$, this algorithm gives another such pseudo-matrix in HNF $(H', \mathfrak{c}_j)$ with $H' = (h'_{i,j})$ having "reduced" entries.

1. [Compute the ideals $\mathfrak{c}_j^{-1}$] For $1 \leq j \leq n$, compute the ideal $\mathfrak{b}_j \leftarrow \mathfrak{c}_j^{-1}$ using, for example, [Coh0, Algorithm 4.8.21]. Then set $i \leftarrow n$ and $H' \leftarrow H$.
2. [Loop on rows] Set $i \leftarrow i - 1$. If $i = 0$, output $(H', \mathfrak{c}_j)$ with $H' = (h'_{i,j})$ and terminate the algorithm. Otherwise, set $j \leftarrow n + 1$.
3. [Loop on columns] Set $j \leftarrow j - 1$. If $j = i$, go to step 2.
4. [Main reduction step] Set $\mathfrak{a} \leftarrow \mathfrak{c}_i\mathfrak{b}_j$. Using Algorithm 1.4.13 (with partial LLL-reduction), compute $\lambda \in \mathfrak{c}_i\mathfrak{b}_j$ such that $h_{i,j} - \lambda$ is "small" in the sense of that algorithm.
5. [Update column j] For $1 \leq k \leq i$, set $h'_{k,j} \leftarrow h'_{k,j} - \lambda h'_{k,i}$ and go to step 3.

Definition and Proposition 2.3.4. *Let I be a fractional ideal of $\mathbb{Z}_L$.*

(1) *We will say that I is a* primitive ideal *if I is an integral ideal of $\mathbb{Z}_L$ and if for any integral ideal $\mathfrak{a}$ of $\mathbb{Z}_K$ different from $\mathbb{Z}_K$, $\mathfrak{a}^{-1}I$ is not an integral ideal.*

(2) *If I is a fractional ideal of $\mathbb{Z}_L$, there exists a unique fractional ideal $\mathfrak{a}$ of K such that $\mathfrak{a}^{-1}I$ is a primitive ideal. This ideal $\mathfrak{a}$ will be called the* content *of the ideal I.*

Proof. Consider the set E of all fractional ideals $\mathfrak{a}$ of K such that $I \subset \mathfrak{a}\mathbb{Z}_L$. This set is nonempty since if $d \in \mathbb{Z}$ is a denominator for I, we may choose $\mathfrak{a} = (1/d)\mathbb{Z}_K$. Set $\mathfrak{c} = \bigcap_{\mathfrak{a}\in E} \mathfrak{a}$. Then $\mathfrak{c}$ is an ideal of K that clearly still belongs to E, so it is the unique minimal element of E. It follows that $\mathfrak{c}^{-1}I$ is an integral ideal. In addition, if $\mathfrak{b}$ is an integral ideal different from $\mathbb{Z}_K$ such that $\mathfrak{b}^{-1}\mathfrak{c}^{-1}I$ is integral, then $\mathfrak{bc} \in E$ and $\mathfrak{bc}$ is a strict subset of $\mathfrak{c}$ — which is absurd since $\mathfrak{c}$ is minimal — so $\mathfrak{c}^{-1}I$ is primitive. Finally, if $\mathfrak{a}^{-1}I$ is primitive, then $\mathfrak{a} \in E$, hence $\mathfrak{c} \subset \mathfrak{a}$, so $\mathfrak{ca}^{-1}$ is an integral ideal such that $(\mathfrak{ca}^{-1})^{-1}\mathfrak{a}^{-1}I$ is integral. Hence $\mathfrak{ca}^{-1} = \mathbb{Z}_K$, so $\mathfrak{a} = \mathfrak{c}$, proving uniqueness. □

Proposition 2.3.5. *Keep the notation of Proposition 2.3.1, in particular that I is an ideal of L, $(\gamma_j, \mathfrak{c}_j)$ is a pseudo-basis of I, and the matrix $(h_{i,j})$ of the γ_j on the ω_i is in HNF.*

(1) *The content $c(I)$ of I is given by*

$$c(I) = \sum_{1\le i\le j\le n} h_{i,j}\mathfrak{c}_j\mathfrak{q}_i = \sum_{1\le i\le j\le n} h_{i,j}\mathfrak{c}_j\mathfrak{a}_i^{-1} \ .$$

(2) *The ideal I is an integral ideal of $\mathbb{Z}_L$ if and only if $c(I)$ is an integral ideal of $\mathbb{Z}_K$.*
(3) *The ideal I is primitive if and only if $c(I) = \mathbb{Z}_K$.*

Proof. The proof follows immediately from Proposition 2.3.1 and is left to the reader (Exercise 28). □

The other privileged representation of an ideal I is a two-element representation $I = \alpha\mathbb{Z}_L + \beta\mathbb{Z}_L$, which is independent of the relative structure. Considering the definition of a pseudo-matrix, it is more natural to give the following definition in the relative case.

Definition 2.3.6. *Let I be an ideal of L. We say that $((\alpha, \mathfrak{a}), (\beta, \mathfrak{b}))$ is a* pseudo–two-element representation *of the ideal I if $\mathfrak{a}$ and $\mathfrak{b}$ are ideals* of $\mathbb{Z}_K$ *(not necessarily integral) and if α and β are elements of L such that*

$$I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L \ .$$

In other words, a pseudo–two-element representation of I is a representation of I by two pseudo-elements in the sense of Definition 1.4.1.

When $\mathbb{Z}_K$ is a principal ideal domain, $\mathfrak{a}$ and $\mathfrak{b}$ are principal ideals, so this does give a two-element representation. In the general case, however, this definition is more flexible than the representation $I = \alpha\mathbb{Z}_L + \beta\mathbb{Z}_L$. Note that if $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$, then $\alpha\mathfrak{a} \subset I$ and $\beta\mathfrak{b} \subset I$, but we do not necessarily have α or β in I. By abuse of language, we will sometimes talk about "two-element representations" instead of "pseudo–two-element representations".

It is important to be able to go back and forth between HNF and two-element representations. As usual, in one direction this is straightforward. Let $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$ be a pseudo–two-element representation of I, and let $(\omega_i, \mathfrak{a}_i)_{1\le i\le n}$ be an integral pseudo-basis. Then $(\alpha\omega_i, \beta\omega_i, \mathfrak{a}\mathfrak{a}_i, \mathfrak{b}\mathfrak{a}_i)_i$ is a $2n$-element pseudo-generating set for I. Hence we obtain the HNF of I by computing the HNF of the corresponding pseudo-matrix, using one of the HNF algorithms of Chapter 1. We see here that the introduction of the extra data $\mathfrak{a}$ and $\mathfrak{b}$ has not added any complexity to the problem.

Given the HNF $(\gamma_i, \mathfrak{c}_i)_{1\le i\le n}$ of I, finding a pseudo–two-element representation is slightly trickier. Of course, using the approximation theorem ([Coh0, Propositions 4.7.7 and 4.7.8]), we can give a deterministic algorithm for doing so, but this will be costly in general. A better way is to use a simple-minded generalization of [Coh0, Algorithm 4.7.10] based on the following lemma.

Lemma 2.3.7. *Let I be an integral ideal of $\mathbb{Z}_L$. Let $\alpha \in K$, let $\mathfrak{a}$ be a fractional ideal of $\mathbb{Z}_K$ such that $\alpha\mathfrak{a} \subset I$, and assume that*

$$\mathcal{N}_{L/K}(I) + \mathcal{N}_{L/K}(\alpha\mathfrak{a})(\mathcal{N}_{L/K}(I))^{-1} = \mathbb{Z}_K \ .$$

Then $I = \mathcal{N}_{L/K}(I)\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$; in other words,

$$((1, \mathcal{N}_{L/K}(I)), (\alpha, \mathfrak{a}))$$

is a pseudo–two-element representation of I.

Proof. We have a $\mathbb{Z}_K$-module isomorphism

$$\mathbb{Z}_L/I \simeq \bigoplus_i \mathbb{Z}_K/\mathfrak{d}_i \ ,$$

with $\mathcal{N}_{L/K}(I) = \prod_i \mathfrak{d}_i$. Since the $\mathfrak{d}_i$ are integral ideals, we have

$$\mathcal{N}_{L/K}(I) \cdot (\mathbb{Z}_K/\mathfrak{d}_i) = \{\overline{0}\}$$

for all i, hence $\mathcal{N}_{L/K}(I) \cdot (\mathbb{Z}_L/I) = \{\overline{0}\}$, in other words $\mathcal{N}_{L/K}(I) \in I$. Since $\alpha\mathfrak{a} \subset I$ and I is an ideal, it follows that $\mathcal{N}_{L/K}(I)\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L \subset I$.

Conversely, let $J = \mathcal{N}_{L/K}(I)\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$, let $\mathfrak{P}$ be a prime ideal of $\mathbb{Z}_L$, and let $\mathfrak{p} = \mathfrak{P} \cap \mathbb{Z}_K$ be the prime ideal of $\mathbb{Z}_K$ below $\mathfrak{P}$. We will show that $v_{\mathfrak{P}}(J) \le v_{\mathfrak{P}}(I)$ for all $\mathfrak{P}$, which will show that $I \subset J$, and hence the equality $I = J$, as claimed in the lemma. If $v_{\mathfrak{P}}(J) = 0$, there is nothing to prove since I is an integral ideal, so assume that $v_{\mathfrak{P}}(J) > 0$. Since

$$v_{\mathfrak{P}}(J) = \min(v_{\mathfrak{P}}(\mathcal{N}_{L/K}(I)), v_{\mathfrak{P}}(\alpha\mathfrak{a})) \ ,$$

we have $\mathfrak{P} \mid \mathcal{N}_{L/K}(I)$ and hence $\mathfrak{p} \mid \mathcal{N}_{L/K}(I)$. By assumption, this implies that $\mathfrak{p} \nmid \mathcal{N}_{L/K}(\alpha\mathfrak{a})(\mathcal{N}_{L/K}(I))^{-1}$, or in other words that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha\mathfrak{a}I^{-1})) = 0$. This means that $v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = v_{\mathfrak{Q}}(I)$ for all prime ideals $\mathfrak{Q}$ above $\mathfrak{p}$, and in particular for $\mathfrak{Q} = \mathfrak{P}$. Thus, when $\mathfrak{P} \mid J$,

$$v_{\mathfrak{P}}(J) = \min(v_{\mathfrak{P}}(\mathcal{N}_{L/K}(I)), v_{\mathfrak{P}}(I)) \leq v_{\mathfrak{P}}(I) \ ,$$

proving the lemma. □

To obtain an algorithm for a pseudo–two-element representation from this lemma, we must do two things. First, compute the ideal $\mathcal{N}_{L/K}(I)$, which is easily done from the HNF representation using Proposition 2.3.1 (4). Second, we must look for α and $\mathfrak{a}$ satisfying the required properties. For this, we write $\alpha = \sum_{1 \leq i \leq n} x_i \gamma_i$ on the pseudo-basis $(\gamma_i, \mathfrak{c}_i)$, with $x_i \in \mathfrak{c}_i$, and try small coefficients x_i until a suitable element α is obtained. In practice, it will be obtained very rapidly; in fact, very frequently we can take $(\alpha, \mathfrak{a}) = (\gamma_i, \mathfrak{c}_i)$ for some i. Thus, the following algorithm is reasonable.

Algorithm 2.3.8 (Pseudo–Two-Element Representation of an Ideal). Given a relative extension L/K of degree n and an integral ideal I of L given by a pseudo-generating set $(\gamma_i, \mathfrak{c}_i)_{1 \leq i \leq k}$, this algorithm computes $\mathcal{N}_{L/K}(I)$ and $(\alpha, \mathfrak{a})$ such that $((1, \mathcal{N}_{L/K}(I)), (\alpha, \mathfrak{a}))$ is a pseudo–two-element representation of I. We let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis of $\mathbb{Z}_L$.

1. [Compute HNF] If necessary, using one of the algorithms for HNF in Dedekind domains, compute the HNF corresponding to the pseudo-generating set $(\gamma_i, \mathfrak{c}_i)$, and replace $(\gamma_i, \mathfrak{c}_i)$ by this HNF. Set $\mathfrak{n} \leftarrow \prod_{1 \leq i \leq n} \mathfrak{c}_i \mathfrak{a}_i^{-1}$ (thus $\mathcal{N}_{L/K}(I) = \mathfrak{n}$).
2. [Check generators] For $i = 1, \ldots, n$, do the following. Compute the ideal sum $\mathfrak{n} + \mathcal{N}_{L/K}(\gamma_i \mathfrak{c}_i)\mathfrak{n}^{-1}$. If it is equal to $\mathbb{Z}_K$, output the pseudo–two-element representation $((1, \mathfrak{n}), (\gamma_i, \mathfrak{c}_i))$ and terminate the algorithm.
3. [Choose random elements of $\mathfrak{c}_i$] Using Algorithm 1.3.13, for $i = 2, \ldots, k$ choose random elements $\lambda_i \in \mathfrak{c}_i$, and let $\alpha \leftarrow \sum_{2 \leq i \leq k} \lambda_i \gamma_i$.
4. [Check α] Compute the ideal sum $\mathfrak{n} + \mathcal{N}_{L/K}(\alpha)\mathfrak{n}^{-1}$. If it is equal to $\mathbb{Z}_K$, output the pseudo–two-element representation $((1, \mathfrak{n}), (\alpha, \mathbb{Z}_K))$ and terminate the algorithm; otherwise, go to step 3.

If I is not an integral ideal, we simply multiply I by a suitable denominator d to make it integral, and divide by d the pseudo–two-element representation found by this algorithm.

In the case where I is a *prime* ideal, there is a simpler variant of this algorithm which we give below (Algorithm 2.3.11). We postpone to that algorithm the discussion of the above algorithm.

2.3.2 Representation of Prime Ideals

As in the absolute case, the case of prime ideals of $\mathbb{Z}_L$ is particularly important — and also simpler.

We first note that [Coh0, Theorem 4.8.13] can trivially be extended to the relative case, and we leave the proof to the reader:

Proposition 2.3.9. *Let L/K be a relative extension, with $L = K(\theta)$ and θ an algebraic integer whose minimal monic polynomial in $K[X]$ is denoted $T(X)$, and let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$. Let $\overline{T(X)} = \prod_{1\le i\le g} \overline{T_i(X)}^{e_i}$ be the factorization of $\overline{T(X)}$ into a product of powers of distinct, monic, irreducible polynomials in $(\mathbb{Z}_K/\mathfrak{p})[X]$. If $\mathfrak{p}$ does not divide the index-ideal $\mathfrak{f} = [\mathbb{Z}_L : \mathbb{Z}_K[\theta]]$, the prime ideal decomposition of $\mathfrak{p}\mathbb{Z}_L$ is given by*

$$\mathfrak{p}\mathbb{Z}_L = \prod_{1\le i\le g} \mathfrak{P}_i^{e_i} \ ,$$

with

$$\mathfrak{P}_i = ((1,\mathfrak{p}),(T_i(\theta),\mathbb{Z}_K)) = \mathfrak{p}\mathbb{Z}_L + T_i(\theta)\mathbb{Z}_L$$

and $f_i = f(\mathfrak{P}_i/\mathfrak{p}) = \deg(T_i)$.

In the general case, we have the following lemma, which gives more precise information than Lemma 2.3.7 in the case of prime ideals.

Lemma 2.3.10. *Let $\mathfrak{P}$ be a prime ideal of L above a prime ideal $\mathfrak{p}$ of K, let $f = f(\mathfrak{P}/\mathfrak{p}) = \dim_{\mathbb{Z}_K/\mathfrak{p}}(\mathbb{Z}_L/\mathfrak{P})$ be the residual degree of $\mathfrak{P}$, and finally let $\alpha \in L$ and $\mathfrak{a}$ be a fractional ideal of K such that $\alpha\mathfrak{a} \subset \mathfrak{P}$. Let π be any element of $\mathfrak{p}\mathfrak{a}^{-1}$ such that $v_{\mathfrak{p}}(\pi\mathfrak{a}) = 1$. Then*

$$\mathfrak{P} = ((1,\mathfrak{p}),(\alpha,\mathfrak{a})) = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$$

if and only if $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha\mathfrak{a})) = f$ or $v_{\mathfrak{p}}(\mathcal{N}_{L/K}((\alpha+\pi)\mathfrak{a}) = f$.

Proof. Assume first that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha\mathfrak{a})) = f$. Since $\mathfrak{P} \mid \alpha\mathfrak{a}$, we can write $\alpha\mathfrak{a} = \mathfrak{P}I$ for some integral ideal I of $\mathbb{Z}_L$. Since $\mathcal{N}_{L/K}(\mathfrak{P}) = \mathfrak{p}^f$ by Lemma 2.2.14, we have $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(I)) = 0$. This means that $v_{\mathfrak{Q}}(I) = 0$ for all prime ideals $\mathfrak{Q}$ above $\mathfrak{p}$, including $\mathfrak{P}$. Thus $v_{\mathfrak{P}}(\alpha\mathfrak{a}) = 1$ and $v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = 0$ for the other $\mathfrak{Q}$ above $\mathfrak{p}$. Since $v_{\mathfrak{P}}(\mathfrak{p}) \ge 1$, this implies that $\min(v_{\mathfrak{Q}}(\mathfrak{p}), v_{\mathfrak{Q}}(\alpha\mathfrak{a})) = 0$ for all prime ideals $\mathfrak{Q}$ different from $\mathfrak{P}$ (and not only for those above $\mathfrak{p}$) and is equal to 1 for $\mathfrak{Q} = \mathfrak{P}$, thus showing the equality $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$.

Assume now that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}((\alpha+\pi)\mathfrak{a}) = f$. Since $(\alpha+\pi)\mathfrak{a} \subset \alpha\mathfrak{a} + \pi\mathfrak{a} \subset \mathfrak{P} + \mathfrak{p} \subset \mathfrak{P}$, we conclude as above that $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + (\alpha+\pi)\mathfrak{a}\mathbb{Z}_L$, and this is equal to $\mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$ since $\pi\mathfrak{a} \subset \mathfrak{p}$.

Conversely, assume that $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$. This again means that for each prime ideal $\mathfrak{Q}$ above $\mathfrak{p}$ other than $\mathfrak{P}$ we have $v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = 0$ and that $\min(v_{\mathfrak{P}}(\alpha\mathfrak{a}), v_{\mathfrak{P}}(\mathfrak{p})) = 1$. Note that $v_{\mathfrak{P}}(\mathfrak{p}) = e(\mathfrak{P}/\mathfrak{p})$ is the ramification index of $\mathfrak{P}$. We consider two cases. Assume first that $\mathfrak{P}$ is ramified, so that $e(\mathfrak{P}/\mathfrak{p}) > 1$. Then $v_{\mathfrak{P}}(\alpha\mathfrak{a}) = 1$, and since $v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = 0$ for the other prime ideals $\mathfrak{Q}$ above $\mathfrak{p}$, we have

$$v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha\mathfrak{a})) = \sum_{\mathfrak{Q}\mid\mathfrak{p}} f(\mathfrak{Q}/\mathfrak{p}) v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = f(\mathfrak{P}/\mathfrak{p}) = f \ ,$$

as claimed.

Assume now that $\mathfrak{P}$ is unramified, so that $e(\mathfrak{P}/\mathfrak{p}) = 1$. In this case, we can only assert that $v_{\mathfrak{P}}(\alpha\mathfrak{a}) \geq 1$. If $v_{\mathfrak{P}}(\alpha\mathfrak{a}) = 1$, we conclude as before that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha\mathfrak{a})) = f$. So assume $v_{\mathfrak{P}}(\alpha\mathfrak{a}) > 1$. Since $v_{\mathfrak{P}}(\alpha\mathfrak{a}) > v_{\mathfrak{P}}(\pi\mathfrak{a})$, we have $v_{\mathfrak{P}}(\alpha+\pi)\mathfrak{a} = v_{\mathfrak{P}}(\pi\mathfrak{a}) = 1$, and for $\mathfrak{Q}$ above $\mathfrak{p}$ but different from $\mathfrak{P}$, we have

$$0 = v_{\mathfrak{Q}}(\alpha\mathfrak{a}) < v_{\mathfrak{Q}}(\mathfrak{p}) \leq v_{\mathfrak{Q}}(\pi\mathfrak{a}) \ ,$$

hence $v_{\mathfrak{Q}}((\alpha+\pi)\mathfrak{a}) = v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = 0$, and it follows as before that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}((\alpha+\pi)\mathfrak{a})) = f$, as claimed. □

Thus, finding a pseudo–two-element representation of a prime ideal is easier than in the general case (Algorithm 2.3.8). Thanks to Proposition 2.3.9 a prime ideal is, however, usually obtained directly together with a pseudo–two-element representation. We will always assume that a prime ideal is represented in this way. When a pseudo–two-element representation is not known, the following algorithm allows us to compute such a representation, using Lemma 2.3.10.

Algorithm 2.3.11 (Pseudo–Two-Element Representation of a Prime Ideal). Given a relative extension L/K and a prime ideal $\mathfrak{P}$ of L given by a pseudo-generating set $(\gamma_i, \mathfrak{c}_i)_{1\leq i\leq k}$, this algorithm computes a pseudo-element $(\alpha, \mathfrak{a})$ such that $((1,\mathfrak{p}),(\alpha,\mathfrak{a}))$ is a pseudo–two-element representation of $\mathfrak{P}$. We assume that the relative norm $\mathfrak{p}^f$ of $\mathfrak{P}$ is known (this is always the case in practice and can always be obtained from the $(\gamma_i, \mathfrak{c}_i)$) and that $\gamma_1 = 1$, $\mathfrak{c}_1 = \mathfrak{p}$ (if this is not the case, add it to the generating set). We let $\mathfrak{p} = p\mathbb{Z}_K + \pi\mathbb{Z}_K$ with $v_{\mathfrak{p}}(\pi) = 1$ (if this is not the case, replace π by $\pi + p$).

1. [Check generators] For $i = 1, \ldots, k$, do the following. Compute

$$v \leftarrow v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\gamma_i\mathfrak{c}_i)) \ .$$

If $v = f$, output $(\gamma_i, \mathfrak{c}_i)$ and terminate. Otherwise, compute

$$v \leftarrow v_{\mathfrak{p}}(\mathcal{N}_{L/K}((\gamma_i+\pi)\mathfrak{c}_i)) \ .$$

If $v = f$, output $(\gamma_i, \mathfrak{c}_i)$ and terminate.

2. [Choose random elements of $\mathfrak{c}_i$] Using Algorithm 1.3.13, for $i = 2, \ldots, k$, choose random elements $\lambda_i \in \mathfrak{c}_i$, and let $\alpha \leftarrow \sum_{2\leq i\leq k} \lambda_i\gamma_i$.

3. [Check α] Compute $v \leftarrow v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha))$. If $v = f$, output $(\alpha, \mathbb{Z}_K)$ and terminate. Otherwise, compute $v \leftarrow v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\alpha+\pi))$. If $v = f$, output $(\alpha, \mathbb{Z}_K)$ and terminate; otherwise, go to step 2.

Remarks (Also valid for Algorithm 2.3.8)

(1) If $k > n$, to speed up the algorithm, it may be worthwhile first to find a pseudo-basis of $\mathfrak{P}$ using one of the algorithms for HNF in Dedekind domains, as we have done systematically in Algorithm 2.3.8.

(2) The manner in which the random elements of $\mathfrak{c}_i$ are chosen is not important, since the algorithm is expected to find a result very rapidly (of course, it is preferable to take elements that are small in some sense). In fact, we want $\min(e(\mathfrak{P}/\mathfrak{p}), v_{\mathfrak{P}}(\alpha)) = 1$, and $v_{\mathfrak{Q}}(\alpha) = 0$ for all other prime ideals $\mathfrak{Q}$ dividing $\mathfrak{p}$. The probability that a random $\alpha \in \mathfrak{P}$ satisfies these conditions can be estimated to be

$$\prod_{\mathfrak{Q}|\mathfrak{p}} (1 - 1/\mathcal{N}_{L/\mathbb{Q}}(\mathfrak{Q})) \ ,$$

where the product is over all prime ideals above $\mathfrak{p}$ if $\mathfrak{P}$ is unramified, and all prime ideals above $\mathfrak{p}$ except $\mathfrak{P}$ if $\mathfrak{P}$ is ramified. This quantity is not small, so very few trials should be necessary.

(3) We have not used the systematic backtracking method of [Coh0, Algorithm 4.7.10], since this would be in general much more costly and is essentially equivalent to using absolute instead of relative representations. In fact, even in the absolute case, it is probably preferable to use random elements of $\mathbb{Z}$ instead of a systematic backtracking procedure.

(4) There is a completely different method to find a two-element representation (valid also in the absolute case), by directly using the approximation theorem in Dedekind domains. Indeed, we have $\mathfrak{P} = ((1, \mathfrak{p}), (\alpha, \mathfrak{a}))$ if and only if $v_{\mathfrak{Q}}(\alpha\mathfrak{a}) = 0$ for all $\mathfrak{Q} \mid \mathfrak{p}$ and different from $\mathfrak{P}$, and $\min(e(\mathfrak{P}/\mathfrak{p}), v_{\mathfrak{P}}(\alpha\mathfrak{a})) = 1$, which is, for example, the case if $v_{\mathfrak{P}}(\alpha\mathfrak{a}) = 1$. We can plug this in the deterministic version of the approximation theorem with $\mathfrak{a} = \mathbb{Z}_K$ (Proposition 1.3.8), and obtain in this way a two-element representation. In practice, however, this method is less efficient than the random search method described above.

2.3.3 Computing Valuations

As in the absolute case, we also want to compute $\mathfrak{P}$-adic valuations, and for this we proceed in a similar way. Assume that we have computed $(\beta, \mathfrak{b})$ such that

$$\mathfrak{p}\mathfrak{P}^{-1} = \mathfrak{p}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$$

(this is possible; see Algorithm 2.3.14). Let I be an integral ideal of $\mathbb{Z}_L$. Then $v_{\mathfrak{P}}(I)$ is the largest nonnegative integer v such that $\mathfrak{P}^{-v} I \subset \mathbb{Z}_L$ or, equivalently, $(\beta\mathfrak{b}\mathfrak{p}^{-1})^v I \subset \mathbb{Z}_L$, and this can easily be tested. This is the natural generalization of [Coh0, Lemma 4.8.16]. The following lemma generalizes [Coh0, Lemma 4.8.15] for the maximal order.

Lemma 2.3.12. *We have* $\mathfrak{p}\mathfrak{P}^{-1} = \mathfrak{p}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$ *if and only if* $\beta\mathfrak{b}\mathfrak{P} \subset \mathfrak{p}\mathbb{Z}_L$ *and* $\beta\mathfrak{b} \not\subset \mathfrak{p}\mathbb{Z}_L$.

Proof. Assume first that $\mathfrak{p}\mathfrak{P}^{-1} = \mathfrak{p}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$. Thus $\beta\mathfrak{b} \subset \mathfrak{p}\mathfrak{P}^{-1}$; hence $\beta\mathfrak{b}\mathfrak{P} \subset \mathfrak{p}\mathbb{Z}_L$. Furthermore, if we had $\beta\mathfrak{b} \subset \mathfrak{p}\mathbb{Z}_L$, we would have $\mathfrak{p}\mathfrak{P}^{-1} = \mathfrak{p}\mathbb{Z}_L$, hence $\mathfrak{P}^{-1} = \mathbb{Z}_L$, which is impossible since $\mathfrak{P} \neq \mathbb{Z}_L$. Thus, $\beta\mathfrak{b} \not\subset \mathfrak{p}\mathbb{Z}_L$.

Conversely, assume that $\beta\mathfrak{b}\mathfrak{P} \subset \mathfrak{p}\mathbb{Z}_L$ and $\beta\mathfrak{b} \not\subset \mathfrak{p}\mathbb{Z}_L$. Then $\mathfrak{P} \subset \mathfrak{P} + \beta\mathfrak{b}\mathfrak{P}\mathfrak{p}^{-1} \subset \mathbb{Z}_L$, and since $\mathfrak{P}$ is a maximal ideal, it follows that $\mathfrak{P} + \beta\mathfrak{b}\mathfrak{P}\mathfrak{p}^{-1}$ is equal either to $\mathfrak{P}$ or to $\mathbb{Z}_L$. But $\mathfrak{P} + \beta\mathfrak{b}\mathfrak{P}\mathfrak{p}^{-1} = \mathfrak{P}$ implies $\beta\mathfrak{b}\mathfrak{P}\mathfrak{p}^{-1} \subset \mathfrak{P}$; hence $\beta\mathfrak{b} \subset \mathfrak{p}\mathbb{Z}_L$ since $\mathfrak{P}$ is invertible, contrary to our assumption. Thus, $\mathfrak{P} + \beta\mathfrak{b}\mathfrak{P}\mathfrak{p}^{-1} = \mathbb{Z}_L$, and hence $\mathfrak{p}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L = \mathfrak{p}\mathfrak{P}^{-1}$, as claimed. □

Thus, as in the absolute case, we will represent a prime ideal $\mathfrak{P}$ by $\mathfrak{P} = (\mathfrak{p}, (\alpha, \mathfrak{a}), e, f, (\beta, \mathfrak{b}))$, and we will be able conveniently to perform all operations involving $\mathfrak{P}$. We give below algorithms for computing $(\beta, \mathfrak{b})$ and for computing valuations using $(\beta, \mathfrak{b})$.

Algorithm 2.3.13 (Valuation at a Prime Ideal). Let I be an integral ideal of $\mathbb{Z}_L$. Let $(\omega_i, \mathfrak{a}_i)$ be a pseudo-basis in HNF of $\mathbb{Z}_L$, let $(\gamma_i, \mathfrak{c}_i)$ be a pseudo-basis of I, and let $(H, \mathfrak{c}_i)$ be the pseudo-matrix giving this pseudo-basis on the ω_i, where H is in HNF. Finally, let $\mathfrak{P} = (\mathfrak{p}, (\alpha, \mathfrak{a}), e, f, (\beta, \mathfrak{b}))$ be a prime ideal of $\mathbb{Z}_L$ given as above. This algorithm computes the $\mathfrak{P}$-adic valuation $v_{\mathfrak{P}}(I)$ of I.

1. [Make integral] If I is not an integral ideal, let $d \in \mathbb{Z}$ such that dI is integral, set $I \leftarrow dI$ (in other words, set $\mathfrak{c}_i \leftarrow d\mathfrak{c}_i$ for all i), and set $v \leftarrow -v_p(d)e(\mathfrak{P}/p)$, where p is the prime number below $\mathfrak{P}$ and $e(\mathfrak{P}/p)$ is the absolute ramification index of $\mathfrak{P}$. Otherwise, set $v \leftarrow 0$.
2. [Check if $\mathfrak{p} \nmid \mathcal{N}_{L/K}(I)$] If $v_{\mathfrak{p}}\left(\prod_i \mathfrak{a}_i\right) = v_{\mathfrak{p}}\left(\prod_i \mathfrak{c}_i\right)$, output v and terminate the algorithm. Otherwise, set $A \leftarrow H$.
3. [Multiply] Set $\mathfrak{c}_i \leftarrow \mathfrak{b}\mathfrak{c}_i$ for all i, and set $A \leftarrow \beta A$ in the following sense. Each column of A corresponds to an element of L in the K-basis ω_i, and these elements are multiplied by β and expressed again on the ω_i.
4. [Simple test] Using Algorithm 1.6.2, replace $(A, \mathfrak{c}_i)$ by its HNF. If for some j, we have $v_{\mathfrak{p}}(\mathfrak{c}_j) = v_{\mathfrak{p}}(\mathfrak{a}_j)$, output v and terminate the algorithm.
5. [Complete test] If $A = (a_{i,j})$, check whether there exist i and j with $i < j$ such that $v_{\mathfrak{p}}(a_{i,j}) = v_{\mathfrak{p}}(\mathfrak{a}_i) - v_{\mathfrak{p}}(\mathfrak{c}_j)$. If such a pair (i, j) exists, output v and terminate the algorithm. Otherwise, for all j set $\mathfrak{c}_j \leftarrow \mathfrak{p}^{-1}\mathfrak{c}_j$, set $v \leftarrow v + 1$, and go to step 3.

Proof. Step 1 is clear, since

$$v_{\mathfrak{P}}(I) = v_{\mathfrak{P}}(dI) - v_{\mathfrak{P}}(d) = v_{\mathfrak{P}}(dI) - e(\mathfrak{P}/p)v_p(d) \ .$$

On the other hand, by Proposition 2.3.1, we have $\mathcal{N}_{L/K}(I) = \prod_i \mathfrak{c}_i\mathfrak{a}_i^{-1}$. Thus, if I is an integral ideal such that $\mathfrak{p} \nmid \mathcal{N}_{L/K}(I)$, we have $v_{\mathfrak{P}}(I) = 0$, giving step 2. At the end of step 3, we have replaced I by $\beta\mathfrak{b}I$, and by the definition of $(\beta, \mathfrak{b})$, we have $\mathfrak{P} \mid I$ if and only if $\beta\mathfrak{b}I \subset \mathfrak{p}\mathbb{Z}_L$; in other words, if and only if $\beta\mathfrak{b}\mathfrak{p}^{-1}I$ is an integral ideal. By Proposition 2.3.1, this will be true if and only if for all i, $\mathfrak{c}_i \subset \mathfrak{p}\mathfrak{a}_i$ and for all $i \le j$, $a_{i,j} \in \mathfrak{p}\mathfrak{a}_i\mathfrak{c}_j^{-1}$. Since I is an integral ideal, we have $\mathfrak{c}_i \subset \mathfrak{a}_i$ and $a_{i,j} \in \mathfrak{a}_i\mathfrak{c}_j^{-1}$; hence $\mathfrak{P} \nmid I$ if and only if there exists i such that $v_{\mathfrak{p}}(\mathfrak{c}_i) = v_{\mathfrak{p}}(\mathfrak{a}_i)$ or if there exist i and j such that $v_{\mathfrak{p}}(a_{i,j}) = v_{\mathfrak{p}}(\mathfrak{a}_i) - v_{\mathfrak{p}}(\mathfrak{c}_j)$, proving the algorithm's validity. □

Finally, it remains to see how to compute $(\beta, \mathfrak{b})$ satisfying the conditions of Lemma 2.3.12. This is done by using the following algorithm.

Algorithm 2.3.14 (Prime Ideal Inversion). Given a prime ideal $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$, this algorithm computes a pair $(\beta, \mathfrak{b})$ satisfying the conditions of Lemma 2.3.12 — in other words such that $\mathfrak{P}^{-1} = \mathbb{Z}_L + \mathfrak{p}^{-1}\beta\mathfrak{b}\mathbb{Z}_L$ — so as to be able to compute valuations at $\mathfrak{P}$.

1. [Change $\alpha\mathfrak{a}$] If $\mathfrak{a}$ is not an integral ideal, compute an integer d (for example, the denominator of the HNF of $\mathfrak{a}$) such that $d\mathfrak{a}$ is integral, and set $\mathfrak{a} \leftarrow d\mathfrak{a}$ and $\alpha \leftarrow \alpha/d$. Using [Coh0, Algorithm 4.8.17], compute a uniformizer π of $\mathfrak{p}^{-1}$ (see Corollary 1.2.10) and the valuation $v \leftarrow v_{\mathfrak{p}}(\mathfrak{a})$. If $v \neq 0$, set $\mathfrak{a} \leftarrow \pi^v\mathfrak{a}$, $\alpha \leftarrow \pi^{-v}\alpha$ (now $\mathfrak{a}$ is an integral ideal coprime to $\mathfrak{p}$).
2. [Find basis of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$] Using Algorithm 1.5.2 on the pseudo-basis $(\omega_i, \mathfrak{a}_i)$ of $\mathbb{Z}_L$, compute elements $\eta_i \in \mathbb{Z}_L$ such that $\left(\overline{\eta_i}\right)_{1\le i\le n}$ is a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$.
3. [Compute structure constants] Using Algorithm 1.3.2, compute $a \in \mathfrak{a}$ and $b \in \mathfrak{p}$ such that $a + b = 1$ (we now know that $a\alpha \in \mathbb{Z}_L$ and that $a\alpha \equiv \alpha \pmod{\mathfrak{p}\mathfrak{a}^{-1}\mathbb{Z}_L}$). Compute constants $a_{i,j} \in \mathbb{Z}_K$ such that
$$a\alpha\eta_i \equiv \sum_{1\le j\le n} a_{i,j}\eta_j \pmod{\mathfrak{p}\mathbb{Z}_L} \ .$$
4. [Solve system] By ordinary Gaussian elimination in the field $\mathbb{Z}_K/\mathfrak{p}$, find a nontrivial solution to the system of congruences $\sum_{1\le i\le n} a_{i,j}x_i \equiv 0 \pmod{\mathfrak{p}}$.
5. [Terminate] Set $\beta \leftarrow \sum_{1\le i\le n} x_i\eta_i$, set $\mathfrak{b} \leftarrow a\mathfrak{a}^{-1}$, output $(\beta, \mathfrak{b})$, and terminate the algorithm.

Proof. We must show that this algorithm is valid. Step 1 is standard and reduces $\mathfrak{a}$ to the case of an integral ideal coprime to $\mathfrak{p}$. Thus, there exist $a \in \mathfrak{a}$ and $b \in \mathfrak{p}$ such that $a + b = 1$, so $a \equiv 1 \pmod{\mathfrak{p}}$. Since $\mathfrak{a}\alpha \subset \mathfrak{P} \subset \mathbb{Z}_L$, we have $\alpha \in \mathfrak{a}^{-1}\mathbb{Z}_L$ and so $a\alpha \in \mathbb{Z}_L$ and $\alpha - a\alpha = b\alpha \in \mathfrak{p}\mathfrak{a}^{-1}\mathbb{Z}_L$, as claimed. Thus, $a\alpha\eta_i \in \mathbb{Z}_L$, and reducing modulo $\mathfrak{p}$ we can compute the constants $a_{i,j}$.

By definition, the matrix $(a_{i,j})$ is congruent modulo $\mathfrak{p}\mathbb{Z}_L$ (hence modulo $\mathfrak{p}$ since $a_{i,j} \in \mathbb{Z}_K$) to the matrix of multiplication by $a\alpha$ on the basis of the η_i. Thus, its determinant is congruent modulo $\mathfrak{p}$ to $\mathcal{N}_{L/K}(a\alpha)$. Since $a\alpha \in \mathfrak{a}\alpha \subset \mathfrak{P}$, we have $\mathfrak{p}^{f(\mathfrak{P}/\mathfrak{p})} = \mathcal{N}_{L/K}(\mathfrak{P}) \mid \mathcal{N}_{L/K}(\mathfrak{a}\alpha)$, hence $\det((a_{i,j})) \in \mathfrak{p}$.

It follows that the matrix $\left(\overline{a_{i,j}}\right)$ is singular in $\mathbb{Z}_K/\mathfrak{p}$.Hence there exists a nontrivial solution to the system of congruences of step 4. If β and $\mathfrak{b}$ are chosen as in step 5, we have $\alpha\beta\mathfrak{a}\mathfrak{b} = a\alpha\beta\mathbb{Z}_K \subset \mathfrak{p}\mathbb{Z}_L$, $\beta\mathfrak{b} = \beta a\mathfrak{a}^{-1} \subset \beta\mathbb{Z}_K \subset \mathbb{Z}_L$, and $\beta \notin \mathfrak{p}\mathbb{Z}_L$, so $\beta\mathfrak{b} \not\subset \mathfrak{p}\mathbb{Z}_L$ since $\mathfrak{b} = a\mathfrak{a}^{-1}$ and both a and $\mathfrak{a}$ are coprime to $\mathfrak{p}$. □

2.3.4 Operations on Ideals

To add two ideals, we use both HNFs and concatenate them to form an $n \times 2n$ pseudo-matrix M, and use one of the HNF algorithms in $\mathbb{Z}_K$ to compute the

HNF of M, which is the HNF of the ideal sum. Since the determinantal ideal of the individual ideals is easily computed, we can use any one of them in the modular HNF algorithm.

To multiply two ideals, we could use both HNFs, and form an $n \times (n^2)$ pseudo-matrix of basis element products, and HNF-reduce this matrix. This is costly, and it is better to represent one of the ideals by its HNF and the other one by a pseudo–two-element representation $((\alpha, \mathfrak{a}), (\beta, \mathfrak{b}))$. By multiplying α and β in L by each basis element of the HNF, and multiplying the corresponding ideals of the pseudo-matrix by $\mathfrak{a}$ and $\mathfrak{b}$, we obtain an $n \times 2n$ pseudo-matrix whose columns form a pseudo-generating set of the ideal product, and we can then obtain the HNF of this product by HNF-reduction. When n is large, this is much faster than the use of both HNF representations. Of course, to be able to use this method, we must be able to go back and forth between the two types of representations, and this is done using the methods explained in Section 2.3.1, especially Algorithm 2.3.8.

Remarks

(1) This method for computing ideal products is evidently also very useful in the absolute case and is not stressed enough in [Coh0] since a two-element representation was thought to be costly to compute at the time of writing. Practice has shown that this is not the case. In fact, it is not difficult to give complexity estimates for this problem, which show that finding a two-element representation is rather fast (see Algorithm 1.3.15).

(2) To compute an ideal product, one could also think of using *both* pseudo–two-element representations. This would involve computing only $2 \times 2 = 4$ products. Unfortunately, this is only superficially attractive since it then becomes costly to obtain either a pseudo–two-element representation or an HNF from this four-element representation. For example, to obtain the HNF, one would need to multiply each of the four products by the n pseudo-basis elements, and HNF-reducing an $n \times 4n$ pseudo-matrix. This is twice as expensive as the method that we have suggested. Indeed, the use of at least one of the HNFs for the ideals avoids extra multiplications by pseudo-basis elements (see Exercise 29). It seems, however, that suitably implemented, this method can be very slightly faster than the one we suggest for very large degrees (see [Hop]).

Another important algorithm not mentioned in [Coh0], but essential for many applications, is that of raising an ideal I to an integer power. We could, of course, use one of the binary powering algorithms, using the method explained above for ideal multiplication. There is, however, a much better method based on the following proposition, whose absolute counterpart was also not sufficiently stressed in [Coh0].

Proposition 2.3.15. *Let $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$ be a pseudo–two-element representation of an ideal I, and let k be a nonnegative integer. Then*

$$I^k = \alpha^k \mathfrak{a}^k \mathbb{Z}_L + \beta^k \mathfrak{b}^k \mathbb{Z}_L \ .$$

In the special case where $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$ *is a prime ideal, we even have*

$$\mathfrak{P}^k = \mathfrak{p}^s \mathbb{Z}_L + \alpha^k \mathfrak{a}^k \mathbb{Z}_L \quad \textit{with} \quad s = \left\lceil \frac{k}{e(\mathfrak{P}/\mathfrak{p})} \right\rceil \ .$$

Proof. The equality $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$ is equivalent to

$$v_{\mathfrak{P}}(I) = \min(v_{\mathfrak{P}}(\alpha\mathfrak{a}), v_{\mathfrak{P}}(\beta\mathfrak{b}))$$

for all prime ideals $\mathfrak{P}$ of L, and hence

$$v_{\mathfrak{P}}(I^k) = k v_{\mathfrak{P}}(I) = \min(v_{\mathfrak{P}}(\alpha^k \mathfrak{a}^k), v_{\mathfrak{P}}(\beta^k \mathfrak{b}^k)) \ ,$$

proving our first claim. In the case of a prime ideal $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$, we have $\min(v_{\mathfrak{Q}}(\mathfrak{p}), v_{\mathfrak{Q}}(\alpha\mathfrak{a})) = 0$ for any prime ideal $\mathfrak{Q}$ different from $\mathfrak{P}$, while $\min(v_{\mathfrak{P}}(\mathfrak{p}), v_{\mathfrak{P}}(\alpha\mathfrak{a})) = \min(e(\mathfrak{P}/\mathfrak{p}), v_{\mathfrak{P}}(\alpha\mathfrak{a})) = 1$.

It follows that $\min(v_{\mathfrak{Q}}(\mathfrak{p}^s), v_{\mathfrak{Q}}(\alpha^k\mathfrak{a}^k)) = 0$ for any such prime ideal $\mathfrak{Q}$ and any strictly positive s. For the prime ideal $\mathfrak{P}$, we may assume that $\mathfrak{P}$ is ramified; otherwise the result is not any stronger than the general claim. If $e(\mathfrak{P}/p) > 1$, we necessarily have $v_{\mathfrak{P}}(\alpha\mathfrak{a}) = 1$, and so

$$\min(v_{\mathfrak{P}}(\mathfrak{p}^s), v_{\mathfrak{P}}(\alpha^k\mathfrak{a}^k)) = \min(s \cdot e(\mathfrak{P}/p), k) = k$$

if and only if $s \geq k/e(\mathfrak{P}/\mathfrak{p})$, proving the proposition. □

Thus, to compute I^k, we first compute a pseudo–two-element representation $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$ using Algorithm 2.3.8. We then compute $I^k = \alpha^k\mathfrak{a}^k\mathbb{Z}_L + \beta^k\mathfrak{b}^k\mathbb{Z}_K$ (or $\mathfrak{P}^k = \mathfrak{p}^s\mathbb{Z}_L + \alpha^k\mathfrak{a}^k\mathbb{Z}_L$ in the case of a prime ideal) by a binary powering method. Finally, if desired we transform this into an HNF representation as usual by doing an HNF reduction of an $n \times 2n$ pseudo-matrix. Evidently this method is much less costly than the naive method, since we simply have to compute powers of two ideals *in the base field* K, and only powers of two *elements* of L.

Finally, consider the question of computing the *inverse* of an ideal. We proceed essentially as in [Coh0, Section 4.8.4].

Definition 2.3.16. *Let* L/K *be an extension of number fields. The* relative different $\mathfrak{D}(L/K)$ *is the ideal of* $\mathbb{Z}_L$ *defined as the inverse of the ideal* (*called the* relative codifferent)

$$\mathfrak{D}(L/K)^{-1} = \{x \in L, \ \mathrm{Tr}_{L/K}(x\mathbb{Z}_L) \subset \mathbb{Z}_K\} \ .$$

As in the absolute case, the relative different is an integral ideal of $\mathbb{Z}_L$ whose relative norm is the relative discriminant ideal $\mathfrak{d}(L/K)$, and the prime ideals that divide $\mathfrak{D}(L/K)$ are exactly the prime ideals of L that are ramified.

We will also need the following easy, but important, result.

Proposition 2.3.17 (Transitivity of the Different). *Let L/K be a relative extension of number fields, and let k be a subfield of K (for example, $k = \mathbb{Q}$). Then $\mathfrak{D}(L/k) = \mathfrak{D}(L/K)\mathfrak{D}(K/k)$.*

Proof. Let $\mathfrak{a}$ be an ideal of L. Using the transitivity of the trace, by definition of the codifferent, we have

$$\begin{aligned}
\mathfrak{a} \subset \mathfrak{D}(L/K)^{-1} &\iff \mathrm{Tr}_{L/K}(\mathfrak{a}) \subset \mathbb{Z}_K \\
&\iff \mathfrak{D}(K/k)^{-1}\,\mathrm{Tr}_{L/K}(\mathfrak{a}) \subset \mathfrak{D}(K/k)^{-1} \\
&\iff \mathrm{Tr}_{K/k}(\mathfrak{D}(K/k)^{-1}\,\mathrm{Tr}_{L/K}(\mathfrak{a})) \subset \mathbb{Z}_k \\
&\iff \mathrm{Tr}_{K/k}(\mathrm{Tr}_{L/K}(\mathfrak{D}(K/k)^{-1}\mathfrak{a})) \subset \mathbb{Z}_k \\
&\iff \mathrm{Tr}_{L/k}(\mathfrak{D}(K/k)^{-1}\mathfrak{a}) \subset \mathbb{Z}_k \\
&\iff \mathfrak{D}(K/k)^{-1}\mathfrak{a} \subset \mathfrak{D}(L/k)^{-1} \\
&\iff \mathfrak{a} \subset \mathfrak{D}(K/k)\mathfrak{D}(L/k)^{-1}\ .
\end{aligned}$$

It follows that $\mathfrak{D}(L/K)^{-1} = \mathfrak{D}(K/k)\mathfrak{D}(L/k)^{-1}$, proving the proposition. □

As we shall see in Theorem 2.5.1, an important consequence of this proposition is the transitivity property for relative discriminants.

The analog of [Coh0, Proposition 4.8.19] is the following.

Proposition 2.3.18. *Let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis of $\mathbb{Z}_L$, and let I be an ideal of $\mathbb{Z}_L$ given by a pseudo-matrix $(M, \mathfrak{c}_i)$, where the columns of $(M, \mathfrak{c}_i)$ give the coordinates of a pseudo-basis $(\gamma_i, \mathfrak{c}_i)$ on the ω_i.*

If $T = (\mathrm{Tr}_{L/K}(\omega_i\omega_j))$, the pseudo-matrix $((M^tT)^{-1}, \mathfrak{c}_i^{-1})$ represents a pseudo-basis of the ideal $I^{-1}\mathfrak{D}(L/K)^{-1}$ on the ω_i.

Proof. The proof is almost identical to the absolute case. By definition of M, the entry of row i and column j in M^tT is equal to $\mathrm{Tr}_{L/K}(\gamma_i\omega_j)$. If $V = (v_i)$ is a column vector with $v_i \in K$, then V belongs to the image of the pseudo-matrix $((M^tT)^{-1}, \mathfrak{c}_i^{-1})$ if and only if M^tTV is a vector (x_i) with $x_i \in \mathfrak{c}_i^{-1}$. This implies that for all i,

$$\mathrm{Tr}_{L/K}\left(\gamma_i\mathfrak{c}_i\left(\sum_j v_j\omega_j\right)\right) \subset \mathbb{Z}_K\ ,$$

hence that $\mathrm{Tr}_{L/K}(xI) \subset \mathbb{Z}_K$ with $x = \sum_j v_j\omega_j$. Since $xI = xI\mathbb{Z}_L$, the proposition follows. Note that, in the same way that $\mathfrak{D}(L/K)^{-1}$ is the dual of $\mathbb{Z}_L$ for the trace form, the ideal $I^{-1}\mathfrak{D}(L/K)^{-1}$ is the dual of the ideal I for the trace form. □

The analog of [Coh0, Algorithm 4.8.21] is thus as follows.

Algorithm 2.3.19 (Ideal Inversion). Given a relative integral pseudo-basis $(\omega_i, \mathfrak{a}_i)$ of $\mathbb{Z}_L$ and an integral ideal I of $\mathbb{Z}_L$ given by an $n \times n$ pseudo-matrix $(M, \mathfrak{c}_i)$ whose columns give the coordinates of a relative pseudo-basis $(\gamma_i, \mathfrak{c}_i)$ of I on the ω_i, this algorithm computes the HNF of the inverse ideal I^{-1}.

1. [Compute $\mathfrak{d}(L/K)\mathfrak{D}(L/K)^{-1}$] Compute the $n \times n$ matrix $T = (\mathrm{Tr}_{L/K}(\omega_i\omega_j))$. Let $\mathfrak{d} \leftarrow \det(T)\prod_{1\le i\le n}\mathfrak{a}_i^2$ (this is the relative ideal-determinant $\mathfrak{d}(L/K)$ of L and hence is usually available with the ω_i). Finally, call δ_j the elements of L whose coordinates on the ω_i are the columns of T^{-1} (thus, $(\delta_j, \mathfrak{d}\mathfrak{a}_j^{-1})$ will be a pseudo-basis of the integral ideal $\mathfrak{d}(L/K)\mathfrak{D}(L/K)^{-1}$).
2. [Find a pseudo–two-element representation] Using Algorithm 2.3.8, compute a pseudo–two-element representation $((\alpha, \mathfrak{a}), (\beta, \mathfrak{b}))$ of $\mathfrak{d}(L/K)\mathfrak{D}(L/K)^{-1}$ corresponding to the pseudo-basis $(\delta_j, \mathfrak{d}\mathfrak{a}_j^{-1})$ computed in step 1.
3. [Compute $\mathfrak{d}(L/K)\mathfrak{D}(L/K)^{-1}I$] Let $(N, \mathfrak{b}_i)$ be the HNF of the $n \times 2n$ pseudo-matrix whose columns are the coordinates on the integral basis of the products $\gamma_i\alpha$ and $\gamma_i\beta$ with corresponding ideals $\mathfrak{c}_i\mathfrak{a}$ and $\mathfrak{c}_i\mathfrak{b}$ (this will be a pseudo-basis of $\mathfrak{d}(L/K)\mathfrak{D}(L/K)^{-1}I$).
4. [Compute I^{-1}] Set $P \leftarrow (N^tT)^{-1}$. Output the HNF of the pseudo-matrix $(P, \mathfrak{d}\mathfrak{b}_i^{-1})$ and terminate the algorithm.

The proof of this algorithm's validity is left to the reader. We have included in this algorithm one of the remarks made after [Coh0, Algorithm 4.8.21] to speed up step 3. The other remarks are also applicable here. In particular, the computations in steps 1 and 2 are independent of the ideal I.

There exists a completely different and faster method for computing ideal inverses, both in the absolute and in the relative case, that can be used when a two-element representation is known (as we have seen in Algorithm 2.3.8, this is in general quite easy to compute). It is based on the following easy lemma.

Lemma 2.3.20. *Let $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$ be a pseudo–two-element representation of an ideal of L. Then*

$$I^{-1} = (\alpha^{-1}\mathfrak{a}^{-1}\mathbb{Z}_L) \cap (\beta^{-1}\mathfrak{b}^{-1}\mathbb{Z}_L) \ .$$

Proof. Indeed, by looking at valuations, it is clear that for any two ideals I and J in a Dedekind domain we have $(I + J)(I \cap J) = IJ$ (which is the generalization to ideals of the formula $\gcd(a, b)\,\mathrm{lcm}(a, b) = ab$). Applying this to the two ideals $\alpha^{-1}\mathfrak{a}^{-1}\mathbb{Z}_L$ and $\beta^{-1}\mathfrak{b}^{-1}\mathbb{Z}_L$ of $\mathbb{Z}_L$ and multiplying by $\alpha\beta\mathfrak{a}\mathfrak{b}$ immediately gives the desired result. □

The main operation which must be done is thus ideal intersection, which is computed using Algorithm 1.5.1. This gives the following algorithm.

Algorithm 2.3.21 (Ideal Inversion). Given a fractional ideal I of L, this algorithm computes the HNF of the inverse ideal I^{-1}.

1. [Compute two-element] If not already in this form, using Algorithm 2.3.8 compute a pseudo–two-element representation $I = \alpha\mathfrak{a}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$.
2. [Compute inverses] Using the present algorithm in the absolute case, using true two-element representations of $\mathfrak{a}$ and $\mathfrak{b}$, compute $I_1 \leftarrow (\mathfrak{a}^{-1}/\alpha)\mathbb{Z}_L$ and $I_2 \leftarrow (\mathfrak{b}^{-1}/\beta)\mathbb{Z}_L$.
3. [Compute intersection] Using Algorithm 1.5.1, compute an HNF pseudo-basis $(\gamma_i, \mathfrak{c}_i)$ for the intersection $I_3 \leftarrow I_1 \cap I_2$.
4. [Terminate] For each i set $\alpha_i \leftarrow \alpha\beta\gamma_i$ and $\mathfrak{a}_i \leftarrow \mathfrak{a}\mathfrak{b}\mathfrak{c}_i$, output the HNF of the pseudo-basis $(\alpha_i, \mathfrak{a}_i)$ of I^{-1}, and terminate the algorithm.

2.3.5 Ideal Factorization and Ideal Lists

For future use, we describe here some algorithms for computing ideal factorizations and ideal lists. These algorithms have nothing to do with relative extensions and could have been included in [Coh0].

The following is an algorithm for computing prime ideal factorizations, whose proof is immediate.

Algorithm 2.3.22 (Ideal Factorization). Let K be a number field and I be a fractional ideal of K. This algorithm computes the prime ideal factorization $I = \prod_{\mathfrak{p}} \mathfrak{p}^{v_{\mathfrak{p}}}$ of I.

1. [Remove denominator] Let d be a positive integer such that $dI = J$ is an integral ideal ($d = 1$ if I was already integral). If $d \neq 1$, apply recursively this algorithm to the ideals $d\mathbb{Z}_K$ and J, output the prime ideal factorization of $I = Jd^{-1}$ by subtraction of the exponents in the factorization of J and d, and terminate the algorithm.
2. [Compute $N = \mathcal{N}(J)$] (Here $J = I$ is an integral ideal.) If J is not given in HNF, perform an HNF reduction to reduce to that case, and let H be the HNF matrix of J. Let N be the product of the diagonal entries of H (so $N = \det(H) = \mathcal{N}(J)$).
3. [Factor N] Factor N as $N = \prod_p p^{a_p}$, with $a_p \geq 0$.
4. [Compute prime ideals] Using [Coh0, Algorithm 6.2.9], for each prime p such that $a_p > 0$, compute the prime ideal decomposition of $p\mathbb{Z}_K$ as $p\mathbb{Z}_K = \prod_{\mathfrak{p}|p} \mathfrak{p}^{e_{\mathfrak{p}}}$.
5. [Compute valuations] For each p such that $a_p > 0$, and each $\mathfrak{p} \mid p$ found in step 4, use [Coh0, Algorithm 4.8.17] to compute $v_{\mathfrak{p}} \leftarrow v_{\mathfrak{p}}(J)$. Output the factorization $I = \prod_{\mathfrak{p}} \mathfrak{p}^{v_{\mathfrak{p}}}$ and terminate the algorithm.

Remarks

(1) It is essential to remove the denominator of I in step 1, since otherwise the factorization of the norm would not include all the possible $\mathfrak{p}$ (for example, if $\mathfrak{p}$ and $\mathfrak{p}'$ are prime ideals of the same residual degree over the same prime p, $\mathfrak{p}'\mathfrak{p}^{-1}$ is an ideal of norm 1).

(2) We can speed up the computation in step 5 by noticing that the factorization of the norm gives us the equality $a_p = \sum_{\mathfrak{p}|p} v_{\mathfrak{p}}(J)f(\mathfrak{p}/p)$. Since all the coefficients are nonnegative, as soon as this equality is achieved, we know that all the other $\mathfrak{p}$ above p (if any remain) will not divide J.

We will also need to find the list of integral ideals of K of norm less than or equal to some bound B, and perhaps satisfying some additional conditions. Although easily done, there are some slightly subtle tricks involved. The basic algorithm where no conditions are imposed is the following.

Algorithm 2.3.23 (Ideal List). Let K be a number field and B be a positive integer. This algorithm outputs a list of lists $\mathcal{L}$ such that for each $n \leq B$, $\mathcal{L}_n$ is the list of all integral ideals of absolute norm equal to n.

1. [Initialize] For $2 \leq n \leq B$ set $\mathcal{L}_n \leftarrow \emptyset$, then set $\mathcal{L}_1 \leftarrow \{\mathbb{Z}_K\}$ and $p \leftarrow 0$.
2. [Next prime] Replace p by the smallest prime strictly larger than p. If $p > B$, output $\mathcal{L}$ and terminate the algorithm.
3. [Factor $p\mathbb{Z}_K$] Using [Coh0, Algorithm 6.2.9], factor $p\mathbb{Z}_K$ as $p\mathbb{Z}_K = \prod_{1\leq i\leq g} \mathfrak{p}_i^{e_i}$ with $e_i \geq 1$, and let $f_i = f(\mathfrak{p}_i/p)$. Set $j \leftarrow 0$.
4. [Next prime ideal] Set $j \leftarrow j + 1$. If $j > g$, go to step 2. Otherwise, set $q \leftarrow p^{f_j}$, $n \leftarrow 0$.
5. [Loop through all multiples of q] Set $n \leftarrow n + q$. If $n > B$, go to step 4. Otherwise, set $\mathcal{L}_n \leftarrow \mathcal{L}_n \cup \mathfrak{p}_j\mathcal{L}_{n/q}$, where $\mathfrak{p}_j\mathcal{L}_{n/q}$ is the list of products by the ideal $\mathfrak{p}_j$ of the elements of $\mathcal{L}_{n/q}$ and go to step 5.

Remark. The only subtle point of this algorithm is step 5. Since we loop by *increasing* multiples n of q, if $\mathcal{L}'$ denotes the list at the end of step 4, then step 5 is equivalent to setting

$$\mathcal{L}_n \leftarrow \bigcup_{1\leq k\leq v_q(n)} \mathfrak{p}_j^k \mathcal{L}'_{n/q^k}$$

for all n.

In the sequel, we will need two modifications of this algorithm. In the first modification, we want only *squarefree* ideals $\mathfrak{a}$ — in other words ideals whose prime ideal factorization has only exponents 0 or 1. To do this, it is sufficient in step 5 to loop through the multiples of q in *decreasing* order, giving the following algorithm.

Algorithm 2.3.24 (Squarefree Ideal List). Let K be a number field and B be a positive integer. This algorithm outputs a list of lists $\mathcal{L}$ such that for each $n \leq B$, $\mathcal{L}_n$ is the list of all squarefree integral ideals of absolute norm equal to n.

1. [Initialize] For $2 \leq n \leq B$ set $\mathcal{L}_n \leftarrow \varnothing$, then set $\mathcal{L}_1 \leftarrow \{\mathbb{Z}_K\}$ and $p \leftarrow 0$.
2. [Next prime] Replace p by the smallest prime strictly larger than p. If $p > B$, output $\mathcal{L}$ and terminate the algorithm.
3. [Factor $p\mathbb{Z}_K$] Using [Coh0, Algorithm 6.2.9], factor $p\mathbb{Z}_K$ as $p\mathbb{Z}_K = \prod_{1 \leq i \leq g} \mathfrak{p}_i^{e_i}$ with $e_i \geq 1$, and let $f_i = f(\mathfrak{p}_i/p)$. Set $j \leftarrow 0$.
4. [Next prime ideal] Set $j \leftarrow j+1$. If $j > g$, go to step 2. Otherwise, set $q \leftarrow p^{f_j}$ and set $n \leftarrow q(\lfloor B/q \rfloor + 1)$.
5. [Loop through multiples of q] Set $n \leftarrow n-q$. If $n < 1$, go to step 4. Otherwise, set $\mathcal{L}_n \leftarrow \mathcal{L}_n \cup \mathfrak{p}_j \mathcal{L}_{n/q}$ and go to step 5.

Let ℓ be a fixed prime. In the second more technical modification, we need the list of ideals that we will call "conductors at ℓ" for reasons we will see in class field theory, that is, ideals $\mathfrak{a}$ such that $v_{\mathfrak{p}}(\mathfrak{a}) = 0$ or 1 for all $\mathfrak{p} \nmid \ell$, while $2 \leq v_{\mathfrak{p}}(\mathfrak{a}) \leq \lfloor \ell e(\mathfrak{p}/\ell)/(\ell-1) + 1 \rfloor$ if $\mathfrak{p} \mid \ell$. For this, we must loop differently depending on whether or not $\mathfrak{p} \mid \ell$, giving the following algorithm.

Algorithm 2.3.25 (Conductor at ℓ Ideal List). Let K be a number field, let ℓ be a prime number, and let B be a positive integer. This algorithm outputs a list of lists $\mathcal{L}$ such that for each $n \leq B$, $\mathcal{L}_n$ is the list of all integral ideals of absolute norm equal to n which are conductors at ℓ in the above sense.

1. [Initialize] For $2 \leq n \leq B$ set $\mathcal{L}_n \leftarrow \varnothing$, then set $\mathcal{L}_1 \leftarrow \{\mathbb{Z}_K\}$ and $p \leftarrow 0$.
2. [Next prime] Replace p by the smallest prime strictly larger than p. If $p > B$, output $\mathcal{L}$ and terminate the algorithm.
3. [Factor $p\mathbb{Z}_K$] Using [Coh0, Algorithm 6.2.9], factor $p\mathbb{Z}_K$ as $p\mathbb{Z}_K = \prod_{1 \leq i \leq g} \mathfrak{p}_i^{e_i}$ with $e_i \geq 1$, and let $f_i = f(\mathfrak{p}_i/p)$. Set $j \leftarrow 0$.
4. [Next prime ideal] Set $j \leftarrow j+1$. If $j > g$, go to step 2. Otherwise, set $q \leftarrow p^{f_j}$, set $q_1 \leftarrow q$ if $p \neq \ell$, $q_1 \leftarrow q^2$ if $p = \ell$, and set $n \leftarrow q_1(\lfloor B/q_1 \rfloor + 1)$.
5. [Loop through multiples of q_1] Set $n \leftarrow n - q_1$. If $n < 1$, go to step 4. Otherwise, do as follows. If $p \neq \ell$, set $\mathcal{L}_n \leftarrow \mathcal{L}_n \cup \mathfrak{p}_j \mathcal{L}_{n/q}$. On the other hand, if $p = \ell$, set
$$k_s \leftarrow \lfloor \min(v_\ell(n)/f_j, \ell e_j/(\ell-1) + 1) \rfloor$$
and
$$\mathcal{L}_n \leftarrow \mathcal{L}_n \cup \bigcup_{2 \leq k \leq k_s} \mathfrak{p}_j^k \mathcal{L}_{n/q^k} \ ,$$
where $\mathfrak{p}_j^k \mathcal{L}_{n/q^k}$ is the list of products by the ideal $\mathfrak{p}_j^k$ of the elements of $\mathcal{L}_{n/q^k}$, and go to step 5.

2.4 The Relative Round 2 Algorithm and Related Algorithms

Let L/K be a relative extension. The ring of integers $\mathbb{Z}_L$ of L is a $\mathbb{Z}_K$-module in a natural way, and thus we want to compute $\mathbb{Z}_L$ as a $\mathbb{Z}_K$-module, not only as a $\mathbb{Z}$-module. This has three advantages. First, the $\mathbb{Z}_K$-module structure is richer than the $\mathbb{Z}$-module structure. Second, as we shall see, it is much easier to compute the $\mathbb{Z}_K$-module structure than the $\mathbb{Z}$-module structure, since the relative degree $n = [L : K]$ is much smaller than the absolute degree. Finally, if the $\mathbb{Z}$-module structure is really desired, it is trivial to obtain it from the $\mathbb{Z}_K$-module structure (see Section 2.5.1).

In case a number field L is given as a relative extension, this allows us to compute integral bases and discriminants for much larger degrees than would be possible otherwise. For example, in [Dab1] such computations are made for an extension of degree 33 of a base field of degree 32, whereas directly computing the discriminant of an absolute number field of degree $32 \times 33 = 1056$ is almost impossible using current algorithms.

2.4.1 The Relative Round 2 Algorithm

In this section, we explain how to generalize the round 2 algorithm to the relative case. We assume that we know everything needed about the base ring $\mathbb{Z}_K$, and we must find an algorithm for computing a $\mathbb{Z}_K$-pseudo-basis of $\mathbb{Z}_L$, in other words a relative integral pseudo-basis. We set $R = \mathbb{Z}_K$, $n = [L : K]$, $m = [K : \mathbb{Q}]$, and we assume that L is given as $L = K(\theta)$ for some algebraic integer θ whose minimal monic polynomial over K is denoted by $T(X) \in \mathbb{Z}_K[X]$.

We follow closely the exposition given in [Coh0, Chapter 6].

Definition 2.4.1. *Let $\mathcal{O} \subset \mathbb{Z}_L$ be an order in L and let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$.*

(1) *We will say that $\mathcal{O}$ is* $\mathfrak{p}$-maximal *if the order-ideal (see Definition 1.2.33) of the torsion module $\mathbb{Z}_L/\mathcal{O}$ is not divisible by the ideal $\mathfrak{p}$ or, equivalently, if $\mathfrak{p}\mathbb{Z}_L + \mathcal{O} = \mathbb{Z}_L$.*

(2) *We define the* $\mathfrak{p}$-radical $I_{\mathfrak{p}}$ *of $\mathcal{O}$ as follows:*

$$I_{\mathfrak{p}} = \{x \in \mathcal{O} \mid \exists m \geq 1 \text{ such that } x^m \in \mathfrak{p}\mathcal{O}\} \ .$$

Then, as in the absolute case, it is easy to prove that $I_{\mathfrak{p}}$ is an ideal of $\mathcal{O}$ equal to the product of all distinct prime ideals of $\mathcal{O}$ lying above $\mathfrak{p}$. To compute $I_{\mathfrak{p}}$ (or more precisely $I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$) explicitly, we may use the following proposition, which is also proved as in the absolute case.

Proposition 2.4.2. *Let $q = \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{p}) = |\mathbb{Z}_K/\mathfrak{p}|$ and let $j \geq 1$ be such that $q^j \geq n$, where $n = [L : K]$. Then $I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$ is the kernel of the $\mathbb{Z}_K/\mathfrak{p}$-linear map $x \mapsto x^{q^j}$ from $\mathcal{O}/\mathfrak{p}\mathcal{O}$ into itself.*

When $\mathfrak{p}$ is large, however, there is a more efficient method for computing the $\mathfrak{p}$-radical based on the following proposition, which should have been included in [Coh0] for the absolute case.

Proposition 2.4.3. *Let p be the prime number below $\mathfrak{p}$, assume that $p > n = [L : K]$, and let $\alpha \in \mathcal{O}$. The following three properties are equivalent.*

(1) $\alpha \in I_{\mathfrak{p}}$.
(2) *The characteristic polynomial $C_\alpha(X)$ of α over K satisfies $C_\alpha(X) \equiv X^n \pmod{\mathfrak{p}}$.*
(3) *For all $\beta \in \mathcal{O}$ we have $\mathrm{Tr}_{L/K}(\alpha\beta) \in \mathfrak{p}$.*

Proof. The proof that (1) implies (2) is the same as the proof of Proposition 2.4.2 (see [Coh0], Lemma 6.1.6): if $\alpha \in I_{\mathfrak{p}}$, multiplication by α induces a nilpotent map from the $\mathbb{Z}_K/\mathfrak{p}$-vector space $\mathcal{O}/\mathfrak{p}\mathcal{O}$ to itself, hence its eigenvalues are all equal to 0, so its characteristic polynomial is equal to X^n modulo $\mathfrak{p}$. Conversely, (2) implies (1) by the Cayley–Hamilton theorem (note that the equivalence of (1) and (2) does not use the condition $p > n$).

Assume now that $\alpha \in I_{\mathfrak{p}}$. Then for all $\beta \in \mathcal{O}$ we have $\alpha\beta \in I_{\mathfrak{p}}$; hence by what we have just proved, $C_{\alpha\beta}(X) \equiv X^n \pmod{\mathfrak{p}}$, and in particular $\mathrm{Tr}_{L/K}(\alpha\beta) \in \mathfrak{p}$.

Conversely, assume (3). If we apply (3) to $\beta = \alpha^{k-1}$ for $k \geq 1$, we deduce that $\mathrm{Tr}_{L/K}(\alpha^k) \in \mathfrak{p}$ for all $k \geq 1$. Let $C_\alpha(X) = X^n + \sum_{j=1}^{n}(-1)^j a_j X^{n-j}$ be the characteristic polynomial of α, where the $a_j \in \mathbb{Z}_K$ are the elementary symmetric functions of α. Newton's relations between elementary symmetric functions and sums of powers give the recursion

$$k a_k = \sum_{j=1}^{k}(-1)^{j-1} a_{k-j}\,\mathrm{Tr}_{L/K}(\alpha^j) \ .$$

Since $\mathrm{Tr}_{L/K}(\alpha^j) \in \mathfrak{p}$ for $j \geq 1$, it follows by induction that $a_k \in \mathfrak{p}$ for $k < p$, since $k < p$ implies $k \notin \mathfrak{p}$. Since p has been assumed to be larger than n, it follows that $a_k \in \mathfrak{p}$ for $1 \leq k \leq n$, proving (2). □

Condition (3) can easily be transformed into an algorithm for computing the $\mathfrak{p}$-radical as follows. Let $(\overline{\omega_i})_{1\leq i\leq n}$ be a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathcal{O}/\mathfrak{p}\mathcal{O}$. It is clear that (3) is equivalent to $\mathrm{Tr}_{L/K}(\alpha\omega_i) \in \mathfrak{p}$ for $1 \leq i \leq n$. Thus, if we write $\alpha \equiv \sum_{1\leq j\leq n} x_j\omega_j \pmod{\mathfrak{p}}$, we obtain the following linear system in $\mathbb{Z}_K/\mathfrak{p}$:

$$\sum_{1\leq i\leq n} x_i\,\mathrm{Tr}_{L/K}(\omega_i\omega_j) \equiv 0 \pmod{\mathfrak{p}} \ ;$$

hence $I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$ is the kernel of the matrix $(\mathrm{Tr}_{L/K}(\omega_i\omega_j))$ over $\mathbb{Z}_K/\mathfrak{p}$, which can easily be found using Gaussian elimination. If p is large, the resulting computation will be much shorter than the computation based on Proposition

2.4.2. We leave the details of the resulting algorithm to the reader (Exercise 32).

Zassenhaus's theorem, being a local statement, goes through without change:

Proposition 2.4.4. *Set*

$$\mathcal{O}' = \{x \in L \mid xI_{\mathfrak{p}} \subset I_{\mathfrak{p}}\} \ .$$

Then

(1) $\mathcal{O}'$ *is an order in* L *containing* $\mathcal{O}$*;*
(2) $\mathcal{O}' = \mathcal{O}$ *if and only if* $\mathcal{O}$ *is* $\mathfrak{p}$*-maximal;*
(3) *if* $\mathcal{O}' \neq \mathcal{O}$*, then the order-ideal of* $\mathcal{O}'/\mathcal{O}$ *is equal to* $\mathfrak{p}^k$ *for some* k *such that* $1 \leq k \leq n$.

To compute $\mathcal{O}'$ algorithmically, we use the following proposition, whose proof is immediate.

Proposition 2.4.5. *Let* U *be the kernel of the* $\mathbb{Z}_K$*-linear map*

$$\alpha \longmapsto \left(\overline{\beta} \mapsto \overline{\alpha\beta}\right)$$

from $\mathcal{O}$ *to* $\mathrm{End}_{\mathbb{Z}_K/\mathfrak{p}}(I_{\mathfrak{p}}/\mathfrak{p}I_{\mathfrak{p}})$*; then* $\mathcal{O}' = \mathfrak{p}^{-1}U$.

Finally, we must explain how to find a $\mathbb{Z}_K/\mathfrak{p}$-basis of $I_{\mathfrak{p}}/\mathfrak{p}I_{\mathfrak{p}}$ knowing one of $I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$. There are slight differences with the absolute case, so we give all the details.

As in [Coh0, Chapter 6], let $\beta_1, \ldots, \beta_l$ in $I_{\mathfrak{p}}$ be such that $\left(\overline{\beta_i}\right)_{1 \leq i \leq l}$ is a basis of $I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$ as a $\mathbb{Z}_K/\mathfrak{p}$-vector space. Using Algorithm 1.5.2, from a given pseudo-basis $(\omega_i, \mathfrak{a}_i)$ of $\mathcal{O}$, we may compute a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathcal{O}/\mathfrak{p}\mathcal{O}$. Thus, we can use [Coh0, Algorithm 2.3.6] to supplement the $\overline{\beta_i}$ with $\beta_{l+1}, \ldots, \beta_n$ so that the $\left(\overline{\beta_i}\right)_{1 \leq i \leq n}$ form a basis of $\mathcal{O}/\mathfrak{p}\mathcal{O}$.

Let $\pi \in \mathfrak{p} \setminus \mathfrak{p}^2$, so that π is an element of $\mathbb{Z}_K$ whose valuation at $\mathfrak{p}$ is exactly equal to 1 (if $\mathfrak{p} = p\mathbb{Z}_K + \alpha\mathbb{Z}_K$, then π can be taken to be either α or $\alpha + p$). I claim that if we set $\alpha_i = \beta_i$ for $1 \leq i \leq l$ and $\alpha_i = \pi\beta_i$ for $l + 1 \leq i \leq n$, then $\left(\overline{\alpha_i}\right)_{1 \leq i \leq n}$ is a basis of $I_{\mathfrak{p}}/\mathfrak{p}I_{\mathfrak{p}}$ (by abuse of notation we will write $\overline{x}$ for the reduction of x mod $\mathfrak{p}$, $\mathfrak{p}\mathcal{O}$, or $\mathfrak{p}I_{\mathfrak{p}}$).

First, it is clear that $\alpha_i \in I_{\mathfrak{p}}$ for all i, since $\pi \in \mathfrak{p}$ and $\mathfrak{p}\mathcal{O} \subset I_{\mathfrak{p}}$. Furthermore, $I_{\mathfrak{p}}/\mathfrak{p}I_{\mathfrak{p}}$ is a $\mathbb{Z}_K/\mathfrak{p}$-vector space of dimension n. Hence we must simply prove that the $\overline{\alpha_i}$ are $\mathbb{Z}_K/\mathfrak{p}$-linearly independent. So assume that there exist $a_i \in \mathbb{Z}_K$ such that $\sum_{i=1}^{n} \overline{a_i}\overline{\alpha_i} = \overline{0}$ or, equivalently,

$$\sum_{i=1}^{n} a_i\alpha_i \in \mathfrak{p}I_{\mathfrak{p}} \ .$$

Since $I_{\mathfrak{p}} \subset \mathcal{O}$ and $\pi \in \mathfrak{p}$, this implies that $\sum_{i=1}^{l} a_i\beta_i \in \mathfrak{p}\mathcal{O}$; hence $a_i \in \mathfrak{p}$ for $1 \leq i \leq l$, since the $\overline{\beta_i}$ are $\mathbb{Z}_K/\mathfrak{p}$-linearly independent.

Hence we have

$$\pi \sum_{i=l+1}^{n} a_i\beta_i \in \mathfrak{p}I_{\mathfrak{p}} \ .$$

Since $\pi \notin \mathfrak{p}^2$, we can write $\pi\mathbb{Z}_K = \mathfrak{p}\mathfrak{b}$ for some ideal $\mathfrak{b}$ prime to $\mathfrak{p}$. Let $u \in \mathfrak{b}$ such that $u \notin \mathfrak{p}$. Then $(u/\pi)\mathfrak{p} \subset \mathbb{Z}_K$. Multiplying our relation by u/π, we obtain

$$u \sum_{l+1 \leq i \leq n} a_i\beta_i \in (u/\pi)\mathfrak{p}I_{\mathfrak{p}} \subset I_{\mathfrak{p}} \ ;$$

hence $\sum_{l+1 \leq i \leq n} \overline{a_i\beta_i} \in I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$ since $\overline{u}$ is invertible in $\mathbb{Z}_K/\mathfrak{p}$. But since the subspace generated by the $\overline{\beta_i}$ for $l+1 \leq i \leq n$ is in direct sum with $I_{\mathfrak{p}}/\mathfrak{p}\mathcal{O}$, this implies that $a_i \in \mathfrak{p}$ also for $i \geq l+1$, thus proving our claim. □

Finally, an important point must be clarified. Most of the computations are done in the residue field $\mathbb{Z}_K/\mathfrak{p}$. Since this is not simply $\mathbb{Z}/p\mathbb{Z}$, we must explain how elements are represented. The way that we have chosen is based on the following proposition.

Proposition 2.4.6. *Let $(\omega_1, \ldots, \omega_m)$ be a $\mathbb{Z}$-integral basis of a number field K, let $\mathfrak{p}$ be a prime ideal of degree f of $\mathbb{Z}_K$, and let $A = (a_{i,j})_{1 \leq i,j \leq m}$ be its Hermite normal form on the integral basis. Let D_1 (resp., D_p) be the set of indices $i \in [1, m]$ such that $a_{i,i} = 1$ (resp., $a_{i,i} = p$). Then*

(1) *$|D_p| = f$ and $|D_1| = m - f$;*
(2) *if $i \in D_1$, then $a_{i,j} = 0$ for $j > i$ (each off-diagonal entry of* row *i is equal to zero);*
(3) *if $j \in D_p$, then $a_{i,j} = 0$ for $i < j$ (each off-diagonal entry of* column *j is equal to zero).*

Proof. Call α_j the HNF basis elements of $\mathfrak{p}$ given by the matrix A. The determinant of A is equal to the index of $\mathfrak{p}$ in $\mathbb{Z}_K$, hence is equal to $\mathcal{N}(\mathfrak{p}) = p^f$, so the diagonal entries of the HNF matrix must be powers of p. Assume that $a_{j,j} = p^k$. Since $p\mathbb{Z}_K \subset \mathfrak{p}$, we have $p\omega_j = \sum_{1 \leq i \leq m} x_i\alpha_i$ for some integers x_i. Since the matrix A is triangular, we deduce that $x_i = 0$ for $i > j$, and in addition $p = x_j a_{j,j} = x_j p^k$, and it follows that $k = 0$ or $k = 1$, so the diagonal entries are equal to 1 or p. Since the determinant is equal to p^f, we have f diagonal entries equal to p, proving (1). (2) is a trivial consequence of the definition of the HNF.

For (3), let j be such that $a_{j,j} = p$. Then

$$\sum_{i<j} a_{i,j}\omega_i = \alpha_j - p\omega_j \in \mathfrak{p} \ .$$

Let i_0 be the largest index i in the sum (if it exists) such that $a_{i,j} \neq 0$. Thus,

$$\beta = a_{i_0,j}\omega_{i_0} + \sum_{i<i_0} a_{i,j}\omega_i \in \mathfrak{p} \ .$$

Since $a_{i_0,j} \neq 0$, by definition of the HNF we have $a_{i_0,i_0} > 1$, so $a_{i_0,i_0} = p$. But then, once again writing $\beta = \sum x_i\alpha_i$, we obtain $x_i = 0$ for $i > i_0$ and $x_{i_0}a_{i_0,i_0} = x_{i_0}p = a_{i_0,j}$. Since $0 \leq a_{i_0,j} < p$, this implies that $a_{i_0,j} = 0$, a contradiction. It follows that $\alpha_j = p\omega_j$, which is (3). □

Corollary 2.4.7. *Keep the notation of the preceding proposition. The classes modulo* $\mathfrak{p}$ *of the* ω_i *for* $i \in D_p$ *form an* $\mathbb{F}_p$*-basis of* $\mathbb{Z}_K/\mathfrak{p}$.

Proof. Since $|D_p| = f = \dim_{\mathbb{F}_p}(\mathbb{Z}_K/\mathfrak{p})$, we must simply show that the classes $\overline{\omega_i}$ for $i \in D_p$ are $\mathbb{F}_p$-linearly independent. Assume that $\sum_{i\in D_p} \overline{x_i\omega_i} = \overline{0}$, in other words that $\sum_{i\in D_p} x_i\omega_i \in \mathfrak{p}$, where we can assume that $0 \leq x_i < p$. Using the same method as in the proof of the proposition, letting i_0 be the largest index $i \in D_p$ (if it exists) such that $x_i \neq 0$, the triangular form of the HNF implies that $a_{i_0,i_0} \mid x_{i_0}$. Since $0 \leq x_{i_0} < p$ and $a_{i_0,i_0} = p$, we have $x_{i_0} = 0$, which is absurd, proving the corollary. □

We will thus represent an element of $\mathbb{Z}_K/\mathfrak{p}$ as an m-tuple $v = (v_i)$ of elements of $\mathbb{Z}/p\mathbb{Z}$, where $v_i = \overline{0}$ when $i \in D_1$. If $x \in \mathbb{Z}_K$ is represented as an m-tuple on the integral basis, we can then reduce x modulo $\mathfrak{p}$ by subtracting x_iA_i to x for each $i \in D_1$ (where A_i is the ith column of A) and reducing x_i mod p for all other i. Since our HNF matrices are upper-triangular, the subtraction of the x_iA_i must be done from bottom up. Also, by (3) above, we may reduce x_i modulo p for $i \in D_p$ without subtracting a multiple of A_i from x since all the off-diagonal entries of A_i are equal to zero. Note that this is a special case of Algorithm 1.4.12.

We will also want to reduce elements that are not in $\mathbb{Z}_K$ but in $S^{-1}\mathbb{Z}_K$, where $S = \mathbb{Z}_K \setminus \mathfrak{p}$. In this case the above procedure may not work since some of the x_i can have a denominator divisible by p (even though x itself is in $S^{-1}\mathbb{Z}_K$). In that case, we proceed as follows. Using Algorithm 1.3.2, we compute an element α such that $\alpha \equiv 1 \pmod{\mathfrak{p}}$ and $\alpha \in p/\mathfrak{p}^{e(\mathfrak{p}/p)}$. If k is the largest exponent of p appearing in the coefficients of x, it is clear that $x\alpha^k \in \mathbb{Z}_K$ and $x\alpha^k \equiv x \pmod{\mathfrak{p}}$ so we may apply the reduction procedure to $x\alpha^k$ instead of x. It is clear that the result is independent of the choice of α.

Finally, the Dedekind criterion ([Coh0, Theorem 6.1.4]) can also be easily generalized as follows.

Theorem 2.4.8. *Let* L/K *be a relative extension, with* $L = K(\theta)$ *and* θ *an algebraic integer whose minimal monic polynomial in* $K[X]$ *is denoted* $T(X)$*; let* $\mathfrak{p}$ *be a prime ideal of* $\mathbb{Z}_K$*, and let* β *be a uniformizer of* $\mathfrak{p}^{-1}$*, so that* $\beta \in \mathbb{Z}_K \setminus \mathfrak{p}^{-1}$.

Let $\overline{T(X)} = \prod_{1 \le i \le k} \overline{T_i(X)}^{e_i}$ *be the factorization of* $\overline{T(X)}$ *in* $(\mathbb{Z}_K/\mathfrak{p})[X]$ *with the* T_i *monic. Set*

$$g(X) = \prod_{1 \le i \le k} T_i(X), \quad h(X) = \prod_{1 \le i \le k} T_i(X)^{e_i - 1} ,$$

so that $g(X)h(X) - T(X) \in \mathfrak{p}[X]$. *Set*

$$f(X) = \beta \cdot (g(X)h(X) - T(X)) \in \mathbb{Z}_K[X] ,$$

and let U *be a monic lift of* $\overline{T}/\left(\overline{f}, \overline{g}, \overline{h}\right)$ *to* $\mathbb{Z}_K[X]$. *The order given by Zassenhaus's theorem starting with* $\mathcal{O} = \mathbb{Z}_K[\theta]$ *is equal to*

$$\mathcal{O}' = \mathbb{Z}_K[\theta] + \mathfrak{p}^{-1}U(\theta)\mathbb{Z}_K[\theta] .$$

In particular, $\mathcal{O}$ *is* $\mathfrak{p}$*-maximal if and only if* $\left(\overline{f}, \overline{g}, \overline{h}\right) = 1$ *in* $(\mathbb{Z}_K/\mathfrak{p})[X]$.

Remarks

(1) The proof of this theorem is essentially identical to the one in the absolute case and is left to the reader (Exercise 33).
(2) The result does not depend on the uniformizer β that we choose.
(3) A more direct construction of $\mathcal{O}'$ can be obtained by generalizing [Coh0 (third printing), Exercise 3 of Chapter 6]; see Exercise 34.
(4) Because of the presence of the ideal $\mathfrak{p}^{-1}$, $\mathcal{O}'$ is not free in general. Using an HNF algorithm in Dedekind domains, we can obtain a pseudo-basis for $\mathcal{O}'$ if desired.

We can now give the complete relative round 2 algorithm, in a form slightly different from that of [Coh0]. We start with a "driver" algorithm.

Algorithm 2.4.9 (Relative Round 2). Let L/K be a relative extension, with $L = K(\theta)$ and θ an algebraic integer whose minimal monic polynomial in $K[X]$ is denoted $T(X)$. This algorithm computes a pseudo-basis $(\omega_i, \mathfrak{a}_i)$ for $\mathbb{Z}_L$ and the relative discriminant $\mathrm{disc}(L/K) = (\mathfrak{d}(L/K), \overline{d(L/K)})$.

1. [Factor discriminant of polynomial] Using [Coh0, Algorithm 3.3.7], compute $D \leftarrow \mathrm{disc}(T)$, and let $\overline{d(L/K)} \leftarrow \overline{D}$ in K^*/K^{*2}. Using Algorithm 2.3.22, factor $D\mathbb{Z}_K$ as $D\mathbb{Z}_K = \prod_{1 \le i \le k} \mathfrak{p}_i^{v_i}$.
2. [Initialize] Set $j \leftarrow 0$, $\mathcal{O} \leftarrow \mathbb{Z}_K[\theta]$, $\omega_i \leftarrow \theta^{i-1}$, and $\mathfrak{a}_i \leftarrow \mathbb{Z}_K$ for $1 \le i \le n$, and set $\mathfrak{d}(L/K) \leftarrow D\mathbb{Z}_K$.
3. [Finished?] If $j = k$, output $(\omega_i, \mathfrak{a}_i)$, $\mathrm{disc}(L/K) = (\mathfrak{d}(L/K), \overline{d(L/K)})$, and terminate the algorithm. Otherwise, let $j \leftarrow j + 1$, $\mathfrak{p} \leftarrow \mathfrak{p}_j$.
4. [Compute $\mathfrak{p}$-maximal order] If $v_j < 2$, go to step 3. Otherwise, using Algorithm 2.4.11 below, compute a pseudo-basis $(\omega_{i,\mathfrak{p}}, \mathfrak{a}_{i,\mathfrak{p}})$ of a $\mathfrak{p}$-maximal order $\mathcal{O}_\mathfrak{p}$ containing $\mathbb{Z}_K[\theta]$ as well as the integer $s_\mathfrak{p}$ such that the order-ideal of the torsion module $\mathcal{O}_\mathfrak{p}/\mathbb{Z}_K[\theta]$ (in other words, the index-ideal $[\mathcal{O}_\mathfrak{p} : \mathbb{Z}_K[\theta]]$) is equal to $\mathfrak{p}^{s_\mathfrak{p}}$.

5. [Join orders] If $s_{\mathfrak{p}} \neq 0$, use Algorithm 2.4.10 to set $\mathcal{O} \leftarrow \mathcal{O}\mathcal{O}_{\mathfrak{p}}$, let $(\omega_i, \mathfrak{a}_i)$ be the corresponding pseudo-basis, and set $\mathfrak{d}(L/K) \leftarrow \mathfrak{d}(L/K)\mathfrak{p}^{-2s_{\mathfrak{p}}}$. Go to step 3.

Given two orders $\mathcal{O}$ and $\mathcal{O}'$ in $\mathbb{Z}_L$, we define their product $\mathcal{O}\mathcal{O}'$ as the smallest order containing both $\mathcal{O}$ and $\mathcal{O}'$. It is clear that it is the set of linear combinations of products of elements of $\mathcal{O}$ by elements of $\mathcal{O}'$. The following trivial algorithm computes this product.

Algorithm 2.4.10 (Product of Orders). Let $\mathcal{O}$ and $\mathcal{O}'$ be orders of L given by pseudo-bases $(\omega_i, \mathfrak{a}_i)$ and $(\eta_j, \mathfrak{b}_j)$, respectively. This algorithm computes a pseudo-basis for $\mathcal{O}\mathcal{O}'$.

1. [Form products] Let E be the list of element products $\omega_i\eta_j$, and let L be the list of ideal products $\mathfrak{a}_i\mathfrak{b}_j$.
2. [Apply HNF] Using Algorithm 1.6.2, output the pseudo-basis for the module whose pseudo-generating set is (E, L), and terminate the algorithm.

Remark. Contrary to the case of ideal products where the use of a two-element representation considerably speeds up the algorithm, I do not see how to apply a similar method here.

We can now explain the construction of a $\mathfrak{p}$-maximal order, which is the essential part of the relative round 2 algorithm.

Algorithm 2.4.11 (Relative Round 2 at $\mathfrak{p}$). Let L/K be a relative extension of degree n, with $L = K(\theta)$ and θ an algebraic integer whose minimal monic polynomial in $K[X]$ is denoted by $T(X)$. Let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$, and let p be the prime number below $\mathfrak{p}$. This algorithm computes an integral pseudo-basis $(\omega_i, \mathfrak{a}_i)$ for a $\mathfrak{p}$-maximal order $\mathcal{O}_{\mathfrak{p}}$ containing $\mathbb{Z}_K[\theta]$ as well as the $\mathfrak{p}$-adic valuation $s_{\mathfrak{p}}$ of the order-ideal of $\mathcal{O}_{\mathfrak{p}}/\mathbb{Z}_K[\theta]$. We assume given $v_{\mathfrak{p}} = v_{\mathfrak{p}}(\mathrm{disc}(T))$ (otherwise, compute it using [Coh0, Algorithms 3.3.7 and 4.8.17]).

1. [Initialize] For $i = 1, \ldots, n$, set $\omega_i \leftarrow \theta^{i-1}$, $\mathfrak{a}_i \leftarrow \mathbb{Z}_K$, and $s_{\mathfrak{p}} \leftarrow 0$.
2. [Trivial case] If $v_{\mathfrak{p}} < 2$, $\mathbb{Z}_K[\theta]$ is $\mathfrak{p}$-maximal, so output $(\omega_i, \mathfrak{a}_i)$ and $s_{\mathfrak{p}}$ and terminate the algorithm.
3. [Find uniformizers of $\mathfrak{p}$ and $\mathfrak{p}^{-1}$] Using [Coh0, Algorithm 4.7.10] if necessary, compute π such that $\mathfrak{p} = p\mathbb{Z}_K + \pi\mathbb{Z}_K$. If $v_{\mathfrak{p}}(\pi) > 1$, set $\pi \leftarrow \pi + p$. Then, using steps 1 and 2 of [Coh0, Algorithm 4.8.17], find $\beta \in \mathbb{Z}_K$ such that $p\mathfrak{p}^{-1} = p\mathbb{Z}_K + \beta\mathbb{Z}_K$, and set $\beta \leftarrow \beta/p$ (β will be a uniformizer of $\mathfrak{p}^{-1}$ and π a uniformizer of $\mathfrak{p}$).
4. [Factor modulo $\mathfrak{p}$] Using a factorization algorithm in the finite field $\mathbb{Z}_K/\mathfrak{p}$, factor T modulo $\mathfrak{p}$ as $\overline{T} = \prod_i \overline{T_i}^{e_i}$, where the $\overline{T_i}$ are distinct, monic, irreducible polynomials in $(\mathbb{Z}_K/\mathfrak{p})[X]$ and $e_i > 0$ for all i. Set $\overline{g} \leftarrow \prod \overline{T_i}$, $\overline{h} \leftarrow \overline{T}/\overline{g}$, $f \leftarrow \beta \cdot (gh - T)$, $\overline{Z} \leftarrow (\overline{f}, \overline{g}, \overline{h})$, $\overline{U} \leftarrow \overline{T}/\overline{Z}$, and $z \leftarrow \deg(Z)$.

5. [Apply Dedekind] If $z = 0$, then $\mathbb{Z}_K[\theta]$ is $\mathfrak{p}$-maximal, so output $(\omega_i, \mathfrak{a}_i)$ and $s_\mathfrak{p}$ and terminate the algorithm. Otherwise, apply Algorithm 1.6.2 to the pseudo-generating set
$$((\omega_1, \ldots, \omega_n, \omega_1, \ldots, \omega_z), (\mathfrak{a}_1, \ldots, \mathfrak{a}_n, \mathfrak{p}^{-1}, \ldots, \mathfrak{p}^{-1}))$$
(of course at this stage we still have $\omega_i = \theta^{i-1}$ and $\mathfrak{a}_i = \mathbb{Z}_K$), replace $(\omega_i, \mathfrak{a}_i)$ by the new pseudo-basis obtained in this way, and set $s_\mathfrak{p} \leftarrow z$.
6. [Finished?] If $2s_\mathfrak{p} + 1 \geq v_\mathfrak{p}$, output $(\omega_i, \mathfrak{a}_i)$ and $s_\mathfrak{p}$ and terminate the algorithm.
7. [Compute $\mathfrak{p}$-radical] If $p \leq n$, proceed as follows. Set $q \leftarrow \mathcal{N}(\mathfrak{p})$, set $q_1 \leftarrow q$, and while $q_1 < n$, set $q_1 \leftarrow q_1 \cdot q$. Then compute the $n \times n$ matrix $A = (a_{i,j})$ over $\mathbb{Z}_K/\mathfrak{p}$ such that $\omega_j^{q_1} = \sum_{1 \leq i \leq n} a_{i,j}\omega_i$. On the other hand, if $p > n$, compute the $n \times n$ matrix $A = (a_{i,j})$ over $\mathbb{Z}_K/\mathfrak{p}$ such that $a_{i,j} = \mathrm{Tr}_{L/K}(\omega_i\omega_j)$. Finally, using [Coh0, Algorithm 2.3.1], compute a $\mathbb{Z}_K/\mathfrak{p}$-basis $\overline{\beta_1}, \ldots, \overline{\beta_l}$ of the kernel of A (this will be a basis of $I_\mathfrak{p}/\mathfrak{p}\mathcal{O}$).
8. [Compute basis of $\mathcal{O}/\mathfrak{p}\mathcal{O}$] Using Algorithm 1.5.2 on the $(\omega_i, \mathfrak{a}_i)$, compute a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathcal{O}/\mathfrak{p}\mathcal{O}$.
9. [Compute basis of $I_\mathfrak{p}/\mathfrak{p}I_\mathfrak{p}$] Using [Coh0, Algorithm 2.3.6], supplement the $\overline{\beta_i}$ found in step 7 with $\beta_{l+1}, \ldots, \beta_n$ so that the $(\overline{\beta_i})_{1 \leq i \leq n}$ form a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathcal{O}/\mathfrak{p}\mathcal{O}$. Set $\alpha_i = \beta_i$ for $1 \leq i \leq l$ and $\alpha_i = \pi\beta_i$ for $l+1 \leq i \leq n$ (where π was found in step 3), where the β_i are any lifts to $\mathcal{O}$ of $\overline{\beta_i} \in \mathcal{O}/\mathfrak{p}\mathcal{O}$.
10. [Compute big matrix] Compute coefficients $c_{i,j,k} \in \mathbb{Z}_K/\mathfrak{p}$ such that $\omega_k\alpha_j \equiv \sum_{1 \leq i \leq n} c_{i,j,k}\alpha_i \pmod{\mathfrak{p}I_\mathfrak{p}}$. Let C be the $n^2 \times n$ matrix over $\mathbb{Z}_K/\mathfrak{p}$ such that $C_{(i,j),k} = c_{i,j,k}$.
11. [Compute new order] Using [Coh0, Algorithm 2.3.1], compute a basis $\gamma_1, \ldots, \gamma_k$ of the kernel of C, where the γ_i are considered as elements of $(\mathbb{Z}_K/\mathfrak{p})^n$. For $1 \leq i \leq n$, set $v_i \leftarrow \omega_i$, $\mathfrak{b}_i \leftarrow \mathfrak{a}_i$; for $1 \leq i \leq k$, let v_{n+i} be a lift of γ_i to $\mathbb{Z}_K^n$, and set $\mathfrak{b}_{n+i} \leftarrow \mathfrak{p}^{-1}$. Apply Algorithm 1.6.2 to the pseudo-generating set $(v_i, \mathfrak{b}_i)_{1 \leq i \leq n+k}$, and let $(\omega_i', \mathfrak{a}_i')$ be the HNF-pseudo-basis thus obtained.
12. [Finished?] Let $t \leftarrow \sum_{1 \leq i \leq n}(v_\mathfrak{p}(\mathfrak{a}_i) - v_\mathfrak{p}(\mathfrak{a}_i'))$, and for all i set $(\omega_i, \mathfrak{a}_i) \leftarrow (\omega_i', \mathfrak{a}_i')$. If $t = 0$, output $(\omega_i, \mathfrak{a}_i)$ and $s_\mathfrak{p}$ and terminate the algorithm; otherwise set $s_\mathfrak{p} \leftarrow s_\mathfrak{p} + t$ and go to step 6.

Remarks

(1) Since most computations in this algorithm must be performed in the finite field $\mathbb{Z}_K/\mathfrak{p}$, it is important to note that we will represent elements of this field as explained after Corollary 2.4.7 and not in some more abstract manner.

(2) We need to factor polynomials in $(\mathbb{Z}_K/\mathfrak{p})[X]$. Although we have not given the algorithms explicitly, the algorithms given in [Coh0, Chapter 3] for factoring in $(\mathbb{Z}/p\mathbb{Z})[X]$ can easily be extended to the case of general finite fields. The details are left to the reader, who can also read general computer algebra books such as [GCL].

(3) To compute the $\mathfrak{p}$-radical, we have used Proposition 2.4.2 for $p \leq n$ and Proposition 2.4.3 for $p > n$; this seems to be the best choice.
(4) We can, of course, present the algorithm as in [Coh0], by keeping a single order that we enlarge for each $\mathfrak{p}$ until $\mathbb{Z}_L$ is obtained, instead of computing a $\mathfrak{p}$-maximal order for each $\mathfrak{p}$ and putting the orders together only at the end. The method given here is, however, usually faster.

2.4.2 Relative Polynomial Reduction

In most applications, it is essential to reduce polynomials defining a given number field. When the number field is given by an absolute defining polynomial, we use the Polred algorithm or one of its variants (for example, [Coh0, Algorithm 4.4.12]). When the number field is given by a relative defining polynomial, we have more choices. Let L/K be given by a defining polynomial $T_2(X) \in K[X]$. A first possibility is the use of a relative version of the Polred algorithm using a relative version of the LLL algorithm or of the Fincke–Pohst algorithm. This is probably the best approach, but the necessary relative lattice algorithms (see [Fie-Poh]) are for now not good enough to provide excellent reduction, although they do help somewhat.

A second possibility is the use of the absolute Polred algorithm, but using a *relative* integral pseudo-basis instead of an absolute one. This method works quite well and is far superior to the naive method consisting of applying Polred on some absolute defining polynomial. A possible algorithm is as follows.

Algorithm 2.4.12 (Simple Relative Polynomial Reduction). Let $K = \mathbb{Q}(\theta_1)$ be a number field defined by a root θ_1 of an irreducible polynomial $T_1(X) \in \mathbb{Q}[X]$, and let $L = K(\theta_2)$ be a relative extension defined by a root θ_2 of an irreducible polynomial $T_2(X) \in K[X]$. We let $m = [K : \mathbb{Q}] = \deg(T_1)$ and $n = [L : K] = \deg(T_2)$. Finally, if $T_2(X) = \sum_{1 \leq k \leq n} A_k(\theta_1) X^k$ for some polynomials A_k, we set $W(X, Y) \leftarrow \sum_{1 \leq k \leq n} A_k(Y) X^k$. This algorithm computes polynomials $P(X) \in \mathbb{Q}[X]$ defining subfields of L, those defining L usually being simpler than the absolute defining polynomial computed by Algorithm 2.1.11.

1. [Compute roots] Compute the complex roots $\theta_1^{(i)}$ of the polynomial $T_1(X)$ for $1 \leq i \leq m$. For $1 \leq i \leq m$, set $T_2^{(i)}(X) \leftarrow W(X, \theta_1^{(i)})$, and let $\theta_2^{(i,j)}$ be the complex roots of $T_2^{(i)}(X)$ for $1 \leq j \leq n$.
2. [Compute relative pseudo-basis] Using Algorithm 2.4.9, compute a relative integral pseudo-basis $(\omega_j, \mathfrak{a}_j)$. For $1 \leq j \leq n$, compute an LLL-reduced basis $(\alpha_{i,j})_{1 \leq i \leq m}$ of $\mathfrak{a}_j$. Write $\omega_j = W_j(\theta_1, \theta_2)$ and $\alpha_{i,j} = A_{i,j}(\theta_1)$ with $W_j(X, Y) \in \mathbb{Q}[X, Y]$ and $A_{i,j}(X) \in \mathbb{Q}[X]$; for $1 \leq i \leq m$ and $1 \leq j \leq n$, set $B_{i,j}(X, Y) \leftarrow W_j(X, Y) A_{i,j}(Y)$.
3. [Compute absolute T_2-matrix] For all i_1, i_2, j_1, j_2 such that $1 \leq i_1, i_2 \leq m$ and $1 \leq j_1, j_2 \leq n$, set

$$a_{(i_1,j_1),(i_2,j_2)} \leftarrow \sum_{i_3,j_3} B_{i_1,j_1}(\theta_2^{(i_3,j_3)},\theta_1^{(i_3)})\overline{B_{i_2,j_2}(\theta_2^{(i_3,j_3)},\theta_1^{(i_3)})} \ .$$

4. [Apply LLL] Let A be the $nm \times nm$ matrix whose entries are the $a_{(i_1,j_1),(i_2,j_2)}$ (this will be a real, positive-definite, symmetric matrix). Apply one of the LLL algorithms (for example, [Coh0, Algorithm 2.6.3]) to this matrix, thus finding an LLL-reduced basis $\mathbf{b}_k$ for $1 \le k \le nm$.

5. [Compute corresponding polynomial] For $1 \le k \le nm$, proceed as follows. The coordinates of $\mathbf{b}_k$ are indexed by pairs (i,j) with $1 \le i \le m$ and $1 \le j \le n$, so let $\mathbf{b}_k = (u_{i,j})$. Set

$$\gamma_k \leftarrow \sum_{i,j} u_{i,j}\omega_j\alpha_{i,j} \ .$$

   Using [Coh0, Section 4.3], compute the characteristic polynomial $C_k \in \mathbb{Q}[X]$ of γ_k.

6. [Terminate] For each $k \le nm$, compute $P_k(X) \leftarrow C_k(X)/(C_k(X),C_k'(X))$. Output the $P_k(X)$ and terminate the algorithm.

Remarks

(1) As usual in polynomial reduction algorithms, we may not be interested in all the polynomials P_k, but only in those that define the field L, in other words those whose degree is equal to mn, and among those, in the ones with the smallest "size", for example in the sense of [Coh0, Algorithm 4.4.12]. For this, we modify step 6 accordingly.

(2) In step 4, we only find small elements for the T_2 norm in the sense of the LLL algorithm. If desired, we can strengthen the search and look for elements having the smallest T_2 norm, using the Fincke–Pohst algorithm ([Coh0, Algorithm 2.7.7]).

2.4.3 Prime Ideal Decomposition

The Buchmann–Lenstra algorithm for prime decomposition ([Coh0, Algorithm 6.2.9) can also be extended very simply as follows.

Algorithm 2.4.13 (Relative Prime Ideal Decomposition). Let L/K be a relative extension of degree n, with $L = K(\theta)$ and θ an algebraic integer whose minimal monic polynomial in $K[X]$ is denoted $T(X)$. Let $\mathfrak{p} = p\mathbb{Z}_K + \pi\mathbb{Z}_K$ be a prime ideal of $\mathbb{Z}_K$, where $v_\mathfrak{p}(\pi) = 1$ (change π into $\pi + p$ if this is not the case). This algorithm computes the prime ideal factorization $\mathfrak{p}\mathbb{Z}_L = \prod_{1\le i\le g}\mathfrak{P}_i^{e_i}$ by giving for each i the values $e_i = e(\mathfrak{P}_i/\mathfrak{p})$, $f_i = f(\mathfrak{P}_i/\mathfrak{p})$, and a two-element representation $\mathfrak{P}_i = ((1,\mathfrak{p}),(\alpha,\mathfrak{a})) = \mathfrak{p}\mathbb{Z}_L + \alpha\mathfrak{a}\mathbb{Z}_L$. We assume that we have already computed a pseudo-basis $(\omega_i,\mathfrak{a}_i)$ of $\mathbb{Z}_L$ and the relative discriminant ideal $\mathfrak{d}(L/K)$. All the ideals I that we will use (except for the final $\mathfrak{P}_i$) will be represented by $\mathbb{Z}_K/\mathfrak{p}$-bases of $I/\mathfrak{p}\mathbb{Z}_L$.

1. [Check if easy] If $v_{\mathfrak{p}}(\mathrm{disc}(T)) = v_{\mathfrak{p}}(\mathfrak{d}(L/K))$, let $\overline{T(X)} = \prod_{1\le i\le g} \overline{T_i(X)}^{e_i}$ in $(\mathbb{Z}_K/\mathfrak{p})[X]$ be the factorization of $T(X)$ into distinct, monic, irreducible polynomials over the finite field $\mathbb{Z}_K/\mathfrak{p}$ (obtained by straightforward generalizations of the algorithms of [Coh0, Section 3.4]). For each i, let $f_i \leftarrow \deg(T_i)$, $\mathfrak{P}_i \leftarrow \mathfrak{p}\mathbb{Z}_L + T_i(\theta)\mathbb{Z}_L$, output the e_i, f_i, $\mathfrak{P}_i = ((1,\mathfrak{p}),(T_i(\theta),\mathbb{Z}_K))$, and terminate the algorithm.

2. [Compute $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$] Using Algorithm 1.5.2 on the pseudo-basis $(\omega_i,\mathfrak{a}_i)$, compute a $\mathbb{Z}_K/\mathfrak{p}$-basis $(\overline{\eta_i})$ of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$.

3. [Compute $I_{\mathfrak{p}}/\mathfrak{p}\mathbb{Z}_L$] If $p \le n$, proceed as follows. Set $q \leftarrow \mathcal{N}(\mathfrak{p})$ and $q_1 \leftarrow q$, and while $q_1 < n$, set $q_1 \leftarrow q_1 \cdot q$. Then compute the $n \times n$ matrix $A = (a_{i,j})$ over $\mathbb{Z}_K/\mathfrak{p}$ such that $\eta_j^{q_1} = \sum_{1\le i\le n} a_{i,j}\eta_i$. On the other hand, if $p > n$, compute the $n \times n$ matrix $A = (a_{i,j})$ over $\mathbb{Z}_K/\mathfrak{p}$ such that $a_{i,j} = \mathrm{Tr}_{L/K}(\eta_i\eta_j)$. Finally, using [Coh0, Algorithm 2.3.1], compute the kernel $\overline{I_{\mathfrak{p}}}$ of A as a $\mathbb{Z}_K/\mathfrak{p}$-vector space.

4. [Initialize list] Set $\mathcal{L} \leftarrow \{\overline{I_{\mathfrak{p}}}\}$ and $c \leftarrow 1$ ($\mathcal{L}$ will be a list of c ideals of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$).

5. [Finished?] If $c = 0$, terminate the algorithm.

6. [Compute separable algebra $\mathcal{A}$] Let $\overline{H}$ be an element of $\mathcal{L}$. Compute a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathcal{A} = \mathbb{Z}_L/H = (\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L)/(H/\mathfrak{p}\mathbb{Z}_L)$ in the following way. If $\overline{\beta_1},\dots,\overline{\beta_r}$ is the given $\mathbb{Z}_K/\mathfrak{p}$-basis of $\overline{H}$, set $\overline{\beta_{r+1}} \leftarrow \overline{1}$. Using [Coh0, Algorithm 2.3.6] and the given basis $\overline{\eta_i}$ of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$, supplement this family into a basis $(\overline{\beta_i})_{1\le i\le n}$ of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$. Then set $f \leftarrow n-r$, and for $1 \le i \le f$ set $\gamma_i \leftarrow \overline{\beta_{r+i}}$.

7. [Compute multiplication table] (Here the γ_i form a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathcal{A}$ whose first element is $\overline{1}$.) By using [Coh0, Algorithm 2.3.5], compute coefficients $a_{i,j,k}$ and $b_{i,j,k}$ in $\mathbb{Z}_K/\mathfrak{p}$ such that

$$\gamma_i\gamma_j = \sum_{1\le k\le f} a_{i,j,k}\gamma_k + \sum_{1\le k\le r} b_{i,j,k}\overline{\beta_k} \ .$$

   The multiplication table of the γ_i (which will be used implicitly from now on) is given by the $a_{i,j,k}$ (we can discard the $b_{i,j,k}$).

8. [Compute $V = \mathrm{Ker}(\phi)$] Let M be the matrix of the map $\alpha \mapsto \alpha^q - \alpha$ from $\mathcal{A}$ to $\mathcal{A}$ on the $\mathbb{Z}_K/\mathfrak{p}$-basis that we have found (where $q = |\mathbb{Z}_K/\mathfrak{p}| = \mathcal{N}(\mathfrak{p})$ as above). Using [Coh0, Algorithm 2.3.1], compute a basis M_1 of the kernel of M (whatever algorithm is used, ensure that the first column of M corresponds to $\alpha = \overline{1}$).

9. [Do we have a field?] If M_1 has at least two elements (that is, if the kernel of M is not one-dimensional), go to step 10. Otherwise, apply Subalgorithm 2.4.14 below, and output a two-element representation $((1,\mathfrak{p}),(\alpha,\mathfrak{a}))$, the ramification index e, and the residual degree f corresponding to the prime ideal H. Remove the ideal H from the list $\mathcal{L}$, set $c \leftarrow c-1$, and go to step 5.

10. [Find $m(X)$] Let $\alpha \in \mathcal{A}$ be an element of M_1 that is not proportional to $\overline{1}$. By computing the successive powers of α in $\mathcal{A}$, let $m(X) \in (\mathbb{Z}_K/\mathfrak{p})[X]$ be the minimal monic polynomial of α in $\mathcal{A}$.
11. [Factor $m(X)$] (We know that $m(X)$ is a squarefree product of linear polynomials.) By using [Coh0, Section 3.4] or simply by trial and error if q is small, factor $m(X)$ into linear factors as $m(X) = m_1(X) \cdots m_k(X)$ in $(\mathbb{Z}_K/\mathfrak{p})[X]$.
12. [Split H] As above, let $r = \dim_{\mathbb{Z}_K/\mathfrak{p}}(\overline{H})$. For $1 \leq i \leq r$, do as follows. Set $\alpha_i \leftarrow m_i(\alpha)$, let M_i be the $n \times (r+n)$ matrix over $\mathbb{Z}_K/\mathfrak{p}$ whose first r columns give the basis of $\overline{H}$ and the last express the $\alpha_i\eta_j$ on the η_k (recall that (η_j) is a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$ computed in step 2). Finally, let $\overline{H_i}$ be the image of M_i computed using [Coh0, Algorithm 2.3.2].
13. [Update list] Remove $\overline{H}$ and add $\overline{H_1}, \ldots, \overline{H_k}$ to the list $\mathcal{L}$, set $c \leftarrow c + k - 1$ and go to step 6.

The following straightforward algorithm explicitly computes the prime ideal $\mathfrak{P}$ from a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\mathfrak{P}/\mathfrak{p}\mathbb{Z}_L$.

Subalgorithm 2.4.14 (Compute $\mathfrak{P}$ from $\mathfrak{P}/\mathfrak{p}\mathbb{Z}_L$). Given an integral pseudo-basis $(\omega_i, \mathfrak{a}_i)$ of $\mathbb{Z}_L$, a prime ideal $\mathfrak{p}$ of $\mathbb{Z}_K$, and a prime ideal $\mathfrak{P}$ above $\mathfrak{p}$ given modulo $\mathfrak{p}\mathbb{Z}_L$ in the form $\overline{H} = H/\mathfrak{p}\mathbb{Z}_L$ as a $\mathbb{Z}_K/\mathfrak{p}$-vector space, this algorithm computes a pseudo–two-element representation of $\mathfrak{P}$, the ramification index $e(\mathfrak{P}/\mathfrak{p})$, and the residual degree $f(\mathfrak{P}/\mathfrak{p})$.

1. [Lift basis of $\mathfrak{P}$] Set $s \leftarrow \dim_{\mathbb{Z}_K/\mathfrak{p}}(\overline{H})$, let $\beta_1, \ldots, \beta_s$ be lifts to $\mathbb{Z}_L$ of a $\mathbb{Z}_K/\mathfrak{p}$-basis of $\overline{H}$, and set $f = f(\mathfrak{P}/\mathfrak{p}) \leftarrow n - s$.
2. [Compute pseudo-generating set] Set $\gamma_i \leftarrow \omega_i$ and $\mathfrak{c}_i \leftarrow \mathfrak{p}\mathfrak{a}_i$ for $1 \leq i \leq n$, $\gamma_{i+n} \leftarrow \beta_i$ and $\mathfrak{c}_{i+n} \leftarrow \mathbb{Z}_K$ for $1 \leq i \leq s$ ($(\gamma_i, \mathfrak{c}_i)$ is now a pseudo-generating set of $\mathfrak{P}$, with $\gamma_1 = 1$ and $\mathfrak{c}_1 = \mathfrak{p}$ if the pseudo-basis of $\mathbb{Z}_L$ is in HNF).
3. [Compute pseudo–two-element representation] Using Algorithm 2.3.11, compute a pseudo–two-element representation $\mathfrak{P} = ((1, \mathfrak{p}), (\alpha, \mathfrak{a}))$ of the prime ideal $\mathfrak{P}$.
4. [Compute $e(\mathfrak{P}/\mathfrak{p})$] Using Algorithms 2.3.14 and 2.3.13, compute the $\mathfrak{P}$-adic valuation $e = e(\mathfrak{P}/\mathfrak{p})$ of $\mathfrak{p}\mathbb{Z}_L$ (note that $(\omega_i, \mathfrak{p}\mathfrak{a}_i)$ is a pseudo-basis of $\mathfrak{p}\mathbb{Z}_L$). Output $\mathfrak{P}$, e, and f, and terminate the algorithm.

Remarks

(1) As already noted after [Coh0, Algorithm 6.2.9], the method given above is faster than the initial Buchmann–Lenstra method since it avoids costly ideal multiplications and divisions. Apart from that, the algorithm is essentially identical.

(2) Of course, in practice one also keeps the pseudo-element $(\beta, \mathfrak{b})$ computed by Algorithm 2.3.14 in step 4 so as to be able to compute $\mathfrak{P}$-adic valuations with Algorithm 2.3.13.

2.5 Relative and Absolute Representations

In this section, we consider the problem of going back and forth from relative to absolute representations of ideals and orders. Let the base field K be given as $K = \mathbb{Q}(\theta_1)$ (or $k(\theta_1)$ for some subfield k of K, but for simplicity of exposition we will restrict to $k = \mathbb{Q}$), and let L/K be a relative extension given as $L = K(\theta_2)$. In Section 2.1.5 we have seen how to compute an absolute defining polynomial for $L/\mathbb{Q}$, more precisely how to find a small integer k such that $\theta = \theta_2 + k\theta_1$ satisfies $L = \mathbb{Q}(\theta)$, how to find the minimal polynomial of θ over $\mathbb{Q}$, and how to express θ_1 and θ_2 in terms of θ (see Algorithm 2.1.11). Although it is preferable to work systematically with relative extensions, it is sometimes unavoidable to work also with absolute extensions, and in this case the above data are essential. In particular, when we perform operations between elements of K (represented as polynomials in θ_1 with coefficients in $\mathbb{Q}$) and elements of L (represented as polynomials in θ_2 with coefficients in $K = \mathbb{Q}(\theta_1)$), it is necessary to replace the expressions of θ_1 and θ_2 by the polynomials in θ as output by Algorithm 2.1.11.

To simplify the computations and to avoid many possible sources of errors, we suggest *changing* the relative defining polynomial so that $k = 0$ and $\theta = \theta_2$. Indeed, since $\theta = \theta_2 + k\theta_1$ and $\theta \in K$, we clearly have $K(\theta) = K(\theta_2) = L$, and if T_2 is the minimal polynomial of θ_2 in $K[X]$, it is clear that the minimal polynomial of θ over K is $T(X) = T_2(X - k\theta_1) \in K[X]$.

Hence, from now on we assume that $\theta_2 = \theta$.

2.5.1 Relative and Absolute Discriminants

Let $(\omega_i, \mathfrak{a}_i)$ be the HNF pseudo-basis of $\mathbb{Z}_L$ on the power basis θ^{i-1}, and let $T(X)$ be the minimal monic polynomial of θ. Since the matrix of the ω_j in terms of the θ^{i-1} is upper-triangular with 1 in the diagonal, the relative discriminant ideal $\mathfrak{d}(L/K)$ is given by the formula (see Section 2.2.3)

$$\mathfrak{d}(L/K) = d(\omega_1, \ldots, \omega_n) \prod_{1 \le i \le n} \mathfrak{a}_i^2 = \mathrm{disc}(T) \prod_{1 \le i \le n} \mathfrak{a}_i^2 \ .$$

From this, it is easy to compute the absolute discriminant using the following important theorem.

Theorem 2.5.1. *Let L/K be an extension, and as usual let $\mathfrak{d}(L/K)$ be the discriminant ideal of L/K. Denote by (r_1, r_2) (resp., (R_1, R_2)) the signature of K (resp., L). The absolute discriminant $d(L)$ of L is given by the following formula:*

$$d(L) = (-1)^{R_2 - [L:K]r_2} d(K)^{[L:K]} \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \ ,$$

where $R_2 - [L:K]r_2$ is given by Proposition 2.2.5.

Proof. This theorem immediately follows from the transitivity of the different. Indeed, by Proposition 2.3.17, we have $\mathfrak{D}(L/\mathbb{Q}) = \mathfrak{D}(L/K)\mathfrak{D}(K/\mathbb{Q})$. Taking norms, and using $\mathfrak{d}(L/K) = \mathcal{N}_{L/K}(\mathfrak{D}(L/K))$ for any extension L/K, we obtain

$$\begin{aligned} d(L)\mathbb{Z} = \mathfrak{d}(L/\mathbb{Q}) &= \mathcal{N}_{L/\mathbb{Q}}(\mathfrak{D}(L/\mathbb{Q})) = \mathcal{N}_{K/\mathbb{Q}}(\mathcal{N}_{L/K}(\mathfrak{D}(L/K)\mathfrak{D}(K/\mathbb{Q}))) \\ &= \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)\mathfrak{D}(K/\mathbb{Q})^{[L:K]}) = \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K))d(K)^{[L:K]} \ . \end{aligned}$$

It follows that $d(L) = \pm d(K)^{[L:K]}\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K))$. By [Coh0, Proposition 4.8.11], we know that the sign of $d(K)$ is $(-1)^{r_2}$ and that of $d(L)$ is $(-1)^{R_2}$, from which the theorem follows. □

It is possible to give an alternate proof of this theorem using only the expression of the discriminant as a determinant of traces (see [Bou2]), but the proof given here is much more natural and directly follows [Ser]. Of course, as in Proposition 2.3.17, we can replace the base field $\mathbb{Q}$ by an arbitrary number field k.

2.5.2 Relative and Absolute Bases

Consider now the problem of the absolute integral basis of $\mathbb{Z}_L$. As in the proof of Theorem 2.5.1, from an integral pseudo-basis $(\omega_j, \mathfrak{a}_j)$ and $\mathbb{Z}$-bases $(\alpha_{i,j})$ of $\mathfrak{a}_j$, we immediately obtain an integral basis $\alpha_{i,j}\omega_j$ for $\mathbb{Z}_L$. Although in general it is not in HNF, it is usually *not* a good idea to put it in HNF (although using an HNF reduction algorithm, this is trivially done if desired) since an HNF is usually very badly skewed. For example, in applications such as the polynomial reduction algorithm Polred, we want an LLL-reduced basis for the T_2-norm (see Algorithm 2.4.12).

Usually the ω_j will be given on the relative power basis $1, \theta_2, \ldots, \theta_2^{n-1}$, and the $\alpha_{i,j}$ will be given on an absolute power basis $1, \theta_1, \ldots, \theta_1^{m-1}$ of $K/\mathbb{Q}$ (or as an HNF on an absolute integral basis of $\mathbb{Z}_K$ whose columns are easy to transform into polynomials in θ_1). To perform operations such as the element product $\alpha_{i,j}\omega_j$, we must express θ_1 as a polynomial in θ as explained above (or express both θ_1 and θ_2 as polynomials in θ if we have not changed the relative defining polynomial so that $\theta = \theta_2$).

The same method applies to ideals, except that one must be careful that the pseudo-matrix representation of the pseudo-basis $(\beta_j, \mathfrak{b}_j)$ of an ideal is usually given on the K-basis ω_i and not on the power basis θ^{i-1}, so we must first convert the generating elements of the ideal into polynomials in θ before performing the conversion. Once the absolute $\mathbb{Z}$-basis of the ideal is obtained, expressed as a matrix on powers on θ, by using linear algebra we can transform this matrix so as to get the matrix on the absolute integral basis. At this point it is very important to know *which* absolute integral basis is chosen, either the absolute HNF basis or the basis we obtained above

from the relative HNF basis, which is less skewed. The choice is not very important, but evidently it must be consistent.

Conversely, if only an absolute integral basis $(\eta_j)_{1\le j\le mn}$ of $\mathbb{Z}_L$ is known, then I do not see any really better method to find a pseudo-basis of $\mathbb{Z}_L$ than to apply Algorithm 1.6.2 to the pseudo-generating set $(\eta_j, \mathbb{Z}_K)$. For an ideal I of $\mathbb{Z}_L$, we can do as for $\mathbb{Z}_L$, but assuming that a pseudo-basis $(\omega_j, \mathfrak{a}_j)$ of $\mathbb{Z}_L$ is known, we can do much better as follows. From some absolute $\mathbb{Z}$-basis of I, compute an absolute two-element representation (α, β) using Algorithm 1.3.15 (in the absolute case). If we set $\beta_j = \alpha\omega_j$ and $\mathfrak{b}_j = \mathfrak{a}_j$ for $1 \le j \le n$, $\beta_{j+n} = \beta\omega_j$ and $\mathfrak{b}_{j+n} = \mathfrak{a}_j$ for $1 \le j \le n$, we obtain a $2n$-element pseudo-generating set of I, to which we apply Algorithm 1.6.2. This is, of course, much less costly than applying it to an mn-element pseudo-generating set.

Considering its importance, we isolate from the above discussion an algorithm to compute the relative norm of an ideal when one knows only an absolute basis of the ideal and a relative pseudo-basis of $\mathbb{Z}_L$ (if one knows a relative pseudo-basis of the ideal, the answer is given by Proposition 2.3.1).

Algorithm 2.5.2 (Relative Norm of an Ideal). Given an absolute basis of an ideal I of L and a pseudo-basis $(\omega_j, \mathfrak{a}_j)$ of $\mathbb{Z}_L$, this algorithm computes the relative norm $\mathcal{N}_{L/K}(I)$.

1. [Compute two-element representation] Using Algorithm 1.3.15 for the number field L, compute an absolute two-element representation (α, β) of the ideal I.
2. [Compute pseudo-generating set] Set $\beta_j \leftarrow \alpha\omega_j$ and $\mathfrak{b}_j \leftarrow \mathfrak{a}_j$ for $1 \le j \le n$, $\beta_{j+n} \leftarrow \beta\omega_j$ and $\mathfrak{b}_{j+n} \leftarrow \mathfrak{a}_j$ for $1 \le j \le n$.
3. [Apply Hermite] By applying Algorithm 1.6.2 to the $2n$-element pseudo-generating set $(\beta_j, \mathfrak{b}_j)_{1\le j\le 2n}$, compute a pseudo-basis $(\gamma_j, \mathfrak{c}_j)_{1\le j\le n}$ of the ideal I (and output this pseudo-basis if desired).
4. [Terminate] Output $\mathcal{N}_{L/K}(I) \leftarrow \prod_{1\le j\le n} \mathfrak{c}_j\mathfrak{a}_j^{-1}$ and terminate the algorithm.

2.5.3 Ups and Downs for Ideals

An important special case of the above is when a prime ideal $\mathfrak{P}$ of $\mathbb{Z}_L$ is given by an *absolute* two-element representation $\mathfrak{P} = p\mathbb{Z}_L + \beta\mathbb{Z}_L$. The above method gives a pseudo-basis $(\beta_j, \mathfrak{b}_j)$ for $\mathfrak{P}$. We can then apply Algorithm 2.3.11 to find a two-element relative representation $((1, \mathfrak{p}), (\alpha, \mathfrak{a}))$. However, we can obtain this representation more directly as follows.

Algorithm 2.5.3 (Prime Ideal Down). Given an absolute two-element representation of a prime ideal $\mathfrak{P} = p\mathbb{Z}_L + \alpha\mathbb{Z}_L$, this algorithm computes the prime ideal $\mathfrak{p}$ of $\mathbb{Z}_K$ below $\mathfrak{P}$ (the relative two-element representation will then simply be $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \alpha\mathbb{Z}_L$).

1. [Compute $\mathcal{N}_{L/K}(\alpha)$] Using the subresultant algorithm and Section 2.2.2, compute $a \leftarrow \mathcal{N}_{L/K}(\alpha)$.

2. [Easy case] If it is known in advance that $\mathfrak{p}$ is unramified in $K/\mathbb{Q}$ (for example, if $\mathfrak{P}$ itself is unramified in $L/\mathbb{Q}$), output $\mathfrak{p} \leftarrow p\mathbb{Z}_K + a\mathbb{Z}_K$ and terminate the algorithm.
3. [Difficult case] Using [Coh0, Algorithm 6.2.9], compute the prime ideals $\mathfrak{p}_i$ of $\mathbb{Z}_K$ above p.
4. [Loop] Using [Coh0, Algorithm 4.8.17], for each i compute $v_i \leftarrow v_{\mathfrak{p}_i}(a)$ until $v_i > 0$ (if no such i exists, there is an error). For this index i, set $\mathfrak{p} \leftarrow \mathfrak{p}_i$, output $\mathfrak{p}$, and terminate the algorithm.

Proof. We note that for every prime ideal $\mathfrak{Q}$ of $\mathbb{Z}_L$ above p and different from $\mathfrak{P}$, we have $v_{\mathfrak{Q}}(\alpha) = 0$. It follows that the relative norm $\mathcal{N}_{L/K}(\alpha)$ is of the form $\mathcal{N}_{L/K}(\alpha) = \mathfrak{p}^u\mathfrak{a}$, where $\mathfrak{a}$ is an ideal coprime to $p\mathbb{Z}_K$ and $u = v_{\mathfrak{P}}(\alpha)f(\mathfrak{P}/\mathfrak{p}) \geq 1$. Thus, if we know that $\mathfrak{p}$ is unramified, then $\mathfrak{P} \cap \mathbb{Z}_K = \mathfrak{p} = p\mathbb{Z}_K + \mathcal{N}_{L/K}(\alpha)\mathbb{Z}_K$ as can be seen, for example, by computing valuations at all prime ideals, proving the validity of step 2. If $\mathfrak{p}$ is ramified, the loop in step 4 gives a unique index i such that $v_{\mathfrak{p}_i}(\mathcal{N}_{L/K}(\alpha)) > 0$, in other words the index i for which $\mathfrak{p} = \mathfrak{p}_i$, proving the validity of step 4. Finally, we note that, if $\mathfrak{p} = \mathfrak{P} \cap \mathbb{Z}_K$, we have

$$\mathfrak{P} = p\mathbb{Z}_L + \alpha\mathbb{Z}_L \subset \mathfrak{p}\mathbb{Z}_L + \alpha\mathbb{Z}_L \subset \mathfrak{P} + \alpha\mathbb{Z}_L = \mathfrak{P} \ ;$$

hence $\mathfrak{P} = ((1,\mathfrak{p}),(\alpha,\mathbb{Z}_K)) = \mathfrak{p}\mathbb{Z}_L + \alpha\mathbb{Z}_L$ is a relative two-element representation of $\mathfrak{P}$. □

Assume now that we have an integral pseudo-basis $(\omega_i,\mathfrak{a}_i)$ of $\mathbb{Z}_L$ and an ideal I of $\mathbb{Z}_L$ given by a pseudo-basis $(\beta_i,\mathfrak{b}_i)$. Computing the intersection $I \cap \mathbb{Z}_K$ is very easy. If not already in this form, use Algorithm 1.6.2 to compute the HNF of the given pseudo-basis, obtaining a new pseudo-basis $(\beta_i',\mathfrak{b}_i')$. By definition of the HNF we have $\beta_i' = 1$, and so $I \cap \mathbb{Z}_K = \mathfrak{b}_1'$. This of course assumes that the integral pseudo-basis for $\mathbb{Z}_L$ is always chosen such that $\omega_1 = 1$.

If $I = \mathfrak{P}$ is a prime ideal given by a two-element representation $\mathfrak{P} = ((1,\mathfrak{p}),(\alpha,\mathfrak{a}))$, we need not do this since $\mathfrak{P} \cap \mathbb{Z}_K = \mathfrak{p}$.

Finally, if $\mathfrak{P}$ is a prime ideal given by an *absolute* two-element representation $\mathfrak{P} = p\mathbb{Z}_L + \alpha\mathbb{Z}_L$, we apply Algorithm 2.5.3 to compute $\mathfrak{p} = \mathfrak{P} \cap \mathbb{Z}_K$.

Conversely, let $\mathfrak{c}$ be an ideal of $\mathbb{Z}_K$. To compute the ideal $\mathfrak{c}\mathbb{Z}_L$ as a relative HNF representation is trivial thanks to the chosen representation: if $(\omega_i,\mathfrak{a}_i)$ is a pseudo-basis of $\mathbb{Z}_L$, then $(\omega_i,\mathfrak{c}\mathfrak{a}_i)$ is a pseudo-basis of $\mathfrak{c}\mathbb{Z}_L$. In fact, since this ideal is now considered as an ideal of $\mathbb{Z}_L$, it will be represented as a pseudo-matrix on the ω_i (and not on the θ^{i-1}), and the matrix component will simply be the $n \times n$ identity matrix, and the ideals will be the $\mathfrak{c}\mathfrak{a}_i$. If we also want an absolute HNF representation, we use the method explained in Section 2.5.2. In other words, we use the following algorithm.

Algorithm 2.5.4 (Ideal Up in Absolute HNF). Let $(\omega_j,\mathfrak{a}_j)_{1\leq j\leq n}$ be a pseudo-basis of $\mathbb{Z}_L$ and let (η_k) be an absolute $\mathbb{Z}$-basis of $\mathbb{Z}_L$, not necessarily

coming from the pseudo-basis nor in HNF. Given an ideal $\mathfrak{c}$ of $\mathbb{Z}_K$, this algorithm computes the absolute HNF representation of the ideal $\mathfrak{c}\mathbb{Z}_L$ of $\mathbb{Z}_L$ with respect to the (η_k).

1. [Compute relative HNF] For all j set $\mathfrak{c}_j \leftarrow \mathfrak{a}_j\mathfrak{c}$ and let $(\gamma_{i,j})_i$ be a $\mathbb{Z}$-basis of $\mathfrak{c}_j$; then for all i and j set $\alpha_{i,j} \leftarrow \omega_j\gamma_{i,j}$.
2. [Compute absolute HNF] For all i and j, let $M_{(i,j)}$ be the column vector of the coordinates of $\alpha_{i,j}$ on the η_k, and let M be the matrix whose columns, indexed by pairs (i,j), are the $M_{(i,j)}$. Output the HNF of the matrix M and terminate the algorithm.

Finally, if $\mathfrak{c}$ is a prime ideal of $\mathbb{Z}_K$, and we want to know the prime ideals above $\mathfrak{c}$ in $\mathbb{Z}_L$ or the factorization of $\mathfrak{c}\mathbb{Z}_L$, we apply Algorithm 2.4.13.

2.6 Relative Quadratic Extensions and Quadratic Forms

As its title indicates, the aim of this section is to study the special case of relative *quadratic* extensions, and in particular to show that the usual theory of binary quadratic forms can naturally be extended to the relative case. This will give us a powerful computational tool that, as in the absolute case, we will use in Chapter 7 for computing relative class and unit groups.

2.6.1 Integral Pseudo-Basis, Discriminant

Let K be a base field, $D \in K^* \setminus K^{*2}$, and $L = K(\sqrt{D})$, so that L is the general quadratic extension of K. Since $K(\sqrt{D}) = K(\sqrt{Df^2})$ for all $f \in K^*$, we may assume that $D \in \mathbb{Z}_K$, and we make this assumption from now on.

In the absolute case, we can go a step further and assume that D is squarefree (in which case D is the so-called radicand), or that D is the discriminant of L (so D is squarefree congruent to 1 modulo 4 or equal to 4 times a squarefree integer congruent to 2 or 3 modulo 4). Thanks to these reductions, there is a bijection between quadratic extensions of $\mathbb{Q}$ and such integers D.

In the relative case the situation is not that easy when the base field K has a nontrivial class group: we may, of course, assume that D is "squarefree", meaning that it is not divisible by the square of a nonunit of $\mathbb{Z}_K$, but such a reduction is not sufficient since we can still have $D\mathbb{Z}_K = \mathfrak{f}^2\mathfrak{d}$ with a nontrivial ideal $\mathfrak{f}$. In addition, such a "squarefree reduction" would in general not be unique (see Exercise 36). We will see however, in Chapter 9 a way to do this properly (see Lemma 9.2.2 and Algorithm 9.2.3).

Thus, the only assumption we will make is that $D \in \mathbb{Z}_K$, and we write $D\mathbb{Z}_K = \mathfrak{f}^2\mathfrak{d}$ with $\mathfrak{f}$ an integral ideal and $\mathfrak{d}$ a squarefree ideal, which can be done uniquely.

We know that $\mathbb{Z}_L$ has a pseudo-basis over $\mathbb{Z}_K$ in HNF. Thus, in our case Corollary 2.2.9 tells us that the corresponding pseudo-matrix on the basis $(1, \sqrt{D})$ is equal to

$$\left(\begin{pmatrix} 1 & -\delta \\ 0 & 1 \end{pmatrix}, (\mathbb{Z}_K, \mathfrak{q}^{-1})\right) ,$$

where $\delta \in \mathbb{Z}_K$ and $\mathfrak{q}$ is an integral ideal of $\mathbb{Z}_K$.

In addition, the definition of the ideal-discriminant shows that the relative ideal-discriminant $\mathfrak{d}(L/K)$ is given by the formula $\mathfrak{d}(L/K) = 4D\mathfrak{q}^{-2} = (2\mathfrak{f}\mathfrak{q}^{-1})^2\mathfrak{d}$, and the index-ideal $[\mathbb{Z}_L : \mathbb{Z}_K[\sqrt{D}]]$ is equal to $\mathfrak{q}$.

Proposition 2.6.1. *Assume as above that $D \in \mathbb{Z}_K$, and set $D\mathbb{Z}_K = \mathfrak{f}^2\mathfrak{d}$ with $\mathfrak{f}$ integral and $\mathfrak{d}$ squarefree. Then*

(1)

$$2\mathfrak{f} \subset \mathfrak{q} \subset \mathfrak{f} \subset \mathbb{Z}_K$$

or, equivalently,

$$1 \in \mathbb{Z}_K \subset \mathfrak{f}^{-1} \subset \mathfrak{q}^{-1} \subset \frac{1}{2}\mathfrak{f}^{-1} ;$$

(2) $\delta \in \mathfrak{f} \subset \frac{1}{2}\mathfrak{q} \cap \mathbb{Z}_K$;
(3) $D - \delta^2 \in \mathfrak{q}^2$.

Conversely, if these conditions are satisfied, then $\mathcal{O} = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$ is an order of L containing $\mathbb{Z}_K[\sqrt{D}]$.

Proof. Let $\alpha = a + b\sqrt{D}$ with a and b in K, not necessarily integral. Then $\alpha \in \mathbb{Z}_L$ if and only if $2a \in \mathbb{Z}_K$ and $a^2 - b^2D \in \mathbb{Z}_K$. Indeed, if $\alpha \in \mathbb{Z}_L$, then $\sigma(\alpha) = a - b\sqrt{D} \in \mathbb{Z}_L$, where σ denotes the unique nontrivial K-automorphism of L, hence $\alpha + \sigma(\alpha) = 2a \in \mathbb{Z}_L \cap K = \mathbb{Z}_K$ and similarly $\alpha\sigma(\alpha) = a^2 - b^2D \in \mathbb{Z}_K$. Conversely, if these conditions are satisfied, then α is a root of the monic polynomial $X^2 - 2aX + (a^2 - b^2D)$ with coefficients in $\mathbb{Z}_K$, hence is an algebraic integer, proving our claim.

Thus, if $\alpha = a + b\sqrt{D} \in \mathbb{Z}_L$ we have $a \in \frac{1}{2}\mathbb{Z}_K$ and $a^2 - b^2D \in \mathbb{Z}_K$, and since $2a \in \mathbb{Z}_K$, we have $4b^2D \in \mathbb{Z}_K$. This means that for any prime ideal $\mathfrak{p}$,

$$2v_\mathfrak{p}(2b) + 2v_\mathfrak{p}(\mathfrak{f}) + v_\mathfrak{p}(\mathfrak{d}) \geq 0 .$$

Since $\mathfrak{d}$ is squarefree, $v_\mathfrak{p}(\mathfrak{d}) \leq 1$; thus for all prime ideals $\mathfrak{p}$ we have $v_\mathfrak{p}(2b) + v_\mathfrak{p}(\mathfrak{f}) \geq 0$, in other words $2b \in \mathfrak{f}^{-1}$. We have thus proved that $\mathbb{Z}_L \subset \frac{1}{2}(\mathbb{Z}_K + \mathfrak{f}^{-1}\sqrt{D})$.

In the other direction, let us show that $\mathbb{Z}_K + \mathfrak{f}^{-1}\sqrt{D} \subset \mathbb{Z}_L$. For $\mathbb{Z}_K$, this is trivial. Let $u \in \mathfrak{f}^{-1}$. Then

$$(u\sqrt{D})^2 = u^2D \in \mathfrak{f}^{-2}\mathfrak{f}^2\mathfrak{d} \subset \mathfrak{d} \subset \mathbb{Z}_K ,$$

so $u\sqrt{D}$ is an algebraic integer, hence $u\sqrt{D} \in \mathbb{Z}_L$, so $\mathfrak{f}^{-1}\sqrt{D} \subset \mathbb{Z}_L$, as claimed. To summarize, we have shown the double inclusion

$$\mathbb{Z}_K + \mathfrak{f}^{-1}\sqrt{D} \subset \mathbb{Z}_L \subset \frac{1}{2}(\mathbb{Z}_K + \mathfrak{f}^{-1}\sqrt{D}) \ .$$

The precise determination of $\mathbb{Z}_L$ (or, equivalently, of the behavior at the prime ideals dividing 2) can be achieved either by using the relative round 2 algorithm (Algorithm 2.4.9) or by using Hecke's Theorem 10.2.9, but we will not need it in the theoretical analysis, only in the algorithms.

Since we know that $\mathbb{Z}_L = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$, it follows from the above inclusions that $\mathfrak{f}^{-1} \subset \mathfrak{q}^{-1} \subset \frac{1}{2}\mathfrak{f}^{-1}$, or, equivalently, $2\mathfrak{f} \subset \mathfrak{q} \subset \mathfrak{f}$, proving (1).

From the first inclusion above, we also deduce that

$$\mathfrak{f}^{-1}\sqrt{D} \subset \mathbb{Z}_L = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta) \ .$$

Thus, if $u \in \mathfrak{f}^{-1}$, $u\sqrt{D} = u(\sqrt{D} - \delta) + u\delta$ and since

$$u(\sqrt{D} - \delta) \in \mathfrak{f}^{-1}(\sqrt{D} - \delta) \subset \mathfrak{q}^{-1}(\sqrt{D} - \delta) \subset \mathbb{Z}_L \ ,$$

we deduce that $u\delta \in \mathbb{Z}_L \cap K = \mathbb{Z}_K$ for all $u \in \mathfrak{f}^{-1}$; in other words, that $\delta \in \mathfrak{f}$, proving (2).

To prove (3), we have $\mathfrak{q}^{-1}(\sqrt{D} - \delta) \subset \mathbb{Z}_L$, so by applying the automorphism σ (which leaves K pointwise invariant) we also have $\mathfrak{q}^{-1}(-\sqrt{D} - \delta) \subset \mathbb{Z}_L$, and hence by multiplying we obtain $\mathfrak{q}^{-2}(\delta^2 - D) \subset \mathbb{Z}_L \cap K = \mathbb{Z}_K$, proving (3). Note that we could hope for a stronger result such as $D - \delta^2 \in 4\mathfrak{f}^2$ or even $D - \delta^2 \in 2\mathfrak{f}\mathfrak{q}$, but this is not true in general (see Exercise 37). Of course, a stronger result can be obtained by computing $\mathfrak{q}$ explicitly by the round 2 algorithm or by Hecke's theorem, but the statement obtained is not simple.

Conversely assume that these conditions are satisfied and let $\mathcal{O} = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$. Since $1 \in \mathfrak{q}^{-1}$, $\mathbb{Z}_K[\sqrt{D}] \subset \mathcal{O}$. Furthermore, let $q \in \mathfrak{q}^{-1}$, and let $\alpha = q(\sqrt{D} - \delta) = -q\delta + q\sqrt{D}$. By conditions (2) and (3) we have $-2q\delta \in 2\mathfrak{q}^{-1}\delta \in \mathbb{Z}_K$ and

$$(-q\delta)^2 - q^2 D = q^2(\delta^2 - D) \in \mathfrak{q}^{-2}(\delta^2 - D) \in \mathbb{Z}_K \ ;$$

hence by the necessary and sufficient condition proved above, we have $\alpha \in \mathbb{Z}_L$. This shows that $\mathbb{Z}_K[\sqrt{D}] \subset \mathcal{O} \subset \mathbb{Z}_L$. Hence, to show that $\mathcal{O}$ is an order, it is sufficient to show that $\mathcal{O}$ is a ring. Since it is trivially stable by addition and since $1 \in \mathcal{O}$, we must simply show that $\mathcal{O}$ is stable by multiplication.

This is of course equivalent to showing that $\mathfrak{q}^{-2}(\sqrt{D} - \delta)^2 \subset \mathcal{O}$. For this, note that

$$\begin{aligned}\mathfrak{q}^{-2}(\sqrt{D} - \delta)^2 &= \mathfrak{q}^{-2}(D - \delta^2 + 2\delta(\delta - \sqrt{D})) \\ &\subset \mathfrak{q}^{-2}(D - \delta^2) + \mathfrak{q}^{-1}(2\mathfrak{q}^{-1}\delta)(\delta - \sqrt{D}) \\ &\subset \mathbb{Z}_K + \mathfrak{q}^{-1}(\delta - \sqrt{D}) = \mathcal{O}\end{aligned}$$

by conditions (2) and (3), finishing the proof of the proposition. □

Remarks

(1) Corollary 2.2.9, which only uses the $\mathbb{Z}_K[\sqrt{D}]$-module structure, tells us only that $\mathfrak{q} \subset \mathbb{Z}_K$ and that $\delta \in \mathbb{Z}_K$.
(2) The general integral pseudo-basis of the form $\mathbb{Z}_L = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \varepsilon)$ is obtained with $\varepsilon = \delta + \eta$ for an arbitrary $\eta \in \mathfrak{q}$, see Exercise 38.

2.6.2 Representation of Ideals

We assume that $\mathfrak{q}$ and δ are known (obtained, for example, by the relative round 2 algorithm), and we now want to work with ideals in the relative extension L/K. The following proposition gives the result that we need.

Proposition 2.6.2. *Let I be a fractional ideal of L. There exist unique ideals $\mathfrak{n}$ and $\mathfrak{a}$ and an element $b \in \mathbb{Z}_K$ such that*

$$I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b)) \ .$$

In addition, we have the following:

(1) *the ideal $\mathfrak{a}$ is an integral ideal;*
(2) *we have $\delta - b \in \mathfrak{q}$, and in particular $b \in \mathfrak{f}$;*
(3) *the ideal $\mathfrak{c} = (b^2 - D)(\mathfrak{a}\mathfrak{q}^2)^{-1}$ is an integral ideal.*

Conversely, if these conditions are satisfied, then $I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b))$ is an ideal of L.

Proof. By Proposition 2.3.1, I has a pseudo-basis in HNF of the form $((1, \sqrt{D} - b), (\mathfrak{c}_1, \mathfrak{c}_2))$ with $b \in \mathbb{Z}_K$ and $\mathfrak{c}_2\mathfrak{q} \mid \mathfrak{c}_1$. Set $\mathfrak{a} = \mathfrak{c}_1(\mathfrak{c}_2\mathfrak{q})^{-1}$ and $\mathfrak{n} = \mathfrak{c}_2\mathfrak{q} = \mathfrak{c}_1\mathfrak{a}^{-1}$. Then $\mathfrak{a}$ is an integral ideal and

$$I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b)) \ ,$$

proving (1).

Set $J = I\mathfrak{n}^{-1} = \mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b)$. We must express the fact that I (or, equivalently, J) is an ideal of $\mathbb{Z}_L = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$. This is clearly equivalent to $\mathfrak{q}^{-1}(\sqrt{D} - \delta)J \subset J$, hence to $\mathfrak{q}^{-1}(\sqrt{D} - \delta)\mathfrak{a} \subset J$ and $\mathfrak{q}^{-1}(\sqrt{D} - \delta)\mathfrak{q}^{-1}(\sqrt{D} - b) \subset J$.

The first condition is equivalent to $aq(\sqrt{D} - \delta) \in J$ for all $a \in \mathfrak{a}$ and $q \in \mathfrak{q}^{-1}$. We have

$$aq(\sqrt{D} - \delta) = aq(\sqrt{D} - b) + aq(b - \delta) \ ,$$

and since $\mathfrak{a}$ is an integral ideal, $aq(\sqrt{D} - b) \in J$, so the first condition is equivalent to $aq(b - \delta) \in J$, hence to $aq(b - \delta) \in \mathfrak{a}$ for all $a \in \mathfrak{a}$ and $q \in \mathfrak{q}^{-1}$. This is in turn equivalent to $b - \delta \in \mathfrak{q}$, which implies that $b \in \delta + \mathfrak{q} \subset \mathfrak{f} + \mathfrak{f} = \mathfrak{f}$ by Proposition 2.6.1, proving (2). Conversely, it is clear that if (1) and (2)

are satisfied (more precisely, (1) and $b - \delta \in \mathfrak{q}$), then the first condition is satisfied.

Multiplying by $\mathfrak{q}$, the second condition on J means that for all $q \in \mathfrak{q}^{-1}$ we have

$$q(D + b\delta - (b+\delta)\sqrt{D}) \in \mathfrak{aq} \oplus \mathbb{Z}_K(\sqrt{D} - b) \ .$$

This implies that $\mathfrak{q}^{-1}(b+\delta) \subset \mathbb{Z}_K$, hence $b+\delta \in \mathfrak{q}$. However, $b+\delta = b-\delta+2\delta \in \mathfrak{q}$ by (2) and Proposition 2.6.1, so this condition is already satisfied. Thus, we write

$$q(D + b\delta - (b+\delta)\sqrt{D}) = -q(b+\delta)(\sqrt{D} - b) + q(D - b^2) \ .$$

We have $-q(b+\delta)(\sqrt{D} - b) \in J\mathfrak{q}$, so the second condition on J is equivalent to $q(D - b^2) \in \mathfrak{aq}$ for all $q \in \mathfrak{q}^{-1}$, hence to $D - b^2 \in \mathfrak{aq}^2$. This proves (3), and since we have considered only necessary and sufficient conditions, this finishes the proof of the proposition. □

Proposition 2.6.3. *Let* $I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b))$ *as in Proposition 2.6.2.*

(1) *The content of* I *in the sense of Definition 2.3.4 is the ideal* $\mathfrak{n}$.
(2) *The ideal* I *is an integral ideal of* $\mathbb{Z}_L$ *if and only if* $\mathfrak{n}$ *is an integral ideal of* $\mathbb{Z}_K$.
(3) *The ideal* I *is primitive in* L/K *if and only if* $\mathfrak{n} = \mathbb{Z}_K$.
(4) $\mathcal{N}_{L/K}(I) = \mathfrak{an}^2$.

Proof. Note that with the notation of Proposition 2.3.5 we have $h_{1,2} = \delta - b$, hence by that proposition $c(I) = (\mathfrak{n}, \mathfrak{na}, (\delta - b)\mathfrak{nq}^{-1}) = \mathfrak{n}$ since by Proposition 2.6.2, $\mathfrak{a}$ is an integral ideal and $\delta - b \in \mathfrak{q}$, proving (1), and (2) and (3) are immediate consequences. Statement (4) is a restatement of Proposition 2.3.1 (4). □

Definition 2.6.4. *A* pseudo-quadratic form *associated to the ideal* I *is the quadruple* $(\mathfrak{a}, b, \mathfrak{c}; \mathfrak{n})$ *of ideals and element satisfying the conditions of Proposition 2.6.2. When possible, we will often call* $(\mathfrak{a}, b, \mathfrak{c})$ *itself the pseudo-quadratic form associated to* I.

Remarks

(1) The element b is not unique but is clearly defined modulo addition of an arbitrary element of $\mathfrak{aq}$. Thus there is not a single pseudo-quadratic form associated to I, but an equivalence class under the action of a group generalizing the group of integer translations Γ_∞ used in the absolute case. The equivalence relation says that for all $u \in \mathfrak{aq}$ we have

$$(\mathfrak{a}, b, (b^2 - D)(\mathfrak{q}^2\mathfrak{a})^{-1}; \mathfrak{n}) \sim (\mathfrak{a}, b+u, ((b+u)^2 - D)(\mathfrak{q}^2\mathfrak{a})^{-1}; \mathfrak{n}) \ .$$

(2) In the classical case $K = \mathbb{Q}$ and $L = \mathbb{Q}(\sqrt{D})$, we may choose D to be the discriminant of the quadratic field (a choice we almost always make in [Coh0]). When $D \equiv 1 \pmod 4$ we have $\mathfrak{f} = \mathbb{Z}$, $\mathfrak{q} = 2\mathbb{Z}$, and $\delta = 1$, while when $D \equiv 0 \pmod 4$ we have $\mathfrak{f} = 2\mathbb{Z}$, $\mathfrak{q} = 2\mathbb{Z}$, and $\delta = 0$. Thus, an ideal I can be written $I = n(a\mathbb{Z} \oplus ((-b+\sqrt{D})/2)\mathbb{Z})$, and the conditions of the proposition say that I is an ideal if and only if $a \in \mathbb{Z}$, $b \equiv \delta \pmod 2$, and $4a \mid (b^2 - D)$, which are well-known results in the absolute case. The quadratic form associated to the ideal I is the form (a, b, c) with $c = (b^2 - D)/(4a)$, so the above definition is a perfect generalization to the relative case (the number n is not preserved in the absolute case but could be if desired).

(3) There is, however, one important difference between the relative and absolute cases. Since we could work with *primitive* ideals, we could discard the number n. In the relative case, the fact that $\mathfrak{n}$ may not be a principal ideal of $\mathbb{Z}_K$ forbids us to discard it when doing class group computations. This is the main reason for which, instead of simply writing a form as $(\mathfrak{a}, b, \mathfrak{c})$, we have written it as $(\mathfrak{a}, b, \mathfrak{c}; \mathfrak{n})$. We will however, omit the $\mathfrak{n}$ when it is not necessary or when it is equal to $\mathbb{Z}_K$. We will come back to the use of $\mathfrak{n}$ later.

(4) The ideal $\mathbb{Z}_L = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$ is represented by the *unit form* $(\mathbb{Z}_K, \delta, (\delta^2 - D)\mathfrak{q}^{-2}; \mathbb{Z}_K)$.

(5) Although we use the word "pseudo"-quadratic form, there really *is* a quadratic form here, which is the form

$$x^2 + 2bxy + (b^2 - D)y^2 \quad \text{with} \quad (x, y) \in \mathfrak{n}(\mathfrak{a} \times \mathfrak{q}^{-1}) \ .$$

Since we will not need this explicitly, I leave to the reader the study of this form and the generalizations of the correspondence between classes of ideals and classes of forms seen in [Coh0, Chapter 5] (see Exercise 39).

2.6.3 Representation of Prime Ideals

Let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$, and assume first that $\mathfrak{p}$ remains inert in $\mathbb{Z}_L$. Then the quadratic form associated to the ideal $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L$ is

$$(\mathbb{Z}_K, \delta, (\delta^2 - D)\mathfrak{q}^{-2}; \mathfrak{p}) \ .$$

This is already one instance where it is important to keep the fourth component; otherwise all inert prime ideals in L/K would be represented by the same form.

Assume now that $\mathfrak{p}$ is not inert in L/K and let $\mathfrak{P}$ be one of the prime ideals above $\mathfrak{p}$. Thanks to the existence of a pseudo–two-element representation, we know that there exist $\beta \in L$ and an ideal $\mathfrak{b}$ of K such that $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \beta\mathfrak{b}\mathbb{Z}_L$. Since $\mathfrak{P}$ is not inert, $\beta \notin K$ (otherwise $\mathfrak{P} = (\mathfrak{p} + \beta\mathfrak{b})\mathbb{Z}_L = \mathfrak{m}\mathbb{Z}_L$, hence $\mathfrak{m} = \mathfrak{p}$ and so $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L$), and since $L = K(\sqrt{D})$, we can write $\beta\mathfrak{b} = (\sqrt{D} - c_{\mathfrak{p}})\mathfrak{a}$ for some $c_{\mathfrak{p}} \in L$ and some ideal $\mathfrak{a}$ of K, so that

$$\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \mathfrak{a}(\sqrt{D} - c_{\mathfrak{p}})\mathbb{Z}_L$$

(this is, of course, *not* a direct sum).

We want to represent this as a pseudo-quadratic form. This is done by the following proposition.

Proposition 2.6.5. *Let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$ that is not inert in L/K, and let $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \mathfrak{a}(\sqrt{D} - c_{\mathfrak{p}})\mathbb{Z}_L$ be a prime ideal above $\mathfrak{p}$. Then $\mathfrak{p} + \mathfrak{a}\mathfrak{q} + \mathfrak{a}(\delta + c_{\mathfrak{p}}) = \mathbb{Z}_K$. Let $u \in \mathfrak{p}$, $v \in \mathfrak{a}\mathfrak{q}$, and $w \in \mathfrak{a}$ such that $u + v + w(\delta + c_{\mathfrak{p}}) = 1$ (recall that such elements can be found algorithmically by Algorithm 1.3.2). Then*

$$\mathfrak{P} = \mathfrak{p} \oplus \mathfrak{q}^{-1}(\sqrt{D} - (u\delta + vc_{\mathfrak{p}} - w(D + c_{\mathfrak{p}}\delta)))\ ;$$

in other words, if we set $b_{\mathfrak{p}} = u\delta + vc_{\mathfrak{p}} - w(D + c_{\mathfrak{p}}\delta)$, a pseudo-quadratic form associated to $\mathfrak{P}$ is the form $(\mathfrak{p}, b_{\mathfrak{p}}, (b_{\mathfrak{p}}^2 - D)(\mathfrak{q}^2\mathfrak{p})^{-1}; \mathbb{Z}_K)$.

Proof. By Proposition 2.6.2, we can write $\mathfrak{P} = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b_{\mathfrak{p}}))$ for unique ideals $\mathfrak{a}$ and $\mathfrak{n}$. Since $\mathfrak{P}$ is an integral ideal, $\mathfrak{n}$ is integral, and since $\mathfrak{P}$ is not inert, $\mathfrak{n}$ must be equal to $\mathbb{Z}_K$. In addition, $\mathfrak{a} = \mathfrak{P} \cap K = \mathfrak{p}$. Thus, we know a priori that $\mathfrak{P} = \mathfrak{p} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b_{\mathfrak{p}})$ for some element $b_{\mathfrak{p}}$, and any element b such that $\sqrt{D} - b \in \mathfrak{P}\mathfrak{q}$ will be suitable.

Replacing $\mathbb{Z}_L$ by $\mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$, we have

$$\mathfrak{P}\mathfrak{q} = \mathfrak{p}\mathfrak{q} + \mathfrak{p}(\sqrt{D} - \delta) + \mathfrak{a}\mathfrak{q}(\sqrt{D} - c_{\mathfrak{p}}) + \mathfrak{a}(\sqrt{D} - c_{\mathfrak{p}})(\sqrt{D} - \delta)\ .$$

Since we know that $\mathfrak{P}\mathfrak{q} = \mathfrak{p}\mathfrak{q} \oplus \mathbb{Z}_K(\sqrt{D} - b_{\mathfrak{p}})$, we have $\mathfrak{p} + \mathfrak{a}\mathfrak{q} + \mathfrak{a}(\delta + c_{\mathfrak{p}}) = \mathbb{Z}_K$. Thus, let $u \in \mathfrak{p}$, $v \in \mathfrak{a}\mathfrak{q}$ and $w \in \mathfrak{a}$ be such that $u + v + w(\delta + c_{\mathfrak{p}}) = 1$. Then

$$u(\sqrt{D} - \delta) + v(\sqrt{D} - c_{\mathfrak{p}}) - w(D + c_{\mathfrak{p}}\delta - (c_{\mathfrak{p}} + \delta)\sqrt{D}) \in \mathfrak{P}\mathfrak{q}\ ,$$

hence $\sqrt{D} - (u\delta + vc_{\mathfrak{p}} - w(D + c_{\mathfrak{p}}\delta)) \in \mathfrak{P}\mathfrak{q}$, so we can take $b_{\mathfrak{p}} = u\delta + vc_{\mathfrak{p}} - w(D + c_{\mathfrak{p}}\delta)$, proving the proposition. □

Remark. In most cases, the ideals $\mathfrak{p}$ and $\mathfrak{a}\mathfrak{q}$ are already coprime. This is in particular the case when $\mathfrak{p}$ does not divide the index-ideal $[\mathbb{Z}_L : \mathbb{Z}_K[\sqrt{D}]] = \mathfrak{q}$, since in that case Proposition 2.3.9 tells us that we can take $\mathfrak{a} = \mathbb{Z}_K$. It can also occur even when $\mathfrak{p} \mid \mathfrak{q}$. If this happens, the formulas simplify since we can choose $w = 0$.

We now consider the problem of computing *valuations* with respect to a prime ideal $\mathfrak{P}$. In the quadratic case, this is much simpler than applying Algorithms 2.3.13 and 2.3.14. Let $I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b))$ be an ideal and $\mathfrak{P}$ a prime ideal of L. If $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L$ is inert, we clearly have $v_{\mathfrak{P}}(I) = v_{\mathfrak{p}}(\mathfrak{n})$. Otherwise, we still have $v_{\mathfrak{P}}(I) = v_{\mathfrak{P}}(\mathfrak{n}) + v_{\mathfrak{P}}(I\mathfrak{n}^{-1})$ and $v_{\mathfrak{P}}(\mathfrak{n}) = e(\mathfrak{P}/\mathfrak{p})v_{\mathfrak{p}}(\mathfrak{n})$. Thus, we may assume that $\mathfrak{n} = \mathbb{Z}_K$, so that I is a primitive ideal.

If $\mathfrak{p}$ is ramified, so that $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}^2$, then $v_{\mathfrak{P}}(I) = v_{\mathfrak{p}}(\mathcal{N}_{L/K}(I)) = v_{\mathfrak{p}}(\mathfrak{a})$.

Consider finally the case where $\mathfrak{p} = \mathfrak{P}\mathfrak{P}'$ is split. If $I = \mathfrak{P}^v \cdot \mathfrak{P}'^{v'} \cdot J$ with J coprime to $\mathfrak{P}$ and $\mathfrak{P}'$, we have $v_{\mathfrak{p}}(\mathfrak{a}) = v_{\mathfrak{p}}(\mathcal{N}_{L/K}(I)) = v + v'$. Since I is primitive, we cannot have $\mathfrak{p} \mid I$, so v and v' cannot be simultaneously nonzero. Hence we have $v = 0$ or $v = v_{\mathfrak{p}}(\mathfrak{a})$. Assume that $v_{\mathfrak{p}}(\mathfrak{a}) > 0$ (otherwise $v = v_{\mathfrak{P}}(I) = 0$). Then $v = v_{\mathfrak{p}}(\mathfrak{a})$ if and only if $I \subset \mathfrak{P}$. Since $\mathfrak{a} \subset \mathfrak{P}$, $v = v_{\mathfrak{p}}(\mathfrak{a})$ if and only if

$$\mathfrak{q}^{-1}(\sqrt{D} - b) \subset \mathfrak{P} = \mathfrak{p} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b_{\mathfrak{p}}) \ .$$

This condition says that for all $q \in \mathfrak{q}^{-1}$ we have $q(\sqrt{D} - b) = q(\sqrt{D} - b_{\mathfrak{p}}) + q(b_{\mathfrak{p}} - b) \in \mathfrak{P}$, hence that $q(b_{\mathfrak{p}} - b) \in \mathfrak{p}$, so that $b_{\mathfrak{p}} - b \in \mathfrak{p}\mathfrak{q}$. Thus we test if this condition is satisfied. If it is, $v_{\mathfrak{P}}(I) = v_{\mathfrak{p}}(\mathfrak{a})$; otherwise, $v_{\mathfrak{P}}(I) = 0$.

We can write this down as a formal algorithm.

Algorithm 2.6.6 (Valuation at $\mathfrak{P}$ for a Relative Quadratic Extension). Let $\mathfrak{P}$ be a prime ideal of L above $\mathfrak{p}$ and let $I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b))$ be an ideal of L given as explained in Proposition 2.6.2. This algorithm computes the $\mathfrak{P}$-adic valuation $v_{\mathfrak{P}}(I)$. All $\mathfrak{p}$-adic valuations in the base field K are computed using [Coh0, Algorithm 4.8.17].

1. [Inert case] If $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}$ is inert, output $v \leftarrow v_{\mathfrak{p}}(\mathfrak{n})$ and terminate the algorithm.
2. [Ramified case] If $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}^2$, output $v \leftarrow 2v_{\mathfrak{p}}(\mathfrak{n}) + v_{\mathfrak{p}}(\mathfrak{a})$ and terminate the algorithm.
3. [Split case] (Here $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}\mathfrak{P}'$.) If $v_{\mathfrak{p}}(\mathfrak{a}) = 0$ or if $v_{\mathfrak{p}}(b_{\mathfrak{p}} - b) \leq v_{\mathfrak{p}}(\mathfrak{q})$ set $v \leftarrow v_{\mathfrak{p}}(\mathfrak{n})$; otherwise, set $v \leftarrow v_{\mathfrak{p}}(\mathfrak{n}) + v_{\mathfrak{p}}(\mathfrak{a})$. Output v and terminate the algorithm.

2.6.4 Composition of Pseudo-Quadratic Forms

We now consider the problem of computing the compositum of two pseudo-quadratic forms, of course defined as the product of the corresponding ideals. The result is completely analogous to the absolute case, as follows.

Proposition 2.6.7. *For $i = 1$, 2, and 3, let $I_i = \mathfrak{n}_i(\mathfrak{a}_i \oplus \mathfrak{q}^{-1}(\sqrt{D} - b_i))$ be three ideals of L, and assume that $I_3 = I_1 I_2$. Then $\mathfrak{n}_3$, $\mathfrak{a}_3$, b_3 are given by the following formulas. Set $\mathfrak{d} = \mathfrak{a}_1 + \mathfrak{a}_2 + \mathfrak{q}^{-1}(b_1 + b_2)$ and let $a_1 \in \mathfrak{a}_1\mathfrak{d}^{-1}$, $a_2 \in \mathfrak{a}_2\mathfrak{d}^{-1}$, and $q \in \mathfrak{q}^{-1}\mathfrak{d}^{-1}$ such that $a_1 + a_2 + q(b_1 + b_2) = 1$. Then*

$$\mathfrak{n}_3 = \mathfrak{d}\mathfrak{n}_1\mathfrak{n}_2, \quad \mathfrak{a}_3 = \mathfrak{a}_1\mathfrak{a}_2\mathfrak{d}^{-2}, \quad b_3 = b_2 + a_2(b_1 - b_2) + q(D - b_2^2) \ .$$

(We may, of course, reverse the roles of 1 and 2 in the formula for b_3.)

Proof. The proof is identical to that of the absolute case. We have $I_1 I_2 = \mathfrak{n}_1\mathfrak{n}_2 J$ with

$$J = \mathfrak{a}_1\mathfrak{a}_2 + \mathfrak{a}_1\mathfrak{q}^{-1}(\sqrt{D} - b_2) + \mathfrak{a}_2\mathfrak{q}^{-1}(\sqrt{D} - b_1) + \mathfrak{q}^{-2}(D + b_1 b_2 - (b_1 + b_2)\sqrt{D}) \ .$$

The ideal of coefficients of $\sqrt{D}$ is

$$\mathfrak{q}^{-1}(\mathfrak{a}_1+\mathfrak{a}_2+\mathfrak{q}^{-1}(b_1+b_2))=\mathfrak{q}^{-1}\mathfrak{d}=(\mathfrak{n}_3/(\mathfrak{n}_1\mathfrak{n}_2))\mathfrak{q}^{-1}\ ,$$

hence $\mathfrak{n}_3=\mathfrak{n}_1\mathfrak{n}_2\mathfrak{d}$, and by multiplicativity of the norm we know that $\mathcal{N}_{L/K}(J)=\mathfrak{a}_1\mathfrak{a}_2=\mathfrak{d}^2\mathfrak{a}_3$, so $\mathfrak{a}_3=\mathfrak{a}_1\mathfrak{a}_2\mathfrak{d}^{-2}$, as claimed.

Finally, if $a_1\in\mathfrak{a}_1\mathfrak{d}^{-1}$, $a_2\in\mathfrak{a}_2\mathfrak{d}^{-1}$ and $q\in\mathfrak{q}^{-1}\mathfrak{d}^{-1}$ are such that $a_1+a_2+q(b_1+b_2)=1$, then

$$a_1(\sqrt{D}-b_2)+a_2(\sqrt{D}-b_1)-q(D+b_1b_2-(b_1+b_2)\sqrt{D})\in J\mathfrak{q}\ ,$$

which is equal to $\sqrt{D}-b_3$ with

$$\begin{aligned}b_3&=a_1b_2+a_2b_1+q(D+b_1b_2)\\&=b_2(1-a_2-q(b_1+b_2))+a_2b_1+q(D+b_1b_2)\\&=b_2+a_2(b_1-b_2)+q(D-b_2^2)\ ,\end{aligned}$$

proving the proposition. □

In view of this proposition, it is reasonable to set the following definition (keeping in mind that a pseudo-quadratic form is really defined only modulo the equivalence relation mentioned in Remark (1) above).

Definition 2.6.8. *We define the* compositum *of two forms* $(\mathfrak{a}_1,b_1,\mathfrak{c}_1;\mathfrak{n}_1)$ *and* $(\mathfrak{a}_2,b_2,\mathfrak{c}_2;\mathfrak{n}_2)$ *by the following formulas. Set* $\mathfrak{d}=\mathfrak{a}_1+\mathfrak{a}_2+\mathfrak{q}^{-1}(b_1+b_2)$, *let* $a_1\in\mathfrak{a}_1\mathfrak{d}^{-1}$, $a_2\in\mathfrak{a}_2\mathfrak{d}^{-1}$ *and* $q\in\mathfrak{q}^{-1}\mathfrak{d}^{-1}$ *be such that* $a_1+a_2+q(b_1+b_2)=1$, *and, finally,* $b_3=b_2+a_2(b_1-b_2)+q(D-b_2^2)$. *Then*

$$(\mathfrak{a}_1,b_1,\mathfrak{c}_1;\mathfrak{n}_1)\cdot(\mathfrak{a}_2,b_2,\mathfrak{c}_2;\mathfrak{n}_2)=(\mathfrak{a}_1\mathfrak{a}_2\mathfrak{d}^{-2},b_3,(b_3^2-D)\mathfrak{q}^{-2}\mathfrak{d}^2(\mathfrak{a}_1\mathfrak{a}_2)^{-1};\mathfrak{d}\mathfrak{n}_1\mathfrak{n}_2)\ .$$

Corollary 2.6.9. *If* $I=\mathfrak{n}(\mathfrak{a}\oplus\mathfrak{q}^{-1}(\sqrt{D}-b))$, *then* $I^{-1}=\mathfrak{n}^{-1}\mathfrak{a}^{-1}(\mathfrak{a}\oplus\mathfrak{q}^{-1}(\sqrt{D}+b))$. *In other words, in terms of pseudo-quadratic forms we have*

$$(\mathfrak{a},b,\mathfrak{c};\mathfrak{n})^{-1}=(\mathfrak{a},-b,\mathfrak{c};\mathfrak{n}^{-1}\mathfrak{a}^{-1})\ .$$

Proof. By the above proposition, we have

$$(\mathfrak{a}\oplus\mathfrak{q}^{-1}(\sqrt{D}-b))(\mathfrak{a}\oplus\mathfrak{q}^{-1}(\sqrt{D}+b))=\mathfrak{d}(\mathfrak{a}_3\oplus\mathfrak{q}^{-1}(\sqrt{D}-b_3))$$

with $\mathfrak{d}=\mathfrak{a}+\mathfrak{a}+(b-b)\mathfrak{q}^{-1}=\mathfrak{a}$, $\mathfrak{a}_3=\mathfrak{a}^2\mathfrak{d}^{-2}=\mathbb{Z}_K$. In addition, we may choose $a_1=1\in\mathfrak{a}\mathfrak{d}^{-1}$, $a_2=0$, $q=0$, so that $b_3=b_2\equiv\delta\pmod{\mathfrak{q}}$ by Proposition 2.6.2; hence

$$(\mathfrak{a}\oplus\mathfrak{q}^{-1}(\sqrt{D}-b))(\mathfrak{a}\oplus\mathfrak{q}^{-1}(\sqrt{D}+b))=\mathfrak{a}(\mathbb{Z}_K\oplus\mathfrak{q}^{-1}(\sqrt{D}-\delta))=\mathfrak{a}\mathbb{Z}_L\ ,$$

and the first formula of the corollary follows. The second follows from the trivial observation that $(-b)^2=b^2$. □

Note that it is essential to keep the additional factor $\mathfrak{a}^{-1}$ occurring in this corollary, which is discarded in the absolute case.

2.6.5 Reduction of Pseudo-Quadratic Forms

Up to now, the analogy with the absolute case has been perfect. The situation breaks down when we consider the problem of *reduction* of pseudo-quadratic forms.

In any reduction procedure for forms of two variables, there are two completely distinct kinds of reduction steps. First are the *translations*, corresponding essentially to changing x into $x + ky$ while leaving y unchanged. In the absolute case over the integers, this corresponds to using the group Γ_∞ already mentioned above and in [Coh0, Chapter 5] of integer translations.

The second type is the *inversions*, corresponding essentially to the exchange of the variables x and y, perhaps with sign or similar harmless changes.

A third kind can also occur, multiplication of x or y by elements of K.

To take a very typical example, the LLL algorithm, as described in [Coh0, Section 2.6], operates in its primitive form only on *pairs* of vectors and hence is a succession of translations (Subalgorithm RED) and exchanges (Subalgorithm SWAP).

We first consider translations. In the case of ordinary quadratic forms over $\mathbb{Z}$, they correspond to transformations not changing the corresponding ideal. In our case, if $I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b))$, a translation $b \mapsto b + k$ will not change the ideal if and only if $k \in \mathfrak{aq}$. We then need to define what we mean by *translation-reduced*; in other words, we need to choose a "small" value of b modulo $\mathfrak{aq}$ representing the ideal.

As we have seen in Section 1.4.3, there are two ways to do this. One gives a *canonical* representative, using the HNF of the ideal $\mathfrak{aq}$ (see Algorithm 1.4.12). Although not too large (the coefficients are at most equal to the absolute norm of $\mathfrak{aq}$), this can still be large enough to create problems later on. It does have an advantage, which must be used in any implementation, since it allows us to test ideals for *equality* since this representation is unique.

The second method, using an LLL basis of $\mathfrak{aq}$ (see Algorithm 1.4.13), gives much smaller coefficients and so should be used for practical reduction. It has two disadvantages. The first one is that we lose uniqueness. This does not matter much, since for the rare times that we want to test ideal equality, we can always perform an HNF-type reduction. The second, more subtle disadvantage is that it is *slow*, since LLL is quite a sophisticated algorithm. Of course, we are dealing with rather small matrices here (of the size the absolute degree of the base field), but we are talking about an absolutely basic operation that may be performed several thousand times or more. As already mentioned after Algorithm 1.4.13, a good compromise is to use the notion of *partial* reduction introduced by P. Montgomery.

The second type of operations, corresponding to inversions, is simply a swap. We want to transform the pseudo-quadratic form $(\mathfrak{a}, b, \mathfrak{c}; \mathfrak{n})$ into $(\mathfrak{c}, -b, \mathfrak{a}; \mathfrak{m})$ for some $\mathfrak{m}$. This is indeed possible, as the following proposition shows.

Proposition 2.6.10. *Let* $(\mathfrak{a}, b, \mathfrak{c}; \mathfrak{n})$ *be a pseudo-quadratic form and* $I = \mathfrak{n}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b))$ *the corresponding ideal. We have the equality*

$$\mathfrak{n}\mathfrak{a}\mathfrak{q}(\mathfrak{c} \oplus \mathfrak{q}^{-1}(\sqrt{D} + b)) = (\sqrt{D} + b)I \ .$$

In particular, the ideal class in $Cl(L)$ *of the ideal corresponding to* $(\mathfrak{a}, b, \mathfrak{c}; \mathfrak{n})$ *is equal to ideal class of the ideal corresponding to* $(\mathfrak{c}, -b, \mathfrak{a}; \mathfrak{n}\mathfrak{a}\mathfrak{q})$.

Proof. This follows trivially from the equality

$$\mathfrak{c} \oplus \mathfrak{q}^{-1}(\sqrt{D} + b) = (\sqrt{D} + b)(\mathfrak{a}\mathfrak{q})^{-1}(\mathfrak{a} \oplus \mathfrak{q}^{-1}(\sqrt{D} - b)) \ .$$

□

This is the exact analog of the classical swap operation on ordinary quadratic forms except that, as usual, the behavior of the fourth component is important.

The problem starts to become difficult when we want to define what we mean by a *swap-reduced* pseudo-quadratic form. The closest analog of the classical case is the CM-case where the base field K is totally real (for example, a real quadratic field) and D is totally negative. Ideally, as in the imaginary quadratic case, one would like a definition of reducedness that would ensure that each ideal class contains exactly one reduced form. It seems that such a definition is difficult to find, and it would be very nice if one could be given.

If we stupidly copy the definition of the imaginary quadratic case, we can set the following unpleasant definition (but this is all I can offer at present).

Definition 2.6.11. *We say that a form* $(\mathfrak{a}, b, \mathfrak{c}; \mathfrak{n})$ *is* pseudo-reduced *if* b *is LLL-reduced modulo* $\mathfrak{a}\mathfrak{q}$ *in the sense of Algorithm 1.4.13 as explained above (partially LLL-reduced suffices in practice) and if we have the inequality*

$$\mathcal{N}(\mathfrak{a}) \leq \mathcal{N}(\mathfrak{c}) = \left|\mathcal{N}(b^2 - D)\right| / \mathcal{N}(\mathfrak{a}\mathfrak{q}^2) \ .$$

Although it is mathematically unpleasant, experiment has shown that it is usually sufficient for practical applications.

The reduction algorithm of a pseudo-quadratic form is, of course, immediate and need not be written formally: we partially reduce (or LLL-reduce if we agree to spend more time, but this is not a good practical choice) the element b modulo $\mathfrak{a}\mathfrak{q}$. If $\mathcal{N}(\mathfrak{a}) > \mathcal{N}(\mathfrak{c}) = \left|\mathcal{N}(b^2 - D)\right| / \mathcal{N}(\mathfrak{a}\mathfrak{q}^2)$ — in other words, if $\left|\mathcal{N}(b^2 - D)\right| < \mathcal{N}(\mathfrak{a}\mathfrak{q})^2$ — we swap $\mathfrak{a}$ and $\mathfrak{c}$ and change b in $-b$ and modify the fourth component as explained in Proposition 2.6.10, and we iterate this process until the form is reduced.

In Section 7.3.2, we will see a relative ideal reduction method that generalizes the above to arbitrary relative extensions.

2.7 Exercises for Chapter 2

1. Let
$$T(X) = \prod_{1 \le i \le g} T_i(X)^{e_i}$$
be a decomposition of T in $K[X]$ into nonassociate irreducible factors. Using the Chinese remainder theorem, show that
$$K[X]/T(X)K[X] \simeq \prod_{1 \le i \le g} K[X]/T_i(X)^{e_i}K[X] \ .$$

2. Prove the following variant of the primitive element theorem. If $K_1 = \mathbb{Q}(\theta_1)$ and $K_2 = \mathbb{Q}(\theta_2)$ are two number fields, there exists a (small) integer k such that $K_1K_2 = K(\theta_1\theta_2 + k\theta_2)$.

3. Let A be a ring. Show that A has no nonzero nilpotent elements if and only if $x^2 = 0$ implies $x = 0$ in A.

4. Let A be an étale algebra over K. Show that A is an integral domain if and only if A is a field.

5. Compute the Galois group of the separable but reducible polynomials $P_1(X) = (X-2)(X^3-X-1)$, $P_2(X) = (X^2-2)(X^2-3)$, and $P_3(X) = (X^2-2)(X^2-8)$. Compute also the discriminants of the corresponding étale algebras.

6. Classify up to conjugacy (and not only up to abstract isomorphism) all non-transitive subgroups of S_n for $n = 2$, 3, 4, and 5. For each of these groups, give an example of a polynomial in $\mathbb{Z}[X]$ whose Galois group is isomorphic (as a subgroup of S_n) to the given group.

7. Write a complete algorithm for computing the Galois group of separable but not necessarily irreducible polynomials over $\mathbb{Q}[X]$ in degree up to 5, generalizing the algorithms of [Coh0, Section 6.3].

8. Let T_1 and T_2 be two monic irreducible polynomials in $\mathbb{Q}[X]$ of degree n_1 and n_2, respectively, and let θ_1 (resp., θ_2) be a root of T_1 (resp., θ_2). If $\theta = \theta_1\theta_2 + k_1\theta_1 + k_2\theta_2$, show that θ is a root of $R(X, k_1, k_2) = 0$, where $R(X, Z_1, Z_2) = \mathcal{R}_Y(T_1(Y - Z_2), Y^{n_2}T_2((X - Z_1Y + Z_1Z_2)/Y))$. Assuming that $R(X, k_1, k_2)$ is squarefree, express θ_1 and θ_2 in terms of θ using the partial derivatives of $R(X, Z_1, Z_2)$.

9. Continuing the previous exercise, show that the exact analog of Theorem 2.1.14 for $R(X) = R(X, 0, 0)$ (assuming it is squarefree) is true with $V_s(X) = \mathcal{R}_Y(T_2(Y), XY - A_s(Y))$ and
$$U(X) = \frac{\mathcal{R}_Y(T_1(Y), X^{n_1}T_1(Y/X))}{(X-1)^{n_1}T_1(0)} \ ,$$
except that the result must also be multiplied by $T_1(0)^{n_2(n_2-1)}T_2(0)^{n_1(n_1-1)}$. Give the corresponding formulas for V_s and U when $R(X) = R(X, k_1, k_2)$.

10. If T is an irreducible polynomial of degree n, show that $\mathcal{R}_Y(T(Y), Y^nT(X/Y))$ is never squarefree.

11. At the end of the example given after Algorithm 2.1.8, θ_1 and θ_2 are two cube roots of 2 in the number field defined by the polynomial $R_2(X) = X^6 + 108$. Compute the third cube root of 2 in this field. Do the same for the reduced polynomial $S_2(X) = X^6 - 3X^5 + 5X^3 - 3X + 1$.

12. Show that, as claimed in the text, when the resultant in Y of two squarefree polynomials $T_1(Y)$ and $T_2(X - kY)$ is squarefree, the next-to-last polynomial in the polynomial remainder sequence given by the subresultant algorithm is of degree equal to 1 and equal to $R'(X)Y + R'_Z(X,k)$ up to a multiplicative constant.
13. Implement the computation of the compositum of two number fields, including the computation of θ_1 and θ_2, both by using the subresultant algorithm, and by using the other methods mentioned after Algorithm 2.1.8, and compare their relative speed.
14. Let M be an $n \times n$ invertible square matrix, and for $1 \leq i \leq d$ let B_i be a column vector with n entries.
    a) Show how to modify [Coh0, Algorithm 2.2.2] so as to compute the d solutions to $MX_i = B_i$ simultaneously.
    b) Estimate the running time of this algorithm, and compare it with the running time of [Coh0, Algorithms 2.2.1 and 2.2.2].
    c) Modify Algorithm 2.1.9 so that it uses the above method to compute simultaneously $R(X)$, θ_1 and θ_2.
15. (D. Simon) Let $T_1(Y)$ be a polynomial of degree n_1, let $W(X,Y)$ be a polynomial of degree n_2 in X, and let $R(X) = \mathcal{R}_Y(T_1(Y), W(X,Y))$ (so that in the context of Theorem 2.1.10, $R(X)$ is the defining polynomial of the extension L_2/K if Z can be taken equal to 0).
    a) Show that $T_1^{n_2}$ divides $\mathcal{R}_X(R(X), W(X,Y))$.
    b) If, in addition, $W(X,Y)$ is of the special form $W(X,Y) = A(X)Y + B(X)$, show that
$$\mathcal{R}_X(R(X), W(X,Y)) = T_1(Y)^{n_2}\mathcal{R}(A(X), B(X))^{n_1} \ .$$
16. Show that, in Algorithm 2.1.8, the choice of $k = 0$ in step 2 always leads to a nontrivial GCD when $\deg(T_1) > 1$.
17. Write a common generalization to Algorithms 2.1.8 and 2.1.11 in the context of étale algebras.
18. Prove the validity of Algorithm 2.1.12.
19. Let L/K be an extension of number fields, and assume that, in addition to the data for the base field K, we know only an absolute defining polynomial $P(X)$ for $L/\mathbb{Q}$. Write an algorithm for computing a relative defining polynomial for L/K.
20. Generalize Theorem 2.1.14 to the case where T_1 and T_2 are not necessarily monic and have only rational (as opposed to integral) coefficients.
21. Show that the characteristic polynomial is transitive in the following sense. If L/K is a relative extension, if $\alpha \in L$, and if $C_{\alpha,K}(X)$ (resp., $C_{\alpha,\mathbb{Q}}(X)$) denotes the relative (resp., absolute) characteristic polynomial of α, then
$$C_{\alpha,\mathbb{Q}}(X) = \mathcal{N}_{K/\mathbb{Q}}(C_{\alpha,K})(X) \ ,$$
    where the norm of a polynomial is obtained by computing the product of all the polynomials obtained by applying the $[K : \mathbb{Q}]$ embeddings of K into $\mathbb{C}$.
22. Let $K = \mathbb{Q}(\sqrt{10})$ and $L = K(\sqrt{-1})$. Using the relative round 2 algorithm, show that $\mathbb{Z}_L$ is not a free $\mathbb{Z}_K$-module.
23. Let L/K be a relative extension of number fields of degree n, and let $\alpha_1, \ldots, \alpha_n$ be n elements of L. Show that, as claimed in the text, $d(\alpha_1, \ldots, \alpha_n) = 0$ if and only if the α_j are K-linearly dependent.

24. Prove Proposition 2.2.10.
25. Using the explicit parameterization of cyclic cubic fields given in [Coh0, Section 6.4.2], compute the elementary discriminantal divisors of cyclic cubic fields over $\mathbb{Q}$. Do the same for pure cubic fields $\mathbb{Q}(\sqrt[3]{m})$.
26. If I is an ideal of L, show that, as claimed in the text, $\mathcal{N}_{L/K}(I) = \left(\prod_{\sigma_i} \sigma_i(I)\right) \cap K$, where the σ_i are the n K-embeddings of L into $\mathbb{C}$ and the product is considered in the Galois closure of L/K in $\mathbb{C}$.
27. Let $\mathfrak{b}$ be an integral ideal and $\mathfrak{a}$ any ideal of $\mathbb{Z}_K$. Let $\alpha \in \mathfrak{a}$ be such that $v_\mathfrak{p}(\alpha) = v_\mathfrak{p}(\mathfrak{a})$ for all $\mathfrak{p} \mid \mathfrak{b}$. Show that the map $x \mapsto \alpha x$ induces a $\mathbb{Z}_K$-module isomorphism from $\mathbb{Z}_K/\mathfrak{b}$ to $\mathfrak{a}/\mathfrak{a}\mathfrak{b}$.
28. Prove Proposition 2.3.5.
29. Let $I = (\alpha_i, \mathfrak{a}_i)$ be an ideal of $\mathbb{Z}_L$ given by a pseudo-basis (not necessarily in HNF), and let $J = ((\alpha, \mathfrak{a}), (\beta, \mathfrak{b}))$ be an ideal of $\mathbb{Z}_L$ given by a pseudo–two-element representation. Show that, as claimed in the text, $((\alpha_i\alpha, \alpha_i\beta), (\mathfrak{a}_i\mathfrak{a}, \mathfrak{a}_i\mathfrak{b}))$ is a $2n$-element pseudo-generating set of the ideal product IJ.
30. Let $K = \mathbb{Q}(\sqrt{D})$ be a quadratic field of discriminant D, and let p be a prime number that splits in K as $K = \mathfrak{p}\bar{\mathfrak{p}}$. Compute explicitly the HNF of the ideal $\mathfrak{p}^k$ on the usual integral basis $(1, \omega)$ of K, where $\omega = (\delta + \sqrt{D})/2$ with $\delta = D \bmod 2$. In addition, express your result using the truncation of a p-adic number (note that the corresponding exercise for inert or ramified primes is trivial).
31. Generalize Algorithm 2.3.24 to compute the list of all nth power free ideals of norm less than or equal to B — in other words, ideals not divisible by any nth prime power.
32. Write and implement an algorithm for computing the $\mathfrak{p}$-radical based on Proposition 2.4.3, and compare its efficiency with the corresponding algorithm based on Proposition 2.4.2 when p is large.
33. Prove Theorem 2.4.8, following closely the proof given in the absolute case given in [Coh0, Section 6.1.2].
34. (F. Diaz y Diaz) With the notation of Theorem 2.4.8, show that the Dedekind criterion can be restated as follows. Let $r_i(X) \in \mathbb{Z}_K[X]$ be the remainder of the Euclidean division of $T(X)$ by $t_i(X)$. We evidently have $r_i \in \mathfrak{p}[X]$. Set $d_i = 1$ if $e_i \geq 2$ and $r_i \in \mathfrak{p}^2[X]$, $d_i = 0$ otherwise. Then we can take $U(X) = \prod_{1 \leq i \leq k} t_i^{e_i - d_i}$. In particular, $\mathbb{Z}_K[\theta]$ is $\mathfrak{p}$-maximal if and only if $r_i \notin \mathfrak{p}^2[X]$ for every i such that $e_i \geq 2$.
35. Let $L = K(\alpha)$ be a relative extension of number fields, and let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis of $\mathbb{Z}_L$. Let $\beta = B(\alpha)$ with $B \in K[X]$ be an element of L, and let $N = K(\beta)$ be the subfield of L generated by β. Write an algorithm that directly computes an integral pseudo-basis of $\mathbb{Z}_N$ using the polynomial $B(X)$ and the pseudo-basis $(\omega_i, \mathfrak{a}_i)$.
36. Consider $K = \mathbb{Q}(\sqrt{-23})$, $\omega = (-1 + \sqrt{-23})/2$, and $D = 8\omega + 12$. Show that $D = 2^2(2\omega+3) = \omega^2(\omega+3)$ and that this gives two essentially distinct squarefree decompositions of D (thus showing that when the class number is larger than 1, this notion does not make sense for *elements*).
37. By simply considering the case $K = \mathbb{Q}$, show that, as claimed in Section 2.6, we do not have $D - \delta^2 \in 2\mathfrak{f}\mathfrak{q}$ in general.
38. Let $L = K(\sqrt{D})$ be a quadratic extension with $D \in \mathbb{Z}_K$.
    a) Show directly that there exists an integral ideal $\mathfrak{q}$ such that $\mathfrak{d}(L/K) = 4D\mathfrak{q}^{-2}$.

b) Show that $\mathbb{Z}_L = \mathbb{Z}_K \oplus \mathfrak{q}^{-1}(\sqrt{D} - \delta)$ if and only if $\delta \in \frac{1}{2}\mathfrak{q} \cap \mathbb{Z}_K$ and $D - \delta^2 \in \mathfrak{q}^2$, and that two such elements δ can only differ by an arbitrary element of $\mathfrak{q}$.

c) Can the condition $\delta \in \frac{1}{2}\mathfrak{q} \cap \mathbb{Z}_K$ be replaced by the condition $\delta \in \frac{1}{2}\mathfrak{q}$?

39. Generalize the correspondence between classes of ideals and classes of forms seen in [Coh0, Chapter 5] to the relative case, as suggested in Remark (5) after Definition 2.6.4.

3. The Fundamental Theorems of Global Class Field Theory

In this chapter, we give the main results of global class field theory for the case of number fields. We refer the reader to [Art-Tat], [Gras], [Has1], [Jan], or [Mart4] for more detailed statements and proofs. We present the results "à la Hasse", without using ideles. This is more suitable for algorithmic treatment. For an idelic treatment, we refer to [Neu]. I have largely benefited from the notes of J. Martinet [Mart4] in writing this chapter.

This chapter is entirely theoretical, and we defer all algorithms until Chapters 4, 5, and 6. However, as the reader will see, the presentation of the material is very concrete.

Class field theory is one of the most remarkable and important theories in number theory. In fact, a large part of the current trends in number theory (for example, the Langlands program) can be thought of as an attempt to generalize class field theory.

One of its remarkable aspects is that it gives a canonical bijection between rather different objects: on the one hand, *classes of congruence subgroups* (see definitions below), which are nothing more than certain groups of ideals in a base field K; on the other hand, *K-isomorphism classes of finite Abelian extensions of K*. There are two parts to this theorem (in fact, three, as we shall see), the injectivity and the surjectivity, but what is truly spectacular is certainly the surjectivity since it predicts a priori the existence of certain number fields, and it gives their discriminant and signature. Finding these number fields *in practice* is another matter (although in some sense we will simply follow the proof of the theorem), and we will explain in Chapters 5 and 6 how this is done.

3.1 Prologue: Hilbert Class Fields

Before explaining the general theory, we start with a special case that already embodies a large part of the theory. It will be generalized in the subsequent sections.

We will say that an extension L/K of number fields is *unramified* if there are no places of K that ramify in L. This means the following: for every prime ideal $\mathfrak{p}$ of K, we have a decomposition

$$\mathfrak{p}\mathbb{Z}_L = \prod_{\mathfrak{P}_i|\mathfrak{p}} \mathfrak{P}_i^{e_i} ,$$

and we want $\mathfrak{p}$ to be unramified, in other words, we want all e_i to be equal to 1. This must be true for every prime ideal of K. Since the ramified prime ideals are exactly those that divide the relative discriminant, this is equivalent to asking that $\mathfrak{d}(L/K) = \mathbb{Z}_K$. In addition, we also require the embeddings σ_i (or, equivalently, the places at infinity) to be unramified (see Definition 2.2.4).

We are concerned with finite *Abelian* unramified extensions L of a fixed base field K. There are several reasons for the restriction to the case of Abelian extensions, but perhaps the most important one is that very little is known in the non-Abelian case (see [Yam] and Exercise 1).

Hilbert and Furtwängler showed that there exists a *maximal* unramified Abelian extension of K (denoted by $K(1)$) in a strong sense: every Abelian unramified extension of K is isomorphic to a subextension of $K(1)$. This field $K(1)$ is called the *Hilbert class field* of K and has remarkable properties. First and foremost, the Galois group of $K(1)/K$ is isomorphic to the class group $Cl(K)$, hence in particular $[K(1) : K] = h(K)$, the class number of K. This isomorphism is explicitly given (all of this will be described in a more general setting below).

Second, the decomposition in $K(1)$ of a prime ideal of K can easily be described: if $\mathfrak{p}$ is a prime ideal of K, and if f is the least power of $\mathfrak{p}$ such that $\mathfrak{p}^f$ is a principal ideal of K, then $\mathfrak{p}$ splits into $h(K)/f$ distinct prime ideals of $K(1)$ of degree f.

By Galois theory, the subextensions of $K(1)/K$ (and thus all unramified Abelian extensions of K up to isomorphism) correspond in a one-to-one way to subgroups of the Galois group, hence to subgroups of the class group $Cl(K)$ or, equivalently, to subgroups C of the group $I(K)$ of fractional ideals of K containing the group $P(K)$ of principal ideals of K.

Note once again that we are talking about Abelian extensions. The study of general unramified extensions is much more difficult (see, for example, [Yam]).

When there is ramification, the situation is *completely* analogous, except that we must replace the ordinary class group by a more general class group called the *ray class group*, which we study in the next section.

The Hilbert class field extension $K(1)/K$ also possesses the *capitulation* property: every ideal of K becomes principal in $K(1)$; in other words, if $\mathfrak{a}$ is an ideal of K, then $\mathfrak{a}\mathbb{Z}_{K(1)}$ is a principal ideal of $\mathbb{Z}_{K(1)}$ (this is a theorem due to Furtwängler). This does *not* mean that $K(1)$ is itself principal (see Exercise 1). In fact, in the 1960s Golod and Shafarevitch proved that there exist *infinite class field towers*, meaning that there exist number fields H_0 such that if we define H_n to be the Hilbert class field of H_{n-1} for $n \geq 1$, then H_n is never equal to H_{n-1} (or, equivalently, H_{n-1} has a nontrivial class

group). This implies that there exist number fields that are not subfields of a number field whose ring of integers is a principal ideal domain.

The capitulation property was initially one of the main motivations for the study of class fields, but the further development of class field theory by Artin and Takagi has shown that this is more of an additional property than a basic one. We will come back to the study of capitulation in Chapter 7.

3.2 Ray Class Groups

3.2.1 Basic Definitions and Notation

The following definitions summarize most of the notions we will need to study ray class groups.

Definition 3.2.1. (1) *A modulus* $\mathfrak{m}$ *in* K *is a pair* $(\mathfrak{m}_0, \mathfrak{m}_\infty)$, *where* $\mathfrak{m}_0$ *is an integral ideal and* $\mathfrak{m}_\infty$ *is a set of real embeddings of* K *into* $\mathbb{C}$. *We will write this formally as* $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$.

(2) *If* $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$ *and* $\mathfrak{n} = \mathfrak{n}_0\mathfrak{n}_\infty$ *are two moduli, we say that* $\mathfrak{n}$ divides $\mathfrak{m}$ (*and write* $\mathfrak{n} \mid \mathfrak{m}$) *if* $\mathfrak{n}_0 \mid \mathfrak{m}_0$ (*or, equivalently,* $\mathfrak{n}_0 \supset \mathfrak{m}_0$) *and* $\mathfrak{n}_\infty \subset \mathfrak{m}_\infty$.

(3) *We define*

$$(\mathbb{Z}_K/\mathfrak{m})^* = (\mathbb{Z}_K/\mathfrak{m}_0)^* \times \mathbb{F}_2^{\mathfrak{m}_\infty} \ .$$

(4) *If* $\mathfrak{a}$ *is a nonzero fractional ideal of* K, *we say that* $\mathfrak{a}$ *is coprime to* $\mathfrak{m}$ *if* $v_\mathfrak{p}(\mathfrak{a}) = 0$ *for all* $\mathfrak{p} \mid \mathfrak{m}_0$ *or, equivalently, if we can write* $\mathfrak{a} = \mathfrak{b}/\mathfrak{c}$ *with* $\mathfrak{b}$ *and* $\mathfrak{c}$ *integral ideals coprime to* $\mathfrak{m}_0$ *in the usual sense* $(\mathfrak{b}+\mathfrak{m}_0 = \mathfrak{c}+\mathfrak{m}_0 = \mathbb{Z}_K)$. *The set of ideals coprime to* $\mathfrak{m}$ *is a group and is denoted by* $I_\mathfrak{m}(K)$ (*or* $I_\mathfrak{m}$ *if the field* K *is understood*). *If* $\alpha \in K^*$, *we say that* α *is coprime to* $\mathfrak{m}$ *if the principal ideal* $\alpha\mathbb{Z}_K$ *is coprime to* $\mathfrak{m}$.

Remark. When $\mathfrak{a}$ is not an integral ideal of K, the condition that $\mathfrak{a}$ is coprime to $\mathfrak{m}$ is of course *not* equivalent to $\mathfrak{a} + \mathfrak{m}_0 = \mathbb{Z}_K$, since this equality implies that $\mathfrak{a}$ is integral.

We have a natural group homomorphism ρ from the elements of K^* coprime to $\mathfrak{m}$ into $(\mathbb{Z}_K/\mathfrak{m})^*$ defined as follows. Any α coprime to $\mathfrak{m}$ in the above sense can be written as $\alpha = \beta/\gamma$ for β and γ in $\mathbb{Z}_K$ and coprime to $\mathfrak{m}$ (see Algorithm 4.2.22 for an algorithmic way of finding β and γ). Thus, we can define the class $\overline{\alpha} \in (\mathbb{Z}_K/\mathfrak{m}_0)^*$ by setting $\overline{\alpha} = \overline{\beta}/\overline{\gamma}$, and it is clear that this does not depend on the choice of β and γ. We then define ρ by setting

$$\rho(\alpha) = \left(\overline{\alpha}, (\mathrm{sign}(\sigma_i(\alpha))_{\sigma_i \in \mathfrak{m}_\infty})\right) \ ,$$

where $\mathrm{sign}(x)$ is set equal to 0 or 1 in $\mathbb{F}_2$ according to whether x is positive or negative. The strong approximation theorem in Dedekind domains (more precisely, Corollary 1.2.9) tells us that ρ is surjective. Thus, any element of $(\mathbb{Z}_K/\mathfrak{m})^*$ can be represented as $\rho(\alpha)$ for some $\alpha \in \mathbb{Z}_K$.

In algorithmic practice, this is *not* the nicest way to represent an element of $(\mathbb{Z}_K/\mathfrak{m})^*$. We will see that it is much better simply to keep the initial definition and to represent an element as $(\overline{\alpha}, (s_1, \ldots, s_{|\mathfrak{m}_\infty|}))$, where $\overline{\alpha}$ is the class of α modulo $\mathfrak{m}_0$ and $s_i \in \mathbb{F}_2$ (see Sections 4.2.4 and 4.3.2).

Definition 3.2.2. *Let* $\mathfrak{m}$ *be a modulus in* K.

(1) *If* $\alpha \in K^*$ *is coprime to* $\mathfrak{m}$ *and the modulus* $\mathfrak{m}$ *is understood, we write* $\overline{\alpha}$ *instead of* $\rho(\alpha)$ *as defined above.*

(2) *If* $\alpha \in K^*$, *we say that*

$$\alpha \equiv 1 \pmod{{}^*\mathfrak{m}}$$

if for all $\mathfrak{p}$ *dividing* $\mathfrak{m}_0$ *we have* $v_\mathfrak{p}(\alpha - 1) \geq v_\mathfrak{p}(\mathfrak{m}_0)$, *and if for all embeddings* $\sigma_i \in \mathfrak{m}_\infty$ *we have* $\sigma_i(\alpha) > 0$. *We will write* $K_\mathfrak{m}^*$ *for the group of such* α.

(3) *If* α *and* β *are in* K^*, *we say that* $\alpha \equiv \beta \pmod{{}^*\mathfrak{m}}$ *if* α *and* β *are coprime to* $\mathfrak{m}$ *and if* $\alpha/\beta \equiv 1 \pmod{{}^*\mathfrak{m}}$.

Remarks

(1) The condition $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}_0}$ is equivalent to $\alpha \equiv 1 \pmod{\mathfrak{m}_0}$ only if we restrict to $\alpha \in \mathbb{Z}_K$, which would not be usable for our needs.

(2) If α and β are coprime to $\mathfrak{m}$, the condition $\alpha \equiv \beta \pmod{{}^*\mathfrak{m}}$ is clearly equivalent to $\overline{\alpha} = \overline{\beta}$, where $\overline{\alpha}$ is defined above. When $\mathfrak{m}_\infty$ is nonempty, this is *not* a property of the number $\alpha - \beta$ alone.

We will write $P_\mathfrak{m}(K)$ (or $P_\mathfrak{m}$ if the field K is understood) for the set of all (fractional) principal ideals of $\mathbb{Z}_K$ that can be generated by an element α such that $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}}$; in other words, ideals of the form $\alpha\mathbb{Z}_K$ for such an α. It is clear that $P_\mathfrak{m}(K)$ is a subgroup of $I_\mathfrak{m}(K)$, sometimes called the *ray group* of $\mathfrak{m}$.

Let $\mathfrak{a} \in P_\mathfrak{m}(K)$. It is clear that $\mathfrak{a} = \alpha\mathbb{Z}_K = \beta\mathbb{Z}_K$ with α and β in $K_\mathfrak{m}^*$ if and only if β/α is a unit u such that $u \in K_\mathfrak{m}^*$. These units form a subgroup of the unit group $U(K)$, which we will denote by $U_\mathfrak{m}(K) = U(K) \cap K_\mathfrak{m}^*$. From the definitions, it is clear that we have the following exact sequence, which generalizes the corresponding exact sequence for the trivial modulus $\mathfrak{m}$, where $\mathfrak{m}_0 = \mathbb{Z}_K$ and $\mathfrak{m}_\infty = \varnothing$:

$$1 \longrightarrow U_\mathfrak{m}(K) \longrightarrow K_\mathfrak{m}^* \longrightarrow P_\mathfrak{m}(K) \longrightarrow 1 \ .$$

Finally, we define the *ray class group* $Cl_\mathfrak{m}(K)$ by the formula $Cl_\mathfrak{m}(K) = I_\mathfrak{m}(K)/P_\mathfrak{m}(K)$, so that we also have the exact sequence

$$1 \longrightarrow P_\mathfrak{m}(K) \longrightarrow I_\mathfrak{m}(K) \longrightarrow Cl_\mathfrak{m}(K) \longrightarrow 1 \ .$$

The following proposition will be crucial for us in the sequel. By abuse of notation, we write again ρ for the restriction of ρ to the unit group $U(K)$.

Proposition 3.2.3. *We have the following five-term exact sequence*

$$1 \longrightarrow U_{\mathfrak{m}}(K) \longrightarrow U(K) \xrightarrow{\rho} (\mathbb{Z}_K/\mathfrak{m})^* \xrightarrow{\psi} Cl_{\mathfrak{m}}(K) \xrightarrow{\phi} Cl(K) \longrightarrow 1$$

(recall that $U_{\mathfrak{m}}(K) = U(K) \cap K_{\mathfrak{m}}^*$ *is the group of units congruent to* 1 (mod $^*\mathfrak{m}$)).

All the maps are essentially clear, except perhaps for ψ which sends an element $\rho(\alpha) \in (\mathbb{Z}_K/\mathfrak{m})^*$ to the ideal class of $\alpha\mathbb{Z}_K$ in $Cl_{\mathfrak{m}}(K)$ (which is usually not the trivial class, since $\alpha \notin K_{\mathfrak{m}}^*$ in general). Note that this is *not* the map (which could be considered more natural from the algorithmic representation; see Sections 4.2.4 and 4.3.2) that sends $(\overline{\alpha}, s_1, \ldots, s_{|\mathfrak{m}_\infty|})$ to the ideal class of $\alpha\mathbb{Z}_K$. In fact, this map would not even be well-defined.

Proof. The kernel of ρ is by definition the set of units congruent to 1 (mod $^*\mathfrak{m}$) and so is equal to $U_{\mathfrak{m}}(K)$. Furthermore, the map ψ that we have just described is well-defined (for the other maps this is clear). Indeed, if $\rho(\alpha) = \rho(\beta)$, this means that $\alpha \equiv \beta \pmod{\mathfrak{m}_0}$, that α and β are coprime to $\mathfrak{m}_0$, and that $\operatorname{sign}(\sigma_i(\alpha)) = \operatorname{sign}(\sigma_i(\beta))$ for $\sigma_i \in \mathfrak{m}_\infty$. These conditions mean precisely that $\alpha/\beta \equiv 1 \pmod{^*\mathfrak{m}}$, and so the principal ideals $\alpha\mathbb{Z}_K$ and $\beta\mathbb{Z}_K$ are in the same ideal class modulo $P_{\mathfrak{m}}(K)$ or, equivalently, have the same image in $Cl_{\mathfrak{m}}(K)$.

Assume now that $\rho(\alpha) \in (\mathbb{Z}_K/\mathfrak{m})^*$ is sent to the unit element of $Cl_{\mathfrak{m}}(K)$. This means that $\alpha\mathbb{Z}_K \in P_{\mathfrak{m}}(K)$, so there exists $\beta \equiv 1 \pmod{^*\mathfrak{m}}$ such that $\alpha\mathbb{Z}_K = \beta\mathbb{Z}_K$, hence $u = \alpha/\beta$ is a unit; in other words, it belongs to $U(K)$. Since $\beta \equiv 1 \pmod{^*\mathfrak{m}}$, we have $\overline{u} = \overline{\alpha}$ in $(\mathbb{Z}_K/\mathfrak{m}_0)^*$, but also $\operatorname{sign}(\sigma_i(u)) = \operatorname{sign}(\sigma_i(\alpha))$. Thus, $\rho(u) = \rho(\alpha)$, and so the kernel of ψ is indeed equal to the image of ρ.

Now let $\overline{\mathfrak{a}}$ be an ideal class in $Cl_{\mathfrak{m}}(K)$ which is sent to the trivial class in $Cl(K)$. This simply means that $\mathfrak{a} = \alpha\mathbb{Z}_K$ is a principal ideal coprime to $\mathfrak{m}$, hence α is also coprime to $\mathfrak{m}$, and this shows that the kernel of ϕ is the image of ψ.

Finally, the surjectivity of ϕ follows from the approximation theorem in Dedekind domains, more precisely from Corollary 1.2.11, since one can choose as representative of an ideal class an ideal coprime to $\mathfrak{m}$. □

As in the case of ordinary integers, we can define the Euler ϕ-function for moduli by $\phi(\mathfrak{m}) = |(\mathbb{Z}_K/\mathfrak{m})^*|$. If $\mathfrak{m}_0 = \prod_{\mathfrak{p}|\mathfrak{m}_0} \mathfrak{p}^{a_\mathfrak{p}}$, we have the following immediate generalization of the usual formula over $\mathbb{Z}$ (Exercise 4):

$$\phi(\mathfrak{m}) = 2^{|\mathfrak{m}_\infty|} \prod_{\mathfrak{p}|\mathfrak{m}_0} \mathcal{N}(\mathfrak{p})^{a_\mathfrak{p}-1}(\mathcal{N}(\mathfrak{p})-1) = 2^{|\mathfrak{m}_\infty|}\mathcal{N}(\mathfrak{m}_0) \prod_{\mathfrak{p}|\mathfrak{m}_0} (1-\mathcal{N}(\mathfrak{p})^{-1}) \ .$$

Corollary 3.2.4. *The ray class group is finite. Its cardinality, which we will denote by* $h_{\mathfrak{m}}(K)$ *(or simply by* $h_{\mathfrak{m}}$ *when the field is understood), is given by*

$$h_{\mathfrak{m}}(K) = h(K)\frac{\phi(\mathfrak{m})}{[U(K) : U_{\mathfrak{m}}(K)]} \ .$$

In particular,

$$h(K) \mid h_{\mathfrak{m}}(K) \mid h(K)\phi(\mathfrak{m}) \ .$$

Proof. The proof is clear from the proposition. □

This corollary can be seen as a first approach in computing the ray class group, but later we shall see a method that gives the full result (including the structure, as we have done for the ordinary class group). Of course, even if we want only the cardinality, the main problem is the computation of the index $[U(K) : U_{\mathfrak{m}}(K)]$.

3.3 Congruence Subgroups: One Side of Class Field Theory

In this section, the field K is fixed, so we write $I_{\mathfrak{m}}$ instead of $I_{\mathfrak{m}}(K)$, $P_{\mathfrak{m}}$ instead of $P_{\mathfrak{m}}(K)$, $Cl_{\mathfrak{m}}$ instead of $Cl_{\mathfrak{m}}(K)$, and so on.

3.3.1 Motivation for the Equivalence Relation

We will ultimately want a bijection between two sets: the two "sides" of class field theory. We must describe both sets, and we start with the easy side.

Recall that to describe unramified extensions we used subgroups of the group of fractional ideals containing the group of principal ideals. In our more general situation, we do exactly the same. We will say that C is a *congruence subgroup* modulo $\mathfrak{m}$ if C is a group of fractional ideals such that

$$P_{\mathfrak{m}} \subset C \subset I_{\mathfrak{m}} \ .$$

(Some authors call such a C an *ideal group* modulo $\mathfrak{m}$.)

We can also consider the set of classes $\overline{C} = C/P_{\mathfrak{m}} \subset Cl_{\mathfrak{m}}$; so if desired we can consider a congruence subgroup as a subgroup of the ray class group $Cl_{\mathfrak{m}}$. To indicate the modulus to which C corresponds, we will usually write $(\mathfrak{m}, C)$ for a congruence subgroup modulo $\mathfrak{m}$.

Keeping in mind the example of Hilbert class fields, we want to introduce an equivalence relation between congruence subgroups, so that subgroups in the same equivalence class define the same number field. We proceed as follows: if $(\mathfrak{m}, C)$ is a congruence subgroup, one of the main results of class field theory will tell us that, exactly as in the Hilbert class field situation, there exists a generalized Hilbert class field $K(\mathfrak{m})$ (which we will call the *ray class field* for the modulus $\mathfrak{m}$) such that, among other properties, $\mathrm{Gal}(K(\mathfrak{m})/K) \simeq Cl_{\mathfrak{m}}$. Let L be the Abelian extension corresponding to $(\mathfrak{m}, C)$; in other words $L = K(\mathfrak{m})^{\overline{C}}$ is the fixed field of $K(\mathfrak{m})$ by $\overline{C}$, so that

$$\mathrm{Gal}(L/K) \simeq Cl_{\mathfrak{m}}/\overline{C} \simeq I_{\mathfrak{m}}/C \ .$$

The subextensions of $K(\mathfrak{m})$ are in one-to-one correspondence with the congruence subgroups C by Galois theory.

If $\mathfrak{m}$ is a multiple of $\mathfrak{n}$, class field theory tells us that $K(\mathfrak{n})$ can be considered as a subfield of $K(\mathfrak{m})$. Thus, if $\mathfrak{m}_1$ and $\mathfrak{m}_2$ are any moduli, we can consider $K(\mathfrak{m}_1)$ and $K(\mathfrak{m}_2)$ as subfields of a single $K(\mathfrak{m})$, for example, with $\mathfrak{m} = \mathfrak{m}_1\mathfrak{m}_2$ (in fact any common multiple of $\mathfrak{m}_1$ and $\mathfrak{m}_2$ will do). We can then say that two congruence subgroups $(\mathfrak{m}_1, C_1)$ and $(\mathfrak{m}_2, C_2)$ are equivalent if they define the same number field or, equivalently, if $K(\mathfrak{m}_1)^{\overline{C_1}} = K(\mathfrak{m}_2)^{\overline{C_2}}$, considered as subfields of $K(\mathfrak{m})$ (see diagram below). Note that we ask that the fields be *identical*, not only isomorphic, and this is why we need to embed the whole situation in a single number field $K(\mathfrak{m})$. From this definition, it is clear that it is an equivalence relation.

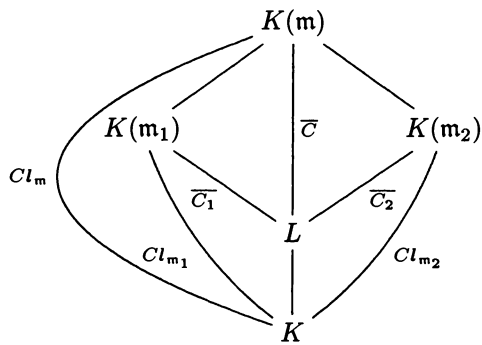

Let us transform this definition into one that does not involve the field $K(\mathfrak{m})$, since after all we do not yet know the results of class field theory. Let $L = K(\mathfrak{m}_1)^{\overline{C_1}} = K(\mathfrak{m}_2)^{\overline{C_2}}$. By Galois theory, $L = K(\mathfrak{m})^{\overline{C}}$ for some congruence subgroup C of $\mathfrak{m}$. The equivalence relation means that the natural maps from $Cl_{\mathfrak{m}}/\overline{C}$ to $Cl_{\mathfrak{m}_i}/\overline{C_i}$ are isomorphisms for $i = 1$ and $i = 2$. It can easily be shown as a consequence of the approximation theorem for Dedekind domains that the maps in question are always surjective (see Exercise 5). Thus, we have $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$ if and only if the maps are injective, and this is easily seen to be equivalent to $I_{\mathfrak{m}} \cap C_i = C$ for $i = 1$ and $i = 2$. Now choose $\mathfrak{m} = \mathfrak{m}_1\mathfrak{m}_2$. Since $C_i \subset I_{\mathfrak{m}_i}$, we clearly have $I_{\mathfrak{m}_1\mathfrak{m}_2} \cap C_1 = I_{\mathfrak{m}_2} \cap C_1$, and similarly for C_2. Thus, $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$ if and only if $I_{\mathfrak{m}_2} \cap C_1 = I_{\mathfrak{m}_1} \cap C_2$. This does not involve any extraneous number fields or moduli, so we can forget about the motivation coming from class field theory and start from scratch the study of the relation $\sim$ between congruence subgroups.

3.3.2 Study of the Equivalence Relation

We begin with the following lemma, which is an immediate consequence of the strong approximation theorem.

Lemma 3.3.1. *Let* $\mathfrak{m}_1$ *and* $\mathfrak{m}_2$ *be two arbitrary moduli, and let* $\mathfrak{a} \in I_{\mathfrak{m}_1}$. *There exists* $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}_1}$ *such that* $\alpha\mathfrak{a}$ *is an integral ideal coprime to* $\mathfrak{m}_1\mathfrak{m}_2$.

Proof. For the infinite places, we of course simply ask that $\sigma(\alpha) > 0$ for all $\sigma \mid \mathfrak{m}_1$. For the finite places, let $\mathfrak{p}$ be a prime ideal. If $\mathfrak{p} \mid \mathfrak{m}_1$, we ask that $v_{\mathfrak{p}}(\alpha - 1) \geq v_{\mathfrak{p}}(\mathfrak{m}_1)$. Then, since $\mathfrak{a} \in I_{\mathfrak{m}_1}$ we necessarily have $v_{\mathfrak{p}}(\alpha) = 0 = -v_{\mathfrak{p}}(\mathfrak{a})$. If $\mathfrak{p} \nmid \mathfrak{m}_1$ and $\mathfrak{p} \mid \mathfrak{m}_2$, we ask that $v_{\mathfrak{p}}(\alpha) = -v_{\mathfrak{p}}(\mathfrak{a})$. Finally, if $\mathfrak{p} \nmid \mathfrak{m}_1\mathfrak{m}_2$, we ask that $v_{\mathfrak{p}}(\alpha) \geq -v_{\mathfrak{p}}(\mathfrak{a})$. Thus, the conditions are compatible. The strong approximation theorem (more precisely, Corollary 1.2.9) shows the existence of an α with the desired properties, and it is clear that such an α satisfies the conditions of the lemma. □

Corollary 3.3.2. *Let* $\mathfrak{m}_1$ *and* $\mathfrak{m}_2$ *be two arbitrary moduli.*

(1) *We have* $I_{\mathfrak{m}_2} \subset I_{\mathfrak{m}_1} P_{\mathfrak{m}_2}$ (*and, of course, also* $I_{\mathfrak{m}_1} \subset I_{\mathfrak{m}_2} P_{\mathfrak{m}_1}$).
(2) *If* $\mathfrak{m}_2 \mid \mathfrak{m}_1$ *and* C_2 *is a congruence subgroup modulo* $\mathfrak{m}_2$ (*for example,* $C_2 = P_{\mathfrak{m}_2}$), *then we have the* equality $I_{\mathfrak{m}_2} = I_{\mathfrak{m}_1} C_2$.

Proof. If $\mathfrak{a} \in I_{\mathfrak{m}_2}$, by the above lemma, we can find $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}_2}$ such that $\alpha\mathfrak{a} \in I_{\mathfrak{m}_1}$. Since $\alpha\mathbb{Z}_K \in P_{\mathfrak{m}_2}$, we thus have $\mathfrak{a} \in I_{\mathfrak{m}_1} P_{\mathfrak{m}_2}$, so

$$I_{\mathfrak{m}_2} \subset I_{\mathfrak{m}_1} P_{\mathfrak{m}_2} \subset I_{\mathfrak{m}_1} C_2$$

for any congruence subgroup C_2 modulo $\mathfrak{m}_2$. If $\mathfrak{m}_2 \mid \mathfrak{m}_1$, then $I_{\mathfrak{m}_1} \subset I_{\mathfrak{m}_2}$ and $C_2 \subset I_{\mathfrak{m}_2}$, so the reverse inclusion is also valid, thus proving the corollary. □

Referring to the discussion of the preceding section, we can now set the following definition.

Definition 3.3.3. *We will say that two congruence subgroups* $(\mathfrak{m}_1, C_1)$ *and* $(\mathfrak{m}_2, C_2)$ *of* K *are* equivalent, *and write* $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$, *if*

$$I_{\mathfrak{m}_2} \cap C_1 = I_{\mathfrak{m}_1} \cap C_2 \ .$$

The following proposition is essential for this definition to make sense.

Proposition 3.3.4. (1) *The relation* $\sim$ *defined above between congruence subgroups is an equivalence relation.*
(2) *If* $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$, *then* $I_{\mathfrak{m}_1}/C_1 \simeq I_{\mathfrak{m}_2}/C_2$; *in other words, we have* $Cl_{\mathfrak{m}_1}/\overline{C_1} \simeq Cl_{\mathfrak{m}_2}/\overline{C_2}$.

Proof. (1). The reflexivity and symmetry are trivial, so the only thing to prove is the transitivity. Assume that $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$ and $(\mathfrak{m}_2, C_2) \sim (\mathfrak{m}_3, C_3)$; in other words, that $I_{\mathfrak{m}_2} \cap C_1 = I_{\mathfrak{m}_1} \cap C_2$ and $I_{\mathfrak{m}_3} \cap C_2 = I_{\mathfrak{m}_2} \cap C_3$. We must prove that $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_3, C_3)$ or, equivalently, that $I_{\mathfrak{m}_3} \cap C_1 = I_{\mathfrak{m}_1} \cap C_3$.

Let $\mathfrak{a} \in I_{\mathfrak{m}_3} \cap C_1$. Since $\mathfrak{a} \in C_1 \subset I_{\mathfrak{m}_1}$ we must only show that $\mathfrak{a} \in C_3$. By Lemma 3.3.1, since $\mathfrak{a} \in I_{\mathfrak{m}_1\mathfrak{m}_3}$, we can find $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}_1\mathfrak{m}_3}$ such that

$\alpha\mathfrak{a}$ is an integral ideal coprime to $\mathfrak{m}_1\mathfrak{m}_2\mathfrak{m}_3$. Since $\alpha\mathbb{Z}_K \in P_{\mathfrak{m}_1} \subset C_1$ and C_1 is a group, it follows that $\alpha\mathfrak{a} \in C_1$, and since $\alpha\mathfrak{a}$ is coprime to $\mathfrak{m}_2$ we have $\alpha\mathfrak{a} \in I_{\mathfrak{m}_2} \cap C_1 = I_{\mathfrak{m}_1} \cap C_2$, so $\alpha\mathfrak{a} \in C_2$. But since $\alpha\mathfrak{a}$ is also coprime to $\mathfrak{m}_3$, we have $\alpha\mathfrak{a} \in I_{\mathfrak{m}_3} \cap C_2 = I_{\mathfrak{m}_2} \cap C_3$, so $\alpha\mathfrak{a} \in C_3$. Finally, since $\alpha\mathbb{Z}_K \in P_{\mathfrak{m}_3}$ and C_3 is a group containing $P_{\mathfrak{m}_3}$, we deduce that $\mathfrak{a} = \alpha\mathfrak{a}\alpha^{-1} \in C_3$, as was to be proved. We have proved the inclusion $I_{\mathfrak{m}_3} \cap C_1 \subset I_{\mathfrak{m}_1} \cap C_3$, and the reverse inclusion follows by symmetry.

(2). Once again by Lemma 3.3.1, for any $\mathfrak{a} \in I_{\mathfrak{m}_1}$ there exists $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}_1}$ such that $(\alpha\mathfrak{a}, \mathfrak{m}_2) = 1$. Although α is not unique, the class of $\alpha\mathfrak{a}$ modulo C_2 is well-defined since if α and α' are two such elements, α'/α is coprime to $\mathfrak{m}_2$ and is in $P_{\mathfrak{m}_1} \subset C_1$ and hence belongs to $C_1 \cap I_{\mathfrak{m}_2} = C_2 \cap I_{\mathfrak{m}_1}$ and hence to C_2.

The same reasoning shows that the map thus defined induces a well-defined map from $I_{\mathfrak{m}_1}/C_1$ to $I_{\mathfrak{m}_2}/C_2$ and that this map is an isomorphism.

□

Note that the isomorphism between $I_{\mathfrak{m}_1}/C_1$ and $I_{\mathfrak{m}_2}/C_2$ is *canonical*, meaning that it does not depend on any special choices we have made.

The following proposition explains what happens in the important special case when one of the moduli divides the other.

Proposition 3.3.5. (1) *Let* $(\mathfrak{m}_1, C_1)$ *be a congruence subgroup, and let* $\mathfrak{m}_2$ *be a divisor of* $\mathfrak{m}_1$ (*see Definition 3.2.1*). *There exists a congruence subgroup* C_2 *modulo* $\mathfrak{m}_2$ *such that* $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$ *if and only if*

$$I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset C_1 \ .$$

If this condition is satisfied, we necessarily have $C_2 = C_1 P_{\mathfrak{m}_2}$.

(2) *Conversely, if* $(\mathfrak{m}_2, C_2)$ *is a congruence subgroup and* $\mathfrak{m}_1$ *is a multiple of* $\mathfrak{m}_2$, *there exists a unique congruence subgroup* C_1 *modulo* $\mathfrak{m}_1$ *such that* $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$ *given by* $C_1 = C_2 \cap I_{\mathfrak{m}_1}$.

Proof. (1). If $\mathfrak{m}_2 \mid \mathfrak{m}_1$, we have $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$ if and only if $I_{\mathfrak{m}_1} \cap C_2 = C_1$. Thus, since $P_{\mathfrak{m}_2} \subset C_2$,

$$I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset I_{\mathfrak{m}_1} \cap C_2 = C_1 \ .$$

Furthermore,

$$C_1 P_{\mathfrak{m}_2} = (I_{\mathfrak{m}_1} \cap C_2) P_{\mathfrak{m}_2} \subset C_2 P_{\mathfrak{m}_2} = C_2 \ .$$

Set $C_2' = C_1 P_{\mathfrak{m}_2}$. Then I claim that $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2')$. Indeed, this means that $C_1 = I_{\mathfrak{m}_2} \cap C_1 = I_{\mathfrak{m}_1} \cap C_1 P_{\mathfrak{m}_2}$ or, equivalently (since the other inclusion is obvious), that $I_{\mathfrak{m}_1} \cap C_1 P_{\mathfrak{m}_2} \subset C_1$. But this follows from the inclusion $I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset C_1$ by multiplying both sides by the group C_1.

Thus, since our equivalence relation is transitive, we have $(\mathfrak{m}_2, C_2) \sim (\mathfrak{m}_2, C_2')$, which of course means that $C_2 = C_2'$, and so that $C_2 = C_1 P_{\mathfrak{m}_2}$, as claimed.

Conversely, if we assume $I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset C_1$ and $C_2 = C_1 P_{\mathfrak{m}_2}$, then by multiplication by C_1 we get as above $I_{\mathfrak{m}_1} \cap C_2 \subset C_1$; since the reverse inclusion is trivial, we have equality, proving (1).

Statement (2) is a trivial consequence of the definition. □

Notation. The following notation will be very useful. If $(\mathfrak{m}, C)$ is a congruence subgroup, we will write $h_{\mathfrak{m},C} = |I_{\mathfrak{m}}/C| = |Cl_{\mathfrak{m}}/\overline{C}|$.

Proposition 3.3.6. *Let* $\mathfrak{m}_1$ *and* $\mathfrak{m}_2$ *be two moduli such that* $\mathfrak{m}_2 \mid \mathfrak{m}_1$, *and let* C_1 *and* C_2 *be two congruence subgroups modulo* $\mathfrak{m}_1$ *and* $\mathfrak{m}_2$, *respectively, such that* $C_1 \subset C_2$.

(1) *We have a canonical isomorphism*

$$I_{\mathfrak{m}_2}/C_2 \simeq \frac{I_{\mathfrak{m}_1}/C_1}{(I_{\mathfrak{m}_1} \cap C_2)/C_1} \ .$$

In particular, we have

$$\frac{h_{\mathfrak{m}_1,C_1}}{h_{\mathfrak{m}_2,C_2}} = |(I_{\mathfrak{m}_1} \cap C_2)/C_1| \ .$$

(2) *We have* $h_{\mathfrak{m}_1,C_1} = h_{\mathfrak{m}_2,C_2}$ *if and only if* $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$.

Proof. Applying Corollary 3.3.2, we have

$$\frac{I_{\mathfrak{m}_2}}{C_2} \simeq \frac{I_{\mathfrak{m}_1} C_2}{C_2} \simeq \frac{I_{\mathfrak{m}_1}}{I_{\mathfrak{m}_1} \cap C_2} \simeq \frac{I_{\mathfrak{m}_1}/C_1}{(I_{\mathfrak{m}_1} \cap C_2)/C_1} \ ,$$

proving (1).

(2). We have already seen in Proposition 3.3.4 that if $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$, then $h_{\mathfrak{m}_1,C_1} = h_{\mathfrak{m}_2,C_2}$ even when $\mathfrak{m}_2$ does not divide $\mathfrak{m}_1$. Conversely, assume that we have this equality. By (1) we have $I_{\mathfrak{m}_1} \cap C_2 = C_1$, and in particular $I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset C_1$, so by Proposition 3.3.5 we deduce that $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$. □

Note that this proposition is clearly *not* true if we do not assume that $\mathfrak{m}_2 \mid \mathfrak{m}_1$.

Corollary 3.3.7. *Let* $\mathfrak{m}_1$ *and* $\mathfrak{m}_2$ *be two moduli such that* $\mathfrak{m}_2 \mid \mathfrak{m}_1$, *let* C_1 *be a congruence subgroup modulo* $\mathfrak{m}_1$, *and let* $C_2 = C_1 P_{\mathfrak{m}_2}$.

Then $|(I_{\mathfrak{m}_1} \cap C_2)/C_1| = h_{\mathfrak{m}_1,C_1}/h_{\mathfrak{m}_2,C_2}$ *divides* $\phi(\mathfrak{m}_1)/\phi(\mathfrak{m}_2)$.

Proof. By the above proposition and Corollary 3.2.4, we have

$$\left|\frac{I_{\mathfrak{m}_1}\cap C_2}{C_1}\right| = \frac{h_{\mathfrak{m}_1,C_1}}{h_{\mathfrak{m}_2,C_2}} = \frac{\phi(\mathfrak{m}_1)}{\phi(\mathfrak{m}_2)}\frac{1}{[U_{\mathfrak{m}_2}(K):U_{\mathfrak{m}_1}(K)](\left|\overline{C_1}\right|/\left|\overline{C_2}\right|)}\ ,$$

where $\overline{C_i} = C_i/P_{\mathfrak{m}_i}$ for $i = 1$ and $i = 2$. Using the same proof as in (1) of the above proposition and the hypothesis $C_2 = C_1P_{\mathfrak{m}_2}$ instead of Corollary 3.3.2, we find a canonical isomorphism

$$\overline{C_2} = C_2/P_{\mathfrak{m}_2} \simeq \frac{C_1/P_{\mathfrak{m}_1}}{(C_1\cap P_{\mathfrak{m}_2})/P_{\mathfrak{m}_1}} = \frac{\overline{C_1}}{(C_1\cap P_{\mathfrak{m}_2})/P_{\mathfrak{m}_1}}\ ,$$

showing in particular that $\left|\overline{C_2}\right|$ divides $\left|\overline{C_1}\right|$, and the corollary follows. □

The next important result we will need about congruence subgroups is the existence of a GCD. Note first that if $\mathfrak{m}_1$ and $\mathfrak{m}_2$ are two moduli, $\gcd(\mathfrak{m}_1,\mathfrak{m}_2)$ is well-defined: we take the sum of the corresponding integral ideals and the intersection of the places at infinity. This is clearly the largest modulus dividing $\mathfrak{m}_1$ and $\mathfrak{m}_2$.

Before giving the result, we need a lemma.

Lemma 3.3.8. *Let $\mathfrak{m}_1$ and $\mathfrak{m}_2$ be two moduli, and let α_1 and α_2 be elements of K^*. A necessary and sufficient condition for the existence of $\beta\in K^*$ such that*

$$\beta\equiv\alpha_1 \pmod{{}^*\mathfrak{m}_1} \quad \textit{and} \quad \beta\equiv\alpha_2 \pmod{{}^*\mathfrak{m}_2}$$

is that $\alpha_1\equiv\alpha_2 \pmod{{}^*\mathfrak{n}}$ with $\mathfrak{n} = \gcd(\mathfrak{m}_1,\mathfrak{m}_2)$.

Proof. Recall that $\beta\equiv\alpha \pmod{{}^*\mathfrak{m}}$ means that for all finite places $\mathfrak{p}$ dividing $\mathfrak{m}$, we have $v_{\mathfrak{p}}(\beta/\alpha-1)\geq v_{\mathfrak{p}}(\mathfrak{m})$, and for infinite places σ dividing $\mathfrak{m}$ we have $\operatorname{sign}(\sigma(\beta/\alpha))>0$. The condition of the lemma is clearly necessary. Conversely, assume that it is satisfied. In particular, it implies that $v_{\mathfrak{p}}(\alpha_1/\alpha_2)=0$ for every $\mathfrak{p}\mid\mathfrak{n}$.

For each finite $\mathfrak{p}$ dividing $\mathfrak{m}_1$ or $\mathfrak{m}_2$, we set the following approximation conditions on β. If $\mathfrak{p}\mid\mathfrak{m}_1$ and $\mathfrak{p}\nmid\mathfrak{m}_2$ (resp., $\mathfrak{p}\mid\mathfrak{m}_2$ and $\mathfrak{p}\nmid\mathfrak{m}_1$), we ask that $v_{\mathfrak{p}}(\beta-\alpha_1)\geq v_{\mathfrak{p}}(\alpha_1)+v_{\mathfrak{p}}(\mathfrak{m}_1)$ (resp., $v_{\mathfrak{p}}(\beta-\alpha_2)\geq v_{\mathfrak{p}}(\alpha_2)+v_{\mathfrak{p}}(\mathfrak{m}_2)$). If $\mathfrak{p}\mid\mathfrak{m}_1$ and $\mathfrak{p}\mid\mathfrak{m}_2$ or, equivalently, if $\mathfrak{p}\mid\mathfrak{n}$, assume first that $v_{\mathfrak{p}}(\mathfrak{m}_1)\leq v_{\mathfrak{p}}(\mathfrak{m}_2)$. We ask that $v_{\mathfrak{p}}(\beta-\alpha_2)\geq v_{\mathfrak{p}}(\alpha_2)+v_{\mathfrak{p}}(\mathfrak{m}_2)$, which implies

$$\begin{aligned} v_{\mathfrak{p}}(\beta-\alpha_1) &= v_{\mathfrak{p}}(\beta-\alpha_2+\alpha_2-\alpha_1)\geq\min(v_{\mathfrak{p}}(\beta-\alpha_2),v_{\mathfrak{p}}(\alpha_2-\alpha_1))\\ &\geq\min(v_{\mathfrak{p}}(\alpha_2)+v_{\mathfrak{p}}(\mathfrak{m}_2),v_{\mathfrak{p}}(\alpha_2)+v_{\mathfrak{p}}(\mathfrak{m}_1))\\ &= v_{\mathfrak{p}}(\alpha_2)+v_{\mathfrak{p}}(\mathfrak{m}_1)=v_{\mathfrak{p}}(\alpha_1)+v_{\mathfrak{p}}(\mathfrak{m}_1)\end{aligned}$$

since $v_{\mathfrak{p}}(\alpha_1)=v_{\mathfrak{p}}(\alpha_2)$ in this case.

If $v_{\mathfrak{p}}(\mathfrak{m}_2)<v_{\mathfrak{p}}(\mathfrak{m}_1))$, we ask that $v_{\mathfrak{p}}(\beta-\alpha_1)\geq v_{\mathfrak{p}}(\alpha_1)+v_{\mathfrak{p}}(\mathfrak{m}_1)$, and in the same way this implies $v_{\mathfrak{p}}(\beta-\alpha_2)\geq v_{\mathfrak{p}}(\alpha_2)+v_{\mathfrak{p}}(\mathfrak{m}_2)$.

Finally, for the infinite places, we ask that $\operatorname{sign}(\sigma(\beta)) = \operatorname{sign}(\sigma(\alpha_i))$ if σ divides $\mathfrak{m}_i$. These conditions are compatible when σ divides both $\mathfrak{m}_1$ and $\mathfrak{m}_2$ since in that case $\operatorname{sign}(\sigma(\alpha_1/\alpha_2)) > 0$.

Thus, we can apply the strong approximation theorem (more precisely, Corollary 1.2.9) to show the existence of β satisfying our conditions, and we will have $\beta \equiv \alpha_1 \pmod{{}^*\mathfrak{m}_1}$ and $\beta \equiv \alpha_2 \pmod{{}^*\mathfrak{m}_2}$. □

This lemma allows us to prove the last statement that we will need about congruence subgroups.

Proposition 3.3.9. *Let* $(\mathfrak{m}_1, C_1)$ *and* $(\mathfrak{m}_2, C_2)$ *be two congruence subgroups such that* $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$, *and let* $\mathfrak{n} = \gcd(\mathfrak{m}_1, \mathfrak{m}_2)$. *There exists a unique congruence subgroup* C *modulo* $\mathfrak{n}$ *such that* $(\mathfrak{n}, C) \sim (\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$, *and* C *is given by* $C = C_1 P_{\mathfrak{n}} = C_2 P_{\mathfrak{n}}$. *The congruence subgroup* $(\mathfrak{n}, C)$ *will be called the GCD of the congruence subgroups* $(\mathfrak{m}_1, C_1)$ *and* $(\mathfrak{m}_2, C_2)$ (*note that the GCD is defined only when the congruence subgroups are equivalent*).

Proof. Set $\mathfrak{m} = \mathfrak{m}_1\mathfrak{m}_2$. By Proposition 3.3.5 (2), if we set $D = I_{\mathfrak{m}} \cap C_1 = I_{\mathfrak{m}} \cap C_2$, we have $(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2) \sim (\mathfrak{m}, D)$. Applying part (1) of the same proposition, we deduce that

$$P_{\mathfrak{m}_1} \cap I_{\mathfrak{m}} \subset D \quad \text{and} \quad P_{\mathfrak{m}_2} \cap I_{\mathfrak{m}} \subset D \ .$$

By the same proposition, to show the existence of C, we must show that $P_{\mathfrak{n}} \cap I_{\mathfrak{m}} \subset D$. Thus, let $\mathfrak{a} \in P_{\mathfrak{n}} \cap I_{\mathfrak{m}}$. Since $\mathfrak{a} \in P_{\mathfrak{n}}$, there exists $\alpha \equiv 1 \pmod{{}^*\mathfrak{n}}$ such that $\mathfrak{a} = \alpha\mathbb{Z}_K$. By Lemma 3.3.8, this implies the existence of $\beta \in K^*$ such that $\beta \equiv \alpha \pmod{{}^*\mathfrak{m}_1}$ and $\beta \equiv 1 \pmod{{}^*\mathfrak{m}_2}$. Since $\mathfrak{a} \in I_{\mathfrak{m}}$, α is coprime to $\mathfrak{m}$; hence β is coprime both to $\mathfrak{m}_1$ and to $\mathfrak{m}_2$ and hence to $\mathfrak{m}$. It follows that $(\beta/\alpha)\mathbb{Z}_K \in P_{\mathfrak{m}_1} \cap I_{\mathfrak{m}} \subset D$. Since $\beta\mathbb{Z}_K \in P_{\mathfrak{m}_2} \cap I_{\mathfrak{m}} \subset D$ and D is a group, we obtain $\alpha\mathbb{Z}_K = (\beta\mathbb{Z}_K)((\beta/\alpha)\mathbb{Z}_K)^{-1} \in D$, as was to be proved. Proposition 3.3.5 thus shows the existence of a unique congruence subgroup C modulo $\mathfrak{n}$ such that $(\mathfrak{n}, C) \sim (\mathfrak{m}, D)$, hence by transitivity $(\mathfrak{n}, C) \sim (\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$, and the uniqueness statement of the same proposition implies that $C = C_1 P_{\mathfrak{n}} = C_2 P_{\mathfrak{n}}$. □

Corollary 3.3.10. *Let* $\mathcal{C}$ *be an equivalence class of congruence subgroups. There exists a congruence subgroup* $(\mathfrak{f}, C_{\mathfrak{f}}) \in \mathcal{C}$ (*called the* conductor *of the class*) *such that* $\mathcal{C}$ *consists exactly of all congruence subgroups of the form* $(\mathfrak{m}, C_{\mathfrak{f}} \cap I_{\mathfrak{m}})$ *for all multiples* $\mathfrak{m}$ *of* $\mathfrak{f}$.

Proof. This immediately follows from the proposition by taking for $\mathfrak{f}$ the GCD of all moduli in the class $\mathcal{C}$ (which will, in fact, be the GCD of only a finite number of moduli) and applying the proposition inductively. □

Definition 3.3.11. (1) *We say that* $\mathfrak{f}$ *is the conductor of a congruence subgroup* $(\mathfrak{m}, C)$ *if there exists a congruence subgroup* $C_\mathfrak{f}$ *modulo* $\mathfrak{f}$ *(necessarily equal to* $CP_\mathfrak{f}$*) such that* $(\mathfrak{f}, C_\mathfrak{f})$ *is the conductor of the equivalence class of* $(\mathfrak{m}, C)$.

(2) *A modulus* $\mathfrak{f}$ *is called a* conductor *if there exists a congruence subgroup of conductor equal to* $\mathfrak{f}$.

Proposition 3.3.12. (1) *If a modulus* $\mathfrak{f}$ *is equal to the conductor of* $(\mathfrak{f}, C)$*, then for all congruence subgroups* $D \subset C$ *modulo* $\mathfrak{f}$*, the conductor of* $(\mathfrak{f}, D)$ *is also equal to* $\mathfrak{f}$.

(2) *A modulus* $\mathfrak{f}$ *is a conductor if and only if the conductor of* $(\mathfrak{f}, P_\mathfrak{f})$ *is equal to* $\mathfrak{f}$.

Proof. (1). Assume that $\mathfrak{f}$ is equal to the conductor of $(\mathfrak{f}, C)$, let $D \subset C$, and let $\mathfrak{n}$ be the conductor of $(\mathfrak{f}, D)$, so that $\mathfrak{n} \mid \mathfrak{f}$. By Proposition 3.3.5, we have $I_\mathfrak{f} \cap P_\mathfrak{n} \subset D \subset C$. Thus, $(\mathfrak{n}, CP_\mathfrak{n}) \sim (\mathfrak{f}, C)$, and since $\mathfrak{f}$ is the conductor of $(\mathfrak{f}, C)$ and $\mathfrak{n} \mid \mathfrak{f}$, we must have $\mathfrak{n} = \mathfrak{f}$, proving (1).

(2). If $\mathfrak{f}$ is the conductor of $(\mathfrak{f}, P_\mathfrak{f})$, then $\mathfrak{f}$ is a conductor, while if $\mathfrak{f}$ is a conductor — that is, if $\mathfrak{f}$ is the conductor of $(\mathfrak{f}, C)$ for some congruence subgroup C — then $\mathfrak{f}$ is the conductor of $(\mathfrak{f}, P_\mathfrak{f})$ by (1). □

Corollary 3.3.13. *A modulus* $\mathfrak{f}$ *is the conductor of the equivalence class of* $(\mathfrak{f}, C)$ *if and only if for any* $\mathfrak{n} \mid \mathfrak{f}$, $\mathfrak{n} \neq \mathfrak{f}$, *we have* $h_{\mathfrak{n},CP_\mathfrak{n}} < h_{\mathfrak{f},C}$. *In particular,* $\mathfrak{f}$ *is a conductor if and only if for all* $\mathfrak{n} \mid \mathfrak{f}$, $\mathfrak{n} \neq \mathfrak{f}$, *we have* $h_\mathfrak{n} < h_\mathfrak{f}$.

Proof. This is an immediate consequence of Proposition 3.3.6 and the above proposition. □

3.3.3 Characters of Congruence Subgroups

We now study the notion of *characters* modulo a modulus, or associated to a congruence subgroup.

Definition 3.3.14. (1) *Let* $\mathfrak{m}$ *be a modulus. A* character χ *modulo* $\mathfrak{m}$ *is a group homomorphism from* $I_\mathfrak{m}$ *to* $\mathbb{C}^*$ *such that* $P_\mathfrak{m} \subset \mathrm{Ker}(\chi)$.

(2) *Let* $(\mathfrak{m}, C)$ *be a congruence subgroup. We say that* χ *is a character of* $(\mathfrak{m}, C)$ *if* χ *is a character modulo* $\mathfrak{m}$ *such that* $C \subset \mathrm{Ker}(\chi)$.

Remarks

(1) We can clearly identify characters χ modulo $\mathfrak{m}$ with characters $\overline{\chi}$ of the finite Abelian group $Cl_\mathfrak{m} = I_\mathfrak{m}/P_\mathfrak{m}$. In particular, there are $h_\mathfrak{m}$ such characters. Similarly, we can identify characters χ of the congruence subgroup $(\mathfrak{m}, C)$ with characters $\overline{\chi}$ of the finite Abelian group $I_\mathfrak{m}/C \simeq Cl_\mathfrak{m}/\overline{C}$; hence, there are $h_{\mathfrak{m},C}$ such characters. This is analogous to the possibility of identifying congruence subgroups C with their quotients $\overline{C}$ by

$P_{\mathfrak{m}}$, and for similar reasons it is preferable to give the basic definitions without taking quotients.

(2) Since $h_{\mathfrak{m}} = Cl_{\mathfrak{m}}$ is finite, the values of $\overline{\chi}$, hence of χ, are roots of unity of order dividing $h_{\mathfrak{m}}$.

Definition 3.3.15. *Let χ be a character modulo* $\mathfrak{m}$.

(1) *The* conductor *of χ, denoted $\mathfrak{f}(\chi)$, is the conductor of the congruence subgroup* $(\mathfrak{m}, \mathrm{Ker}(\chi))$.

(2) *The character χ is said to be* primitive *if* $\mathfrak{f}(\chi) = \mathfrak{m}$.

If ψ is a character modulo some modulus $\mathfrak{m}_2$, and if $\mathfrak{m}_1$ is a multiple of $\mathfrak{m}_2$, we have a canonical homomorphism $s_{\mathfrak{m}_1,\mathfrak{m}_2}$ from $I_{\mathfrak{m}_1}$ to $I_{\mathfrak{m}_2}$, and $\chi = \psi \circ s_{\mathfrak{m}_1,\mathfrak{m}_2}$ is a character modulo $\mathfrak{m}_1$ canonically associated to ψ.

Conversely, if χ is a character modulo $\mathfrak{m}_1$, and if there exists a character ψ modulo $\mathfrak{m}_2$ such that $\chi = \psi \circ s_{\mathfrak{m}_1,\mathfrak{m}_2}$, we will say that χ *can be defined modulo* $\mathfrak{m}_2$ (see diagram).

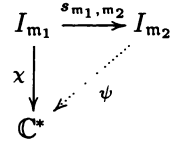

Proposition 3.3.16. *Let χ be a character modulo* $\mathfrak{m}_1$.

(1) *If $\mathfrak{m}_2 \mid \mathfrak{m}_1$, then χ can be defined modulo $\mathfrak{m}_2$ if and only if $I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset \mathrm{Ker}(\chi)$, if and only if there exists a congruence subgroup C_2 modulo $\mathfrak{m}_2$ such that $(\mathfrak{m}_2, C_2) \sim (\mathfrak{m}_1, \mathrm{Ker}(\chi))$.*

(2) *The conductor of χ is equal to $\mathfrak{f}$ if and only if χ can be defined modulo $\mathfrak{f}$ and if for every $\mathfrak{n} \mid \mathfrak{f}$ and different from $\mathfrak{f}$, χ cannot be defined modulo $\mathfrak{n}$.*

(3) *In particular, χ is primitive if and only if for every $\mathfrak{m}_2 \mid \mathfrak{m}_1$ with $\mathfrak{m}_2 \neq \mathfrak{m}_1$, we have $I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \not\subset \mathrm{Ker}(\chi)$.*

Proof. By Proposition 3.3.6, we have the following exact sequence:

$$1 \longrightarrow (I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2})/P_{\mathfrak{m}_1} \longrightarrow Cl_{\mathfrak{m}_1} \longrightarrow Cl_{\mathfrak{m}_2} \longrightarrow 1 \ .$$

Thus, if χ can be defined modulo $\mathfrak{m}_2$, then $\overline{\chi}$ factors through $Cl_{\mathfrak{m}_2}$; hence it is trivial on the kernel of the map from $Cl_{\mathfrak{m}_1}$ to $Cl_{\mathfrak{m}_2}$, that is, on $(I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2})/P_{\mathfrak{m}_1}$. Conversely, if $\overline{\chi}$ is trivial on this kernel, then clearly $\overline{\chi}$ can be lifted to a map from $Cl_{\mathfrak{m}_2}$ to $\mathbb{C}^*$; hence χ can be defined modulo $\mathfrak{m}_2$.

Since $P_{\mathfrak{m}_1} \subset \mathrm{Ker}(\chi)$, we see that χ can be defined modulo $\mathfrak{m}_2$ if and only if $I_{\mathfrak{m}_1} \cap P_{\mathfrak{m}_2} \subset \mathrm{Ker}(\chi)$. By Proposition 3.3.5, this is equivalent to the existence of a congruence subgroup C_2 modulo $\mathfrak{m}_2$ such that $(\mathfrak{m}_2, C_2) \sim (\mathfrak{m}_1, \mathrm{Ker}(\chi))$. Statements (2) and (3) are trivial consequences of (1) and of the definitions. □

Proposition 3.3.17. *Let $(\mathfrak{m}, C)$ be a congruence subgroup, and let $\mathfrak{f}$ be the conductor of $(\mathfrak{m}, C)$. Then*

(1) *we have $C = \bigcap_\chi \mathrm{Ker}(\chi)$, where the intersection is taken over the characters of the congruence subgroup $(\mathfrak{m}, C)$;*
(2) *we have*

$$\mathfrak{f} = \mathrm{lcm}\{\mathfrak{f}(\chi)/\ C \subset \mathrm{Ker}(\chi)\}\ .$$

In other words, the conductor of $(\mathfrak{m}, C)$ is the LCM (or the intersection) of the conductors of the characters of the congruence subgroup $(\mathfrak{m}, C)$.

Proof. (1). By assumption, C is included in the intersection. Conversely, if C was not equal to the intersection, we could find an $\mathfrak{a} \in I_\mathfrak{m}$, $\mathfrak{a} \notin C$, such that $\chi(\mathfrak{a}) = 1$ for all characters χ of the congruence subgroup $(\mathfrak{m}, C)$. But in the finite quotient group $I_\mathfrak{m}/C$, this means that $\overline{\chi}(\overline{\mathfrak{a}}) = 1$ for all characters of the group, hence that $\overline{\mathfrak{a}} = \overline{1}$, so that $\mathfrak{a} \in C$, a contradiction.

(2). If χ is a character of the congruence subgroup $(\mathfrak{m}, C)$, then $C \subset \mathrm{Ker}(\chi)$. By definition of the conductor of a congruence subgroup, we have $I_\mathfrak{m} \cap P_\mathfrak{f} \subset C \subset \mathrm{Ker}(\chi)$. Hence by Proposition 3.3.16, χ can be defined modulo $\mathfrak{f}$, and so $\mathfrak{f}(\chi) \mid \mathfrak{f}$.

Conversely, let $\mathfrak{n}$ be a multiple of all the $\mathfrak{f}(\chi)$ for χ a character of the congruence subgroup $(\mathfrak{m}, C)$. Let χ be such a character. Then by Proposition 3.3.16, since χ can be defined modulo $\mathfrak{f}(\chi)$, we have $I_\mathfrak{m} \cap P_{\mathfrak{f}(\chi)} \subset \mathrm{Ker}(\chi)$. Therefore, since $\mathfrak{f}(\chi) \mid \mathfrak{n}$, we have in particular $I_\mathfrak{m} \cap P_\mathfrak{n} \subset \mathrm{Ker}(\chi)$, so

$$I_\mathfrak{m} \cap P_\mathfrak{n} \subset \bigcap_\chi \mathrm{Ker}(\chi)\ .$$

Thus, by (1) we have $I_\mathfrak{m} \cap P_\mathfrak{n} \subset C$, so by Proposition 3.3.5, there exists a congruence subgroup C' modulo $\mathfrak{n}$ such that $(\mathfrak{n}, C') \sim (\mathfrak{m}, C)$, and hence $\mathfrak{f} \mid \mathfrak{n}$, proving the proposition. □

3.3.4 Conditions on the Conductor and Examples

The following proposition gives a number of *necessary* conditions a conductor must satisfy.

Proposition 3.3.18. *Let $\mathfrak{f}$ be a conductor (in other words, the conductor of some equivalence class of congruence subgroups). Then $\mathfrak{f}$ satisfies the following properties.*

(1) *If $\mathfrak{p} \mid \mathfrak{f}$ and $\mathcal{N}(\mathfrak{p}) = 2$, then $\mathfrak{p}^2 \mid \mathfrak{f}$.*
(2) *If $\mathfrak{f} = \mathfrak{p}^2$ with $\mathcal{N}(\mathfrak{p}) = 2$, then $\mathfrak{p}$ is ramified in $K/\mathbb{Q}$.*
(3) *We cannot have $\mathfrak{f} = \mathfrak{f}_\infty$ with $|\mathfrak{f}_\infty| = 1$ and $\mathfrak{f}_0 = \mathbb{Z}_K$ (in other words $\mathfrak{f}$ cannot be reduced to a single real place).*
(4) *We cannot have $\mathcal{N}(\mathfrak{f}) = 3$.*

Proof. (1). Assume $\mathfrak{p} \mid \mathfrak{f}$, $\mathcal{N}(\mathfrak{p}) = 2$, and $\mathfrak{p}^2 \nmid \mathfrak{f}$. Then $\mathfrak{f}/\mathfrak{p}$ and $\mathfrak{p}$ are coprime ideals, so

$$\phi(\mathfrak{f}) = \phi(\mathfrak{f}/\mathfrak{p})\phi(\mathfrak{p}) = \phi(\mathfrak{f}/\mathfrak{p})(\mathcal{N}(\mathfrak{p}) - 1) = \phi(\mathfrak{f}/\mathfrak{p}) \ .$$

However, $U_{\mathfrak{f}}(K) \subset U_{\mathfrak{f}/\mathfrak{p}}(K)$, so $[U(K) : U_{\mathfrak{f}}(K)] \geq [U(K) : U_{\mathfrak{f}/\mathfrak{p}}(K)]$. Thus, Corollary 3.2.4 implies that $h_{\mathfrak{f}/\mathfrak{p}} \geq h_{\mathfrak{f}}$ (and, since $h_{\mathfrak{f}/\mathfrak{p}} \mid h_{\mathfrak{f}}$, that $h_{\mathfrak{f}/\mathfrak{p}} = h_{\mathfrak{f}}$), so by Corollary 3.3.13 we deduce that $\mathfrak{f}$ is not a conductor.

(2), (3), and (4). First note that $-1 \equiv 1 \pmod{{}^*\mathfrak{m}}$ for a modulus $\mathfrak{m}$ if and only if $\mathfrak{m}_\infty = \varnothing$ and $v_{\mathfrak{p}}(2) \geq v_{\mathfrak{p}}(\mathfrak{m})$ for all $\mathfrak{p} \mid \mathfrak{m}$, hence if and only if $\mathfrak{m} \mid 2\mathbb{Z}_K$. It follows that if $\mathfrak{m} \nmid 2\mathbb{Z}_K$, we have $[U(K) : U_{\mathfrak{m}}(K)] \geq 2$. Thus, if $\mathfrak{m} \nmid 2\mathbb{Z}_K$ and $\phi(\mathfrak{m}) = 2$, then $\phi(\mathfrak{f})/[U(K) : U_{\mathfrak{m}}(K)] = 1$, so $h_{\mathfrak{f}} = h = h_{\mathbb{Z}_K}$, and $\mathfrak{f}$ is not a conductor.

If $\mathfrak{f} = \mathfrak{p}^2$ with $\mathcal{N}(\mathfrak{p}) = 2$, we have $\phi(\mathfrak{f}) = 2$; and if $\mathfrak{p}$ is unramified, then $\mathfrak{f} \nmid 2\mathbb{Z}_K$, and so $\mathfrak{f}$ is not a conductor.

If $\mathfrak{f} = \mathfrak{f}_\infty$ with $|\mathfrak{f}_\infty| = 1$, or if $\mathcal{N}(\mathfrak{f}) = 3$, we have $\phi(\mathfrak{f}) = 2$ and $\mathfrak{f} \nmid 2\mathbb{Z}_K$, so $\mathfrak{f}$ is not a conductor. □

We now specialize to the case where $K = \mathbb{Q}$. Denote by ∞ the unique place at infinity of $\mathbb{Q}$.

Proposition 3.3.19. *The moduli ∞, $3\mathbb{Z}$, $4\mathbb{Z}$, and $m\mathbb{Z}$ and $(m\mathbb{Z})\infty$ for $m \equiv 2 \pmod 4$ are not conductors. All other moduli are conductors.*

Proof. This is an easy consequence of Proposition 3.3.18 and the properties of the ϕ-function, and the details are left to the reader (Exercise 8). □

For the case of imaginary quadratic fields, the result is as follows.

Proposition 3.3.20. *Denote by $\mathfrak{p}_\ell$ (resp., $\mathfrak{p}_\ell$ and $\mathfrak{p}'_\ell$) the prime ideal(s) above ℓ when ℓ is ramified (resp., split) in a quadratic field K. If K is an imaginary quadratic field, all moduli are conductors with the following exceptions, given in completely factored form:*

(1) *If $K = \mathbb{Q}(\sqrt{-3})$, the moduli $\mathfrak{p}_3$, $2\mathbb{Z}_K$, $\mathfrak{p}_7$, $\mathfrak{p}'_7$, $\mathfrak{p}_3^2$, $2\mathfrak{p}_3$;*
(2) *if $K = \mathbb{Q}(\sqrt{-1})$, the moduli $\mathfrak{p}_2^2$, $\mathfrak{p}_2^3$, $\mathfrak{p}_5$, $\mathfrak{p}'_5$, and $\mathfrak{p}_2\mathfrak{n}$, where $\mathfrak{n}$ is not divisible by $\mathfrak{p}_2$;*
(3) *in all other cases, the excluded moduli are exactly those given by Proposition 3.3.18: in other words, $\mathfrak{p}_2^2$ and ${\mathfrak{p}'_2}^2$ if $\mathfrak{p}_2$ and $\mathfrak{p}'_2$ are the unramified ideals of degree 1 above 2 when $d(K) \equiv 1 \pmod 8$, $\mathfrak{p}_2^3$ when $d(K) \equiv 0 \pmod 4$, $\mathfrak{p}_3$ or $\mathfrak{p}_3$ and $\mathfrak{p}'_3$ if $\mathfrak{p}_3$ and $\mathfrak{p}'_3$ are ideals of degree 1 above 3 when $d(K) \not\equiv 2 \pmod 3$, and $\mathfrak{p}_2\mathfrak{n}$, where $\mathfrak{n}$ is not divisible by $\mathfrak{p}_2$, where $\mathfrak{p}_2$ is an ideal of degree 1 above 2, when $d(K) \not\equiv 5 \pmod 8$.*

Proof. Once again the proof is left to the reader (Exercise 9). □

Note that for *real* quadratic fields, or for more general number fields, the situation is more complicated because of the presence of an infinite group of units (see Exercise 10).

If we fix the cardinality of the ray class group $I_{\mathfrak{m}}/C$, the conductor must satisfy more conditions.

Proposition 3.3.21. *Let $(\mathfrak{m}, C)$ be a congruence subgroup, let $\mathfrak{f}$ be its conductor, let $n = h_{\mathfrak{m},C}$, let $\mathfrak{p}$ be a prime ideal dividing $\mathfrak{f}$, and finally let ℓ be the prime number below $\mathfrak{p}$.*

(1) *If $v_{\mathfrak{p}}(\mathfrak{f}) \geq 2$, we necessarily have $\ell \mid n$. In other words, if $\ell \nmid n$, then $v_{\mathfrak{p}}(\mathfrak{f}) = 1$.*
(2) *Conversely, if $v_{\mathfrak{p}}(\mathfrak{f}) = 1$, then $\gcd(n, \mathcal{N}(\mathfrak{p}) - 1) > 1$, or stated otherwise, if $\gcd(n, \mathcal{N}(\mathfrak{p}) - 1) = 1$, and in particular when n is a power of ℓ, then $v_{\mathfrak{p}}(\mathfrak{f}) \geq 2$.*

Proof. Since $h_{\mathfrak{f},CP_{\mathfrak{f}}} = h_{\mathfrak{m},C}$, replacing $(\mathfrak{m}, C)$ by the equivalent congruence subgroup $(\mathfrak{f}, CP_{\mathfrak{f}})$, we may assume that $\mathfrak{f} = \mathfrak{m}$.

For (1), let $\mathfrak{p}$ be such that $v_{\mathfrak{p}}(\mathfrak{m}) \geq 2$, and assume that $\ell \nmid n$. If we set $\mathfrak{n} = \mathfrak{m}/\mathfrak{p}$, it follows in particular that $I_{\mathfrak{n}} = I_{\mathfrak{m}}$. Set $G = CP_{\mathfrak{n}}/C$. We have

$$G \subset CI_{\mathfrak{n}}/C = I_{\mathfrak{m}}/C \simeq Cl_{\mathfrak{m}}/\overline{C} \ ,$$

so $|G| \mid h_{\mathfrak{m},C} = n$.

On the other hand, $P_{\mathfrak{n}} \subset I_{\mathfrak{n}} = I_{\mathfrak{m}}$, so $I_{\mathfrak{m}} \cap CP_{\mathfrak{n}} = CP_{\mathfrak{n}}$. Hence Corollary 3.3.7 tells us that $|G| = h_{\mathfrak{m},C}/h_{\mathfrak{n},CP_{\mathfrak{n}}}$ divides $\phi(\mathfrak{m})/\phi(\mathfrak{n})$. Since $\mathfrak{p}^2 \mid \mathfrak{m}$, we have $\phi(\mathfrak{m})/\phi(\mathfrak{n}) = \mathcal{N}(\mathfrak{p})$, and since we have assumed that $\ell \nmid n$, it follows that $|G|$ divides $\gcd(n, \mathcal{N}(\mathfrak{p})) = 1$.

It follows that $CP_{\mathfrak{n}} = C$, so $P_{\mathfrak{n}} \subset C$. Thus $I_{\mathfrak{m}} \cap P_{\mathfrak{n}} = P_{\mathfrak{n}} \subset C$, so Proposition 3.3.5 shows that the conductor divides $\mathfrak{n} = \mathfrak{m}/\mathfrak{p}$, which is absurd since we have assumed that $\mathfrak{m}$ is the conductor.

For (2), assume that $v_{\mathfrak{p}}(\mathfrak{m}) = 1$ and set $\mathfrak{n} = \mathfrak{m}/\mathfrak{p}$. Then once again by Corollary 3.3.7, we know that $d = h_{\mathfrak{m},C}/h_{\mathfrak{n},CP_{\mathfrak{n}}}$ divides $\phi(\mathfrak{m})/\phi(\mathfrak{n}) = \mathcal{N}(\mathfrak{p}) - 1$ since $\mathfrak{p} \nmid \mathfrak{n}$. On the other hand, since $\mathfrak{m}$ is the conductor, we have $d > 1$, and of course d divides $n = h_{\mathfrak{m},C}$. It follows that $d \mid (n, \mathcal{N}(\mathfrak{p}) - 1)$, so $(n, \mathcal{N}(\mathfrak{p}) - 1) > 1$, as claimed. □

In particular, we deduce from this proposition that if n is a power of a prime ℓ, then for any prime ideal $\mathfrak{p}$ such that $\mathfrak{p} \mid \mathfrak{f}$, we have $v_{\mathfrak{p}}(\mathfrak{f}) = 1$ if $\mathfrak{p}$ is not above ℓ, while $v_{\mathfrak{p}}(\mathfrak{f}) \geq 2$ if $\mathfrak{p}$ is above ℓ (of course, some ideals above ℓ may have $v_{\mathfrak{p}}(\mathfrak{f}) = 0$). In addition, if $\mathfrak{p} \mid \mathfrak{f}$ is above ℓ, we see from Proposition 3.3.21 that $(n, \mathcal{N}(\mathfrak{p}) - 1) = 1$ implies $\ell \mid n$.

The conductor must also satisfy upper bounds. For example, we have the following proposition, which is in fact most easily proved using the "other side" of class field theory; see Corollary 10.1.24.

Proposition 3.3.22. *Keep the notation of Proposition 3.3.21. If $n = \ell$ is prime and $\mathfrak{p}$ is a prime above ℓ dividing $\mathfrak{f}$, then*

$$2 \leq v_{\mathfrak{p}}(\mathfrak{f}) \leq \left\lfloor \frac{\ell e(\mathfrak{p}/\ell)}{\ell - 1} \right\rfloor + 1 \ ,$$

and these bounds are the best possible.

This terminates the description of the "easy" side of class field theory. Although we have used some results of class field theory to *motivate* the definition of equivalence, the definition itself as well as the proofs that we have given are completely self-contained.

In this section on congruence subgroups, we have given complete proofs and details since they are not difficult. In the next section on Abelian extensions and Takagi's theorem, we will omit almost all proofs since they form books by themselves, and we instead refer to [Art-Tat], [Gras], [Has1], [Jan], or [Neu].

3.4 Abelian Extensions: The Other Side of Class Field Theory

We now consider the other — more important — side of class field theory: finite Abelian extensions. The equivalence relation is trivial to define here. We will say that two extensions L/K and L'/K are equivalent (or K-isomorphic) if there exists a K-linear field isomorphism between L and L'; in other words, a field isomorphism from L to L' that leaves K pointwise fixed. If L and L' are K-isomorphic, they are isomorphic as number fields over $\mathbb{Q}$, but the converse is not necessarily true (see Exercise 2 of Chapter 9).

From now on, we let L/K be some (finite) Abelian extension of K of degree n and Abelian Galois group G, and we let $\mathfrak{m}$ be a modulus of K that is assumed always to contain the places of K that ramify in L, that is the prime ideals of K ramified in L/K as well as the real places of K that ramify (see Definition 2.2.4). We will, in fact, need the slightly stronger condition that $\mathfrak{m}$ is a multiple of the *conductor* of L/K (see Definition 3.4.1 below). We are going to define in two completely different ways *two* congruence subgroups attached to $\mathfrak{m}$ (and L, of course, which for the moment is assumed to be fixed). One of the important theorems of class field theory is that these two groups are equal.

3.4.1 The Conductor of an Abelian Extension

We first define the *conductor* of an Abelian extension L/K. For this, it is useful, although not strictly necessary, to use some undefined p-adic terminology

(see, for example, Definition 4.2.5). Let $\mathfrak{p}$ be a prime ideal of K and let $\mathfrak{P}$ be some prime ideal of L above $\mathfrak{p}$. We will say that an element $x \in K_{\mathfrak{p}}^*$ is a *local norm* modulo $\mathfrak{p}$ if there exists $y \in L_{\mathfrak{P}}^*$ such that $x = \mathcal{N}_{L_{\mathfrak{P}}/K_{\mathfrak{p}}}(y)$ (this does not depend on the chosen $\mathfrak{P}$ above $\mathfrak{p}$). If $x \in K^*$, this is equivalent to the requirement that for all $n \geq 0$ there exists $y_n \in L^*$ such that $x \equiv \mathcal{N}_{L/K}(y_n) \pmod{{}^*\mathfrak{p}^n}$, or even $x/\mathcal{N}_{L/K}(y_n) \equiv 1 \pmod{{}^*\mathfrak{p}^n}$.

We define a nonnegative integer $k_{\mathfrak{p}}$ to be the smallest $k \geq 0$ such that any element $x \equiv 1 \pmod{{}^*\mathfrak{p}^k}$ coprime to $\mathfrak{p}$ (this is, of course, necessary only for $k = 0$) is a local norm modulo $\mathfrak{p}$. It can be shown that $k_{\mathfrak{p}}$ exists and that $k_{\mathfrak{p}} = 0$ if and only if $\mathfrak{p}$ is unramified in L/K.

Definition 3.4.1. *With the above notation, let*

$$\mathfrak{f}_0(L/K) = \prod_{\mathfrak{p}} \mathfrak{p}^{k_{\mathfrak{p}}} \ ,$$

and let $\mathfrak{f}_\infty(L/K)$ be the set of real places of K ramified in L. We define the conductor *of the Abelian extension L/K to be the modulus $\mathfrak{f}(L/K) = \mathfrak{f}_0(L/K)\mathfrak{f}_\infty(L/K)$.*

The definition of $\mathfrak{f}_0(L/K)$ involves only a finite number of prime ideals since $k_{\mathfrak{p}} \neq 0$ only for the ramified primes. Thus, the prime ideals that divide the conductor are the ramified primes, and we will see in Theorem 3.5.10 that the finite part of the conductor divides the relative discriminant ideal $\mathfrak{d}(L/K)$.

3.4.2 The Frobenius Homomorphism

In this subsection, we recall some basic facts of algebraic number theory (see Section 10.1.2 and [Marc]).

Let L/K be a normal extension of degree n with Galois group $G = \mathrm{Gal}(L/K)$ (for the moment, not necessarily Abelian), and let $\mathfrak{p}$ be an ideal of K, possibly ramified. Then $\mathfrak{p}$ decomposes in L as a product of prime ideals $\mathfrak{p}\mathbb{Z}_L = \prod_{1 \leq i \leq g} \mathfrak{P}_i^e$. Since the extension is normal, the $\mathfrak{P}_i$ are permuted transitively by the Galois group G and hence all have the same ramification index $e = e(\mathfrak{P}_i/\mathfrak{p})$ and residual degree $f = f(\mathfrak{P}_i/\mathfrak{p})$. Thus, $efg = n$ (Proposition 10.1.3).

Let $\mathfrak{P}$ be one of the ideals $\mathfrak{P}_i$ above $\mathfrak{p}$. Recall that the *decomposition group* $D(\mathfrak{P}/\mathfrak{p})$ of $\mathfrak{P}$ is the group of elements $\sigma \in G$ fixing $\mathfrak{P}$ globally, in other words such that $\sigma(\mathfrak{P}) = \mathfrak{P}$ (see Definition 10.1.4). We have $|D(\mathfrak{P}/\mathfrak{p})| = ef$ and the fixed field L^D of L by $D(\mathfrak{P}/\mathfrak{p})$ is an extension of K of degree g.

Recall also that the *inertia group* $I(\mathfrak{P}/\mathfrak{p})$ of $\mathfrak{P}$ is the group of elements $\sigma \in G$ such that $\sigma(x) \equiv x \pmod{\mathfrak{P}}$ for all $x \in \mathbb{Z}_L$. We have $I(\mathfrak{P}/\mathfrak{p}) \subset D(\mathfrak{P}/\mathfrak{p})$, and $D(\mathfrak{P}/\mathfrak{p})/I(\mathfrak{P}/\mathfrak{p})$ is canonically isomorphic to $\mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$. This has a number of important consequences. First, we have $|I(\mathfrak{P}/\mathfrak{p})| = e$, and the fixed field L^I of L by $I(\mathfrak{P}/\mathfrak{p})$ is an extension of K of degree fg, and

it is the largest subextension of L/K in which $\mathfrak{p}$ is unramified. Since G is not necessarily Abelian, the extension L^D/K is not necessary normal. On the other hand, the extensions L^I/L^D and L/L^I are normal. When G is Abelian, we can say in colorful terms that $\mathfrak{p}$ acquires its splitting behavior in the extension L^D/K of degree g, its residual degrees in the extension L^I/L^D of degree f, and its ramification properties in the extension L/L^I of degree e, all these extensions being Abelian.

Finally, recall the existence (and uniqueness up to conjugation by an element of $I(\mathfrak{P}/\mathfrak{p})$) of a *Frobenius homomorphism* $\sigma_{\mathfrak{P}} \in D(\mathfrak{P}/\mathfrak{p})$ such that for all $x \in \mathbb{Z}_L$ we have $\sigma_{\mathfrak{P}}(x) \equiv x^{\mathcal{N}(\mathfrak{p})} \pmod{\mathfrak{P}}$ (see Proposition 10.1.5).

If $\mathfrak{P}'$ is some other ideal of $\mathbb{Z}_L$ above $\mathfrak{p}$, by transitivity of the Galois action we have $\mathfrak{P}' = \tau(\mathfrak{P})$ for some $\tau \in G$, and we have $D(\mathfrak{P}'/\mathfrak{p}) = \tau D(\mathfrak{P}/\mathfrak{p})\tau^{-1}$ and $I(\mathfrak{P}'/\mathfrak{p}) = \tau I(\mathfrak{P}/\mathfrak{p})\tau^{-1}$ (see Section 10.1.2). In particular, if L/K is Abelian, then $D(\mathfrak{P}/\mathfrak{p})$ and $I(\mathfrak{P}/\mathfrak{p})$ are independent of the choice of $\mathfrak{P}$ above $\mathfrak{p}$.

Let us come back to the situation where L/K is an Abelian extension, and now assume that $\mathfrak{p}$ is unramified in L/K, hence that $I(\mathfrak{P}/\mathfrak{p}) = \{1_G\}$ for all $\mathfrak{P}$ above p. The above discussion shows that there exists a canonical element $\sigma_{\mathfrak{P}} \in G$, called the *Frobenius homomorphism* and characterized by the congruence

$$\sigma_{\mathfrak{P}}(x) \equiv x^{\mathcal{N}(\mathfrak{p})} \pmod{\mathfrak{P}} \quad \text{for all} \quad x \in \mathbb{Z}_L \ .$$

This homomorphism is of order exactly equal to the residual degree $f(\mathfrak{P}/\mathfrak{p})$. Since our group G is Abelian, $\sigma_{\mathfrak{P}}$ only depends on $\mathfrak{p}$ and hence will be denoted $\sigma_{\mathfrak{p}}$. It is characterized by the congruence

$$\sigma_{\mathfrak{p}}(x) \equiv x^{\mathcal{N}(\mathfrak{p})} \pmod{\mathfrak{p}\mathbb{Z}_L} \quad \text{for all} \quad x \in \mathbb{Z}_L \ .$$

3.4.3 The Artin Map and the Artin Group $A_{\mathfrak{m}}(L/K)$

Definition 3.4.2. *Let L/K be an Abelian extension and $\mathfrak{m}$ a modulus of K. We say that $\mathfrak{m}$ is a* suitable *modulus for the extension L/K if $\mathfrak{m}$ is a multiple of the conductor of L/K.*

Let $\mathfrak{m}$ be a suitable modulus of K, so that in particular $\mathfrak{m}$ is divisible by all ramified places. We will now define a group homomorphism from $I_{\mathfrak{m}}$, the group of fractional ideals coprime to $\mathfrak{m}$, into G, the Galois group of L/K. If $\mathfrak{a} \in I_{\mathfrak{m}}$, we can write

$$\mathfrak{a} = \prod_{\mathfrak{p}|\mathfrak{a}} \mathfrak{p}^{v_{\mathfrak{p}}(\mathfrak{a})} \ ,$$

where the prime ideals $\mathfrak{p}$ do not divide $\mathfrak{m}$ and in particular are *unramified* in L/K. We set

$$\mathrm{Art}_{L/K}(\mathfrak{a}) = \prod_{\mathfrak{p}|\mathfrak{a}} \sigma_{\mathfrak{p}}^{v_{\mathfrak{p}}(\mathfrak{a})} \ ,$$

where the product is of course taken in the group G. It is clear that this map is well-defined and is a group homomorphism. This map is called the *Artin reciprocity map*, and $\mathrm{Art}_{L/K}(\mathfrak{a})$ is often denoted by $\left(\frac{L/K}{\mathfrak{a}}\right)$ (and called the Artin symbol) since it generalizes the Jacobi symbol $\left(\frac{D}{n}\right)$ (see Exercise 11). Note that, strictly speaking, we should write $\mathrm{Art}_{L/K,\mathfrak{m}}(\mathfrak{a})$ to indicate that the Artin map is defined on $I_{\mathfrak{m}}$, but since clearly $\mathrm{Art}_{L/K,\mathfrak{m}}(\mathfrak{a})$ does not depend on $\mathfrak{m}$ as long as it is a multiple of the conductor and coprime to $\mathfrak{a}$, we will omit $\mathfrak{m}$. In fact, we shall see in Proposition 3.5.6 that we can even omit the explicit mention of L/K if desired.

The first important theorem of class field theory (called *Artin's reciprocity law* since it implies more or less easily all the usual reciprocity laws) is the following.

Theorem 3.4.3 (Artin reciprocity). (1) *The Artin reciprocity map is a* surjective *group homomorphism from* $I_{\mathfrak{m}}$ *to* $G = \mathrm{Gal}(L/K)$.

(2) *The kernel of the Artin reciprocity map is a congruence subgroup modulo* $\mathfrak{m}$*; in other words, it contains* $P_{\mathfrak{m}}$.

Thanks to this theorem, we can also view the Artin reciprocity map as a surjective map from $Cl_{\mathfrak{m}} = I_{\mathfrak{m}}/P_{\mathfrak{m}}$ to G.

We will denote by $A_{\mathfrak{m}}(L/K)$ the kernel of the Artin reciprocity map, which, by this theorem, is a congruence subgroup modulo $\mathfrak{m}$, and call it the *Artin group* attached to the modulus $\mathfrak{m}$ and the extension L/K.

3.4.4 The Norm Group (or Takagi Group) $T_{\mathfrak{m}}(L/K)$

We now define another congruence subgroup attached to $\mathfrak{m}$ as follows. Denote by $I_{\mathfrak{m},L}$ the group of fractional ideals of L that are coprime to $\mathfrak{m}$, more precisely to the extended ideal $\mathfrak{m}_0\mathbb{Z}_L$ (in other words, $I_{\mathfrak{m},L} = I_{\mathfrak{m}\mathbb{Z}_L}(L)$). The relative norm $\mathcal{N}_{L/K}$ of an ideal belonging to $I_{\mathfrak{m},L}$ is clearly an ideal of K coprime to $\mathfrak{m}$; hence it belongs to $I_{\mathfrak{m}}$. Thus, the image group $\mathcal{N}_{L/K}(I_{\mathfrak{m},L})$ is a subgroup of $I_{\mathfrak{m}}$. However, it is not necessarily a congruence subgroup since it need not contain $P_{\mathfrak{m}}$ (see Exercise 12). Thus, we will set

$$T_{\mathfrak{m}}(L/K) = P_{\mathfrak{m}}\,\mathcal{N}_{L/K}(I_{\mathfrak{m},L})$$

and this is now a congruence subgroup modulo $\mathfrak{m}$ that we will call the *norm group* (or *Takagi group*) for the modulus $\mathfrak{m}$ and the extension L/K.

The following theorem gives an easy way to compute the norm group.

Theorem 3.4.4. *Let* $\mathfrak{p}$ *be a prime ideal of* K *not dividing* $\mathfrak{m}$.

(1) *If* f *is the least positive integer such that* $\mathfrak{p}^f \in T_{\mathfrak{m}}(L/K)$*, then* $f = f(\mathfrak{P}/\mathfrak{p})$ *is the residual degree of* $\mathfrak{P}$*, and hence* $\mathfrak{p}$ *splits into* $g = n/f$ *prime ideals of degree* f *in* L/K.

(2) *The norm group* $T_{\mathfrak{m}}(L/K)$ *is generated by the ideals* $\mathfrak{p}^f = \mathcal{N}_{L/K}(\mathfrak{P})$ (*with* $f = f(\mathfrak{P}/\mathfrak{p})$), *and in fact simply by the ideals of degree* f *equal to* 1.

The second very important, and difficult, theorem of class field theory is the following theorem.

Theorem 3.4.5. *Let* $\mathfrak{m}$ *be a suitable modulus for* L/K (*see Definition 3.4.2*). *Then*

$$A_{\mathfrak{m}}(L/K) = T_{\mathfrak{m}}(L/K) \ .$$

The main point of this theorem (apart from its intrinsic beauty and interest, and its consequences) is that the Artin group $A_{\mathfrak{m}}(L/K)$ is not easy to compute directly, while the norm group $T_{\mathfrak{m}}(L/K)$ is easy to compute thanks to Theorem 3.4.4 (see Algorithm 4.4.3).

Another important result is the following.

Theorem 3.4.6. (1) *If* $\mathfrak{m}$ *and* $\mathfrak{n}$ *are two suitable moduli for* L/K, *the congruence subgroups* $(\mathfrak{m}, A_{\mathfrak{m}}(L/K))$ *and* $(\mathfrak{n}, A_{\mathfrak{n}}(L/K))$ *are equivalent in the sense of Definition 3.3.3.*

(2) *The conductor of the equivalence class of the family of congruence subgroups* $(\mathfrak{m}, A_{\mathfrak{m}}(L/K))$ *is equal to the conductor* $\mathfrak{f}(L/K)$ *of the Abelian extension.*

3.5 Putting Both Sides Together: The Takagi Existence Theorem

3.5.1 The Takagi Existence Theorem

We now state the most important — and most difficult — theorem of classical global class field theory, due to Takagi.

Theorem 3.5.1. (1) *The map that sends an equivalence class of Abelian extensions* L/K *to the equivalence class of the congruence subgroup* $(\mathfrak{m}, A_{\mathfrak{m}}(L/K))$ *for any suitable* $\mathfrak{m}$ *for the extension* L/K *is a bijection* (*by Theorem 3.4.6, this equivalence class is independent of* $\mathfrak{m}$).

(2) *More precisely, if* $(\mathfrak{m}, A_{\mathfrak{m}}(L/K))$ *is equivalent to* $(\mathfrak{m}', A_{\mathfrak{m}'}(L'/K))$ *in the sense of Definition 3.3.3, then the number fields* L *and* L' *are* K*-isomorphic.*

(3) *Conversely, if* $(\mathfrak{m}, C)$ *is any congruence subgroup, there exists an Abelian extension* L/K, *unique up to* K*-isomorphism, such that* $\mathfrak{m}$ *is a suitable modulus for* L/K *and* $C = A_{\mathfrak{m}}(L/K) = T_{\mathfrak{m}}(L/K)$.

The proof that the map is injective is not very difficult. However, the proof of the surjectivity is an *existence* proof and is very hard, like almost all such existence proofs in mathematics. In fact, we will see that this phenomenon

is also reflected in algorithmic practice. The difficulty with the proof lies mainly in the very few tools that we have available to construct Abelian extensions. The known proofs all rely on the method of Kummer extensions (see Chapter 10), which is elementary but heavy to use, and we will do the same in algorithmic practice in Chapter 5.

Thus, given a modulus $\mathfrak{m}$ and a congruence subgroup C modulo $\mathfrak{m}$, we know thanks to Takagi's existence theorem that there exists an Abelian extension L/K corresponding to $(\mathfrak{m}, C)$ under the Takagi map. This extension L/K has the following additional properties (and is uniquely characterized by the first two).

Proposition 3.5.2. *With the above notation, we have the following.*

(1) *The Artin reciprocity map induces a canonical isomorphism from* $Cl_{\mathfrak{m}}/\overline{C}$ *to* $\mathrm{Gal}(L/K)$*; so in particular,* $n = [L : K] = \left|Cl_{\mathfrak{m}}/\overline{C}\right| = h_{\mathfrak{m},C}$.
(2) $C = P_{\mathfrak{m}} N_{L/K}(I_{\mathfrak{m},L})$.
(3) *The conductor* $\mathfrak{f} = \mathfrak{f}(L/K)$ *of the Abelian extension is equal to the conductor of the corresponding congruence subgroup (this is Theorem 3.4.6).*
(4) *The places of* K *that ramify in* L *are exactly the divisors of* $\mathfrak{f}$.

The splitting behavior in L/K of the prime ideals of K is completely described by the following theorem, which generalizes Theorem 3.4.4 (1).

Theorem 3.5.3. *Let* L/K *be an Abelian extension of degree* n *corresponding to a congruence subgroup* $(\mathfrak{m}, C)$ *under the Takagi map (with* $\mathfrak{m}$ *a multiple of the conductor of* $(\mathfrak{m}, C)$ *but not necessarily equal to it), and let* $\mathfrak{p}$ *be a prime ideal of* K. *Let* $\mathfrak{n} = \mathfrak{m}\mathfrak{p}^{-v_{\mathfrak{p}}(\mathfrak{m})}$ *be the prime to* $\mathfrak{p}$ *part of the modulus* $\mathfrak{m}$. *If we let* $\mathfrak{p}\mathbb{Z}_L = \prod_{1 \le i \le g} \mathfrak{P}_i^e$ *be the prime ideal decomposition of* $\mathfrak{p}$ *in the extension* L/K, *we have*

$$e = e(\mathfrak{P}_i/\mathfrak{p}) = \frac{n}{|I_{\mathfrak{n}}/CP_{\mathfrak{n}}|} = \left|\frac{I_{\mathfrak{m}} \cap CP_{\mathfrak{n}}}{C}\right| = \left|\frac{I_{\mathfrak{m}} \cap P_{\mathfrak{n}}}{C \cap P_{\mathfrak{n}}}\right| ,$$

$f = f(\mathfrak{P}_i/\mathfrak{p})$ *is the order of the class of* $\mathfrak{p}$ *in* $I_{\mathfrak{n}}/CP_{\mathfrak{n}}$ *(equivalently, it is the least positive integer* f *such that* $\mathfrak{p}^f \in CP_{\mathfrak{n}}$*), hence* $g = n/ef$ *is equal to the index of the cyclic subgroup generated by the class of* $\mathfrak{p}$ *in the group* $I_{\mathfrak{n}}/CP_{\mathfrak{n}}$.

In particular, if $\mathfrak{p}$ *is unramified in* L/K, *then the common residual degree* f *is the smallest positive integer such that* $\mathfrak{p}^f \in C$, *and* $g = n/f$.

Definition 3.5.4. *Let* $(\mathfrak{m}, C)$ *be a congruence subgroup modulo* $\mathfrak{m}$. *The field extension (or more precisely the equivalence class of field extensions)* L/K *corresponding to* $(\mathfrak{m}, C)$ *by Takagi's theorem is called the* ray class field *for* $(\mathfrak{m}, C)$. *In particular, we denote by* $K(\mathfrak{m})$ *the ray class field for* $(\mathfrak{m}, P_{\mathfrak{m}})$ *and call* $K(1) = K(\mathbb{Z}_K)$ *the* Hilbert class field *of* K.

As mentioned in the prologue, Proposition 3.5.2 shows in particular that the Hilbert class field is the maximal unramified Abelian extension of K and

that $\mathrm{Gal}(K(1)/K) \simeq Cl(K)$. In addition, the ray class field for $(\mathfrak{m}, C)$ is clearly equal to $K(\mathfrak{m})^{\mathrm{Art}(C)}$.

Another easy result we will need is the behavior of class fields under extensions.

Proposition 3.5.5. *As above, let L/K be the Abelian extension of K corresponding to the congruence subgroup $(\mathfrak{m}, C)$. Let K' be any (finite) extension of K. Then LK'/K' is an Abelian extension of K' corresponding to the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K'}, \mathcal{N}_{K'/K}^{-1}(C))$.*

Note that if $\mathfrak{f}$ is the conductor of L/K, then $\mathfrak{f}\mathbb{Z}_{K'}$ is usually *not* equal to the conductor of LK'/K' but is only a multiple of it.

Finally, the following proposition gives the behavior of the Artin map under restriction.

Proposition 3.5.6. *Let N/K be an Abelian extension and let L/K be a subextension of N/K.*

(1) *If $\mathfrak{m}$ is a suitable modulus for the extension N/K, then $\mathfrak{m}$ is a suitable modulus for L/K and the restriction of $\mathrm{Art}_{N/K}$ to the ideals of L coprime to $\mathfrak{m}$ is equal to $\mathrm{Art}_{L/K}$.*
(2) *If $\mathfrak{m}$ is a suitable modulus for the extension N/L, then $\mathcal{N}_{L/K}(\mathfrak{m})$ is a suitable modulus for the extension N/K.*

Thus it is reasonable to drop completely the index L/K from the notation $\mathrm{Art}_{L/K}$. We will usually do this, except when we really want to insist on the specific extension considered.

3.5.2 Signatures, Characters, and Discriminants

This section is taken almost verbatim from joint work of the author with F. Diaz y Diaz and M. Olivier (see [Co-Di-Ol2]).

In this section, we let $(\mathfrak{m}, C)$ be a congruence subgroup, and let L/K be the Abelian extension corresponding to the equivalence class of $(\mathfrak{m}, C)$ by class field theory (well-defined up to K-isomorphism). We do not necessarily assume that $\mathfrak{m}$ is the conductor. We want to compute the signature (R_1, R_2) of L, the relative discriminant ideal $\mathfrak{d}(L/K)$, as well as the absolute discriminant.

As before, we denote by $h_{\mathfrak{m},C}$ the cardinality of the quotient group $I_{\mathfrak{m}}/C \simeq Cl_{\mathfrak{m}}/\overline{C}$. For simplicity, if $\mathfrak{n} \mid \mathfrak{m}$, we write $h_{\mathfrak{n},C}$ instead of $h_{\mathfrak{n},CP_{\mathfrak{n}}}$. Note that by the approximation theorem the natural map $s_{\mathfrak{m},\mathfrak{n}}$ from $Cl_{\mathfrak{m}}$ to $Cl_{\mathfrak{n}}$ is surjective and $s_{\mathfrak{m},\mathfrak{n}}(\overline{C}) = \overline{CP_{\mathfrak{n}}}$.

A reformulation of Corollary 3.3.13 is as follows.

Proposition 3.5.7. *A modulus $\mathfrak{m}$ is the conductor of L/K if and only if for all places $\mathfrak{p} \mid \mathfrak{m}$ (including the places at infinity) we have $h_{\mathfrak{m}/\mathfrak{p},C} < h_{\mathfrak{m},C}$.*

Proof. Indeed, by Corollary 3.3.13 the condition is necessary; but conversely, if this condition is satisfied and if $\mathfrak{n} \mid \mathfrak{m}$, $\mathfrak{n} \neq \mathfrak{m}$, then if $\mathfrak{p} \mid \mathfrak{m}/\mathfrak{n}$, we have $h_{\mathfrak{n},C} \leq h_{\mathfrak{m}/\mathfrak{p},C} < h_{\mathfrak{m},C}$, so we conclude again by Corollary 3.3.13. □

The signature of L is given by the following proposition.

Proposition 3.5.8. *Let (R_1, R_2) be the signature of L, so that $R_1 + 2R_2 = [L : \mathbb{Q}] = [K : \mathbb{Q}] \cdot h_{\mathfrak{m},C}$. Write m_∞ for $|\mathfrak{m}_\infty|$. We have*

$$R_1 = h_{\mathfrak{m},C}\left(r_1 - m_\infty + \sum_{v \in \mathfrak{m}_\infty} \delta(h_{\mathfrak{m},C} - h_{\mathfrak{m}/v,C})\right) ,$$

where $\delta(x) = 1$ if $x = 0$ and $\delta(x) = 0$ otherwise.

In particular, if $\mathfrak{m}$ is the conductor of L/K and $n = [L : K] = h_{\mathfrak{m},C}$, we have

$$R_1 = (r_1 - m_\infty)n \quad \text{and} \quad R_2 = (r_2 + m_\infty/2)n \ .$$

In particular, if m_∞ is odd, then $n = h_{\mathfrak{m},C}$ is even.

Proof. Since L/K is normal, R_1 is equal to $[L : K] = h_{\mathfrak{m},C}$ times the number of real places of K unramified in L. By definition of the ray class group, the $r_1 - m_\infty$ real places not in the modulus $\mathfrak{m}$ are unramified. On the other hand, let $v \in \mathfrak{m}_\infty$. If $h_{\mathfrak{m}/v,C} = h_{\mathfrak{m},C}$, then v does not divide the conductor of L, hence v is unramified in L. On the contrary, if $h_{\mathfrak{m}/v,C} < h_{\mathfrak{m},C}$, then v divides the conductor of L, so v is ramified in L. This gives the first formula of the proposition. The second follows immediately. □

Using the theory of characters of congruence subgroups developed in Section 3.3.3, we now introduce the notion of character associated to an Abelian extension.

Definition 3.5.9. *Let L/K be an Abelian extension of conductor $(\mathfrak{m}, C)$ (so that $\mathfrak{m}$ is the conductor of the extension and C is the corresponding norm group). A character χ of the extension L/K is a character of the congruence subgroup $(\mathfrak{m}, C)$ in the sense of Definition 3.3.14.*

Remark. By class field theory (Proposition 3.5.2) the Galois group $\mathrm{Gal}(L/K)$ is canonically isomorphic to $Cl_{\mathfrak{m}}/\overline{C}$, so we can also consider a character of the extension L/K as being a character of its Galois group. The set of characters of an Abelian extension L/K forms a group of cardinality $n = [L : K]$, isomorphic to $I_{\mathfrak{m}}/C \simeq Cl_{\mathfrak{m}}/\overline{C} \simeq \mathrm{Gal}(L/K)$.

The following result, due to Hasse, is essential for computing discriminants.

Theorem 3.5.10. *Let L/K be an Abelian extension, and denote by $\widehat{G}$ the group of characters of L/K in the sense of Definition 3.5.9.*

(1) *The conductor of L/K is given by $\mathfrak{f}(L/K) = \mathrm{lcm}_{\chi\in\widehat{G}}(\mathfrak{f}(\chi))$, where $f(\chi)$ is the conductor of the character χ (see Definition 3.3.15).*
(2) *The discriminant ideal is given by $\mathfrak{d}(L/K) = \prod_{\chi\in\widehat{G}} \mathfrak{f}(\chi)_0$, where as usual $\mathfrak{f}(\chi)_0$ denotes the finite part of the modulus $\mathfrak{f}(\chi)$.*
(3) *We have $\mathfrak{f}(L/K) \mid \mathfrak{d}(L/K)$, and both ideals are divisible by exactly the same prime ideals: the prime ideals of K ramified in L/K.*

Note that (1) is a reformulation of Proposition 3.3.17, once we know that $\mathfrak{f}(L/K)$ is the conductor of the associated congruence subgroup.

We now give a formula for the relative discriminant ideal $\mathfrak{d}(L/K)$ and hence for the absolute discriminant $d(L)$ of L (see [Co-Di-Ol2]).

Theorem 3.5.11. *Let $(\mathfrak{m}, C)$ be a congruence subgroup, and let L/K be the Abelian extension associated to $(\mathfrak{m}, C)$ by class field theory (defined up to K-isomorphism). Set $n = [L : K] = h_{\mathfrak{m},C}$.*

(1) *The relative discriminant ideal $\mathfrak{d}(L/K)$ is given by $\mathfrak{d}(L/K) = \prod_{\mathfrak{p}|\mathfrak{m}} \mathfrak{p}^{a_\mathfrak{p}}$ with*

$$a_\mathfrak{p} = v_\mathfrak{p}(\mathfrak{m})h_{\mathfrak{m},C} - \sum_{1\le k\le v_\mathfrak{p}(\mathfrak{m})} h_{\mathfrak{m}/\mathfrak{p}^k,C} \ .$$

(2) *Let $\mathfrak{f} = \mathfrak{f}_0\mathfrak{f}_\infty$ be the conductor of the congruence subgroup $(\mathfrak{m}, C)$ (or of L/K), and set $f_\infty = |\mathfrak{f}_\infty|$. The absolute discriminant of L is given by*

$$d(L) = (-1)^{f_\infty n/2} d(K)^n \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \ .$$

Proof. (1). Theorem 3.5.10 tells us that $\mathfrak{d}(L/K) = \prod_{\chi\in\widehat{G}} \mathfrak{f}(\chi)_0$. Set $D(L/K) = \prod_{\chi\in\widehat{G}} \mathfrak{f}(\chi)$, so that $\mathfrak{d}(L/K)$ is the finite part of $D(L/K)$. Note that this can also be taken as the definition of an extended discriminant ideal if desired. Since it is just as simple, we will in fact compute a formula for $D(L/K)$.

For each $\mathfrak{n} \mid \mathfrak{m}$, denote by $a(\mathfrak{n})$ the number of characters of the congruence subgroup $(\mathfrak{m}, C)$ of conductor exactly equal to $\mathfrak{n}$. Since the total number of characters is equal to the order of the group, we have the equation

$$\sum_{\mathfrak{n}|\mathfrak{m}} a(\mathfrak{n}) = \left|Cl_\mathfrak{m}/\overline{C}\right| = h_{\mathfrak{m},C} \ .$$

By Möbius inversion, it follows that

$$a(\mathfrak{n}) = \sum_{\mathfrak{q}|\mathfrak{n}} \mu(\mathfrak{n}/\mathfrak{q})h_{\mathfrak{q},C} \ ,$$

where $\mu(\mathfrak{n})$ is defined as in the case of ordinary integers (this is valid since a modulus can be written as a product of finite or infinite primes in essentially only one way).

Thus, we have

$$D(L/K) = \prod_{\mathfrak{n}|\mathfrak{m}} \prod_{\mathfrak{f}(\chi)=\mathfrak{n}} \mathfrak{f}(\chi) = \prod_{\mathfrak{n}|\mathfrak{m}} \mathfrak{n}^{a(\mathfrak{n})} = \prod_{\mathfrak{n}|\mathfrak{m}} \mathfrak{n}^{\sum_{\mathfrak{q}|\mathfrak{n}} \mu(\mathfrak{n}/\mathfrak{q}) h_{\mathfrak{q},C}}$$

$$= \prod_{\mathfrak{q}|\mathfrak{m}} \left(\prod_{\mathfrak{c}|(\mathfrak{m}/\mathfrak{q})} (\mathfrak{c}\mathfrak{q})^{\mu(\mathfrak{c})} \right)^{h_{\mathfrak{q},C}} = \prod_{\mathfrak{q}|\mathfrak{m}} (p_1(\mathfrak{q}) p_2(\mathfrak{q}))^{h_{\mathfrak{q},C}} ,$$

where

$$p_1(\mathfrak{q}) = \prod_{\mathfrak{c}|(\mathfrak{m}/\mathfrak{q})} \mathfrak{c}^{\mu(\mathfrak{c})} \quad \text{and} \quad p_2(\mathfrak{q}) = \prod_{\mathfrak{c}|(\mathfrak{m}/\mathfrak{q})} \mathfrak{q}^{\mu(\mathfrak{c})} .$$

The product $p_2(\mathfrak{q})$ is trivial to compute: we have

$$p_2(\mathfrak{q}) = \mathfrak{q}^{\sum_{\mathfrak{c}|(\mathfrak{m}/\mathfrak{q})} \mu(\mathfrak{c})} ,$$

and by definition of the μ-function, this exponent is equal to zero unless $\mathfrak{m}/\mathfrak{q} = \mathbb{Z}_K$. Hence $p_2(\mathfrak{q}) = \mathbb{Z}_K$ if $\mathfrak{q} \neq \mathfrak{m}$, and $p_2(\mathfrak{m}) = \mathfrak{m}$.

The product $p_1(\mathfrak{q})$ is computed as follows. Set $L(\mathfrak{c}) = \mathfrak{p}$ if $\mathfrak{c} = \mathfrak{p}^k$ is a nontrivial prime power (including infinite primes, in which case $k = 1$), and $L(\mathfrak{c}) = \mathbb{Z}_K$ otherwise. The existence and uniqueness of the decomposition of $\mathfrak{n}$ into prime powers imply the equality $\prod_{\mathfrak{c}|\mathfrak{n}} L(\mathfrak{c}) = \mathfrak{n}$. By multiplicative Möbius inversion, this gives

$$L(\mathfrak{n}) = \prod_{\mathfrak{c}|\mathfrak{n}} (\mathfrak{n}/\mathfrak{c})^{\mu(\mathfrak{c})} = \prod_{\mathfrak{c}|\mathfrak{n}} \mathfrak{n}^{\mu(\mathfrak{c})} / \prod_{\mathfrak{c}|\mathfrak{n}} \mathfrak{c}^{\mu(\mathfrak{c})} .$$

By definition of μ the numerator is equal to $\mathbb{Z}_K$; hence we obtain the formula

$$\prod_{\mathfrak{c}|\mathfrak{n}} \mathfrak{c}^{\mu(\mathfrak{c})} = L(\mathfrak{n})^{-1} .$$

The reader will certainly have recognized that the function $L(\mathfrak{n})$ is the ideal-theoretic analog of the function $e^{\Lambda(n)}$ of elementary prime number theory.

Using this result in our above formulas, we obtain $p_1(\mathfrak{q}) = L(\mathfrak{m}/\mathfrak{q})^{-1}$; hence,

$$D(L/K) = \mathfrak{m}^{h_{\mathfrak{m},C}} \prod_{\mathfrak{q}|\mathfrak{m}} L(\mathfrak{m}/\mathfrak{q})^{-h_{\mathfrak{q},C}} = \mathfrak{m}^{h_{\mathfrak{m},C}} \prod_{\mathfrak{p}^k|\mathfrak{m}} \mathfrak{p}^{-h_{\mathfrak{m}/\mathfrak{p}^k,C}}$$

$$= \prod_{\mathfrak{p}|\mathfrak{m}} \mathfrak{p}^{v_{\mathfrak{p}}(\mathfrak{m}) h_{\mathfrak{m},C} - \sum_{1 \leq k \leq v_{\mathfrak{p}}(\mathfrak{m})} h_{\mathfrak{m}/\mathfrak{p}^k,C}} ,$$

and (1) follows by taking the finite part.

Note that the infinite part of $D(L/K)$ is equal to

$$\prod_{v \in \mathfrak{m}_\infty} v^{h_{\mathfrak{m},C} - h_{\mathfrak{m}/v,C}} ,$$

which is a restatement of Proposition 3.5.8.

To prove (2), we use the formula giving the absolute discriminant in terms of the relative discriminant ideal (Theorem 2.5.1). Thus,

$$d(L) = (-1)^{R_2-[L:K]r_2} d(K)^n \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \ .$$

Since by Proposition 3.5.8 we know that $R_2 = nr_2 + f_\infty n/2$ with $n = [L : K]$, the theorem is proved. We could have given the result without using the conductor $\mathfrak{f}$, replacing $f_\infty n/2$ by

$$(m_\infty n - \sum_{v \in \mathfrak{m}_\infty} \delta(h_{\mathfrak{m},C} - h_{\mathfrak{m}/v,C}))/2 \ ,$$

but of course this would have been ugly. □

Corollary 3.5.12. *Assume that* $(\mathfrak{m}, C)$ *is the conductor of the Abelian extension* L/K *and that* $\ell = [L : K]$ *is* prime.

(1) *We have* $\mathfrak{d}(L/K) = \mathfrak{m}_0^{\ell-1}$, *where* $\mathfrak{m}_0$ *is the finite part of* $\mathfrak{m}$.
(2) *If* $\mathfrak{p}$ *is a prime ideal dividing* $\mathfrak{m}$ (*i.e., if* $\mathfrak{p}$ *ramifies in* L/K), *then* $v_\mathfrak{p}(\mathfrak{m}) \geq 2$ *if and only if* $\mathfrak{p}$ *is above* ℓ.

Proof. (1). We always have $h_{\mathfrak{n},C} \mid h_{\mathfrak{m},C}$ for all $\mathfrak{n} \mid \mathfrak{m}$, and when $\mathfrak{m}$ is the conductor, we also have $h_{\mathfrak{n},C} < h_{\mathfrak{m},C}$ for all $\mathfrak{n} \mid \mathfrak{m}$ different from $\mathfrak{m}$. Thus, when $\ell = h_{\mathfrak{m},C}$ is prime, we must have $h_{\mathfrak{n},C} = 1$ for all $\mathfrak{n} \mid \mathfrak{m}$ other than $\mathfrak{m}$, and (1) easily follows from the theorem. Note, however, that it can also easily be proved directly (see Exercise 14).

Statement (2) is simply a reformulation of Proposition 3.3.21. □

To conclude, we see that we have quite a good hold on the Abelian extension L/K, except that we do not know an explicit description of L — for example, by a relative defining polynomial. This is the difficult part of Takagi's theorem, so it is not surprising. We will see in Chapters 5 and 6 how this problem is solved in algorithmic practice.

3.6 Exercises for Chapter 3

1. The aim of this exercise is to construct explicitly a non-Abelian unramified extension of a number field. You will need the techniques of Chapters 5 and 6, as well as a package such as Pari/GP, Kant/Kash, or Magma, to perform the computations.
   a) Let $K = \mathbb{Q}(\sqrt{458})$. Show that the class number of K is equal to 2 and that the Hilbert class field of K is the field $H_1 = K(\sqrt{2})$.
   b) Show that the class number of H_1 is equal to 3 and that the Hilbert class field H_2 of H_1 is the field $H_1(\alpha)$, where α is a root of the polynomial $x^3 - 4x - 1$.

c) Show that H_2 is an unramified extension of K, show that H_2/K is not Abelian, and compute relative and absolute defining polynomials for H_2 over K.

2. Perform similar computations with the imaginary quadratic field $K = \mathbb{Q}(\sqrt{-30})$.
3. Let L/K be a *Galois* extension of number fields, and denote as usual by $L(1)$ the Hilbert class field of L. Show that $L(1)/K$ is a Galois extension. More generally, let $\mathfrak{m}$ be a modulus of L stable by $\mathrm{Gal}(L/K)$ (in other words, such that $\sigma(\mathfrak{m}) = \mathfrak{m}$ for all $\sigma \in \mathrm{Gal}(L/K)$) and let C be a congruence subgroup modulo $\mathfrak{m}$ also stable by $\mathrm{Gal}(L/K)$. If N/L denotes the ray class field corresponding to $(\mathfrak{m}, C)$ by the Takagi correspondence, show that N/K is a Galois extension.
4. Prove the formula for $\phi(\mathfrak{m})$ given in the text.
5. Let $(\mathfrak{m}_1, C_1)$ and $(\mathfrak{m}_2, C_2)$ be two congruence subgroups such that $\mathfrak{m}_2 \mid \mathfrak{m}_1$ and $C_2 = C_1 P_{\mathfrak{m}_2}$. Show that there is a natural map from $Cl_{\mathfrak{m}_1}/\overline{C_1}$ to $Cl_{\mathfrak{m}_2}/\overline{C_2}$ and that this map is surjective.
6. Show that if $\mathcal{C}_1$ and $\mathcal{C}_2$ are two classes of congruence subgroups, one can sensibly define the intersection $\mathcal{C}_1 \cap \mathcal{C}_2$ and product $\mathcal{C}_1\mathcal{C}_2$ of these two classes.
7. Denote by $\mathfrak{f}(\mathcal{C})$ the conductor of a class $\mathcal{C}$ of congruence subgroups. Show the following.
   a) If $\mathcal{C}_1 \subset \mathcal{C}_2$, then $\mathfrak{f}(\mathcal{C}_2) \mid \mathfrak{f}(\mathcal{C}_1)$.
   b) $\mathfrak{f}(\mathcal{C}_1 \cap \mathcal{C}_2) = \mathrm{lcm}(\mathfrak{f}(\mathcal{C}_1), \mathfrak{f}(\mathcal{C}_2))$.
   c) $\mathfrak{f}(\mathcal{C}_1\mathcal{C}_2) \mid \gcd(\mathfrak{f}(\mathcal{C}_1), \mathfrak{f}(\mathcal{C}_2))$.
   d) Show that, even in the case $K = \mathbb{Q}$, equality does not necessary hold in this last result.
8. Prove Proposition 3.3.19.
9. Prove Proposition 3.3.20.
10. Give the complete list of possible moduli $\mathfrak{m}$ such that $\mathcal{N}(\mathfrak{m}_0) \leq 50$ for the real quadratic field $\mathbb{Q}(\sqrt{2})$. Do you see a pattern?
11. Let $K = \mathbb{Q}$, D a fundamental discriminant, $L = \mathbb{Q}(\sqrt{D})$, and p a prime number such that $p \nmid D$. Denote by τ the unique nontrivial field automorphism of L.
    a) Show that the Frobenius homomorphism σ_p is equal to the identity if p is split and is equal to τ otherwise.
    b) Deduce that $\mathrm{Art}(n\mathbb{Z}) = \left(\frac{L/K}{n\mathbb{Z}}\right)$ is the identity if $\left(\frac{D}{n}\right) = 1$ and is equal to τ if $\left(\frac{D}{n}\right) = -1$.
    c) Express Artin's reciprocity law (more precisely, Theorem 3.4.3 (2)) using the Jacobi symbol $\left(\frac{D}{n}\right)$, and deduce the quadratic reciprocity law.
12. Show that $\mathcal{N}_{L/K}(I_{\mathfrak{m},L})$ does not necessarily contain $P_{\mathfrak{m}}$. In fact, is it possible that $\mathcal{N}_{L/K}(I_{\mathfrak{m},L})$ contains $P_{\mathfrak{m}}$ when $L \neq K$?
13. Prove Theorem 3.5.3 using Theorem 3.4.4.
14. Let $(\mathfrak{m}, C)$ be the conductor of an Abelian extension L/K of prime degree ℓ. Compute $\mathfrak{f}(\chi)$ for all the characters of the extension L/K (or, equivalently, of the congruence subgroup $(\mathfrak{m}, C)$) and conclude that $\mathfrak{d}(L/K) = \mathfrak{m}_0^{\ell-1}$.
15. Let $(\mathfrak{m}, C)$ be the conductor of the Abelian extension L/K, where we assume that $[L : K] = \ell^r$ with ℓ prime.
    a) Generalizing Corollary 3.5.12, show that

$$\mathfrak{m}_0^{\phi(\ell^r)} = \mathfrak{m}_0^{(\ell-1)\ell^{r-1}} \mid \mathfrak{d}(L/K) \ .$$

b) Show that $\mathfrak{d}(L/K)$ is always the $(\ell-1)$st power of an ideal.

c) Give an example where $\mathfrak{m}_0^{(\ell-1)\ell^{r-1}} \neq \mathfrak{d}(L/K)$.

16. Let K be a number field, and let $\mathfrak{m} = 4\mathbb{Z}_K$ be the modulus whose finite part is the principal ideal generated by 4 and with no infinite part.

a) Assume that K is a quadratic field. Show that $|Cl_{\mathfrak{m}}(K)|$ is odd if and only if K is a real quadratic field of discriminant equal to 8 or to a prime number p (necessarily congruent to 1 modulo 4).

b) Assume that $|Cl_{\mathfrak{m}}(K)|$ is odd. Show that K is necessarily totally real. (I do not have a complete answer to this question.)

4. Computational Class Field Theory

In Chapter 3 we gave the main theoretical results concerning global class field theory over number fields. We are now going to study this subject from the algorithmic point of view. In the present chapter, we give efficient algorithms for computing ray class groups of number fields and for computing the conductor and norm group of the Abelian extensions corresponding to congruence subgroups of these ray class groups by Takagi's Theorem 3.5.1. Thanks to Proposition 3.5.8 and Theorem 3.5.11, this allows us to compute their signature and discriminant.

In the next two chapters, we will explain how to solve the more difficult problem of explicitly constructing relative or absolute defining polynomials for these Abelian extensions, and we will give some applications, particularly to the construction of number fields of small discriminant.

The following exact sequence associated to the ray class group corresponding to a modulus $\mathfrak{m}$ is an immediate consequence of Proposition 3.2.3:

$$1 \longrightarrow (\mathbb{Z}_K/\mathfrak{m})^*/\mathrm{Im}(U(K)) \longrightarrow Cl_{\mathfrak{m}}(K) \longrightarrow Cl(K) \longrightarrow 1 \ .$$

To compute the ray class group $Cl_{\mathfrak{m}}(K)$ from this exact sequence, there are three problems that a priori may seem difficult. First, the exact sequence may not split. Hence, although it may be easy to compute the cardinality $h_{\mathfrak{m}}$ of $Cl_{\mathfrak{m}}(K)$, it may not be easy to compute its structure. Second, we will need to compute the structure of the group $(\mathbb{Z}_K/\mathfrak{m})^*$, and again this may not be simple. Finally, we need to compute the image of the units in this group and compute the quotient.

When doing this by hand, one gets the impression that these tasks are not easy. In fact, this is quite a false impression, and we are going to see that suitable systematic use of the (ordinary) Smith and Hermite normal forms will lead to a nice and complete algorithmic solution to all of the above problems. Using the same tools, we can also compute the group $U_{\mathfrak{m}}(K)$, which also enters in Proposition 3.2.3 (see Exercise 1). Thus, the basic tools we will need are algorithms to compute with Abelian groups.

This chapter is divided as follows. In Section 4.1, we describe the tools necessary for dealing with finitely generated Abelian groups (usually finite). In Section 4.2, we apply these tools to the algorithmic computation of the groups $(\mathbb{Z}_K/\mathfrak{m})^*$ for an arbitrary modulus $\mathfrak{m}$. In Section 4.3, we give a complete algorithm for computing ray class groups of number fields and give the

corresponding principal ideal algorithms. In Section 4.4, we explain how to perform a number of additional explicit computations in class field theory. We defer to Chapters 5 and 6 for algorithms to compute explicit polynomials and for examples.

4.1 Algorithms on Finite Abelian groups

4.1.1 Algorithmic Representation of Groups

In this section, which is an expanded version of [Co-Di-Ol7], we consider finitely generated (in fact, usually finite) Abelian groups, which in view of our applications will be written multiplicatively. When we use the word "group", we will always mean a finitely generated Abelian group.

Even though we will work with Abelian groups, we will denote the group operation multiplicatively since the groups we will consider are, for example, class and unit groups or the group $(\mathbb{Z}_K/\mathfrak{m})^*$, which are all written multiplicatively.

We will systematically use the following matrix notation. If $\mathcal{A}$ is a group and $(\alpha_1, \ldots, \alpha_r)$ are elements of $\mathcal{A}$, we let A be the row vector of the α_i. If X is a column vector with integer entries x_i, we denote by AX the element $\prod_i \alpha_i^{x_i}$ of $\mathcal{A}$. More generally, if M is a matrix with r rows having integer entries, we denote by AM the row vector of the elements $\beta_j = AM_j$, where M_j denotes the jth column of M.

Since the group operation is written multiplicatively, it is necessary to get used to this notation (which is, of course, more natural when the group is written additively), but it is extremely practical.

We will use the following additional notation. If A and B are *row* vectors, or matrices with the same number of rows, we denote by $(A|B)$ the (horizontal) concatenation of A and B. If X and Y are *column* vectors, or matrices with the same number of columns, we will denote by $\binom{X}{Y}$ the (vertical) concatenation of X and Y.

We will always use row vectors to represent lists of elements in some Abelian group, while column vectors and matrices will always have integer (or sometimes rational) entries.

Finally, if $\mathcal{A}$ is a group, we will denote by $1_{\mathcal{A}}$ the unit element of $\mathcal{A}$ and by $\mathbf{1}_{\mathcal{A}}$ a row vector of unit elements of $\mathcal{A}$.

Definition 4.1.1. *Let $\mathcal{A}$ be a group, $G = (g_1, \ldots, g_r)$ be elements of $\mathcal{A}$, and M be an $r \times k$ integral matrix. We say that (G, M) is a* system of generators and relations *for $\mathcal{A}$ if the g_i are generators and if any relation between the g_i is a linear combination with integer coefficients of the columns of the matrix M. In matrix terms, this can be written concisely as follows:*

$$\alpha \in \mathcal{A} \iff \exists X \in \mathbb{Z}^r,\ GX = \alpha\ ,$$
$$GX = 1_{\mathcal{A}} \iff \exists Y \in \mathbb{Z}^k,\ X = MY\ .$$

In particular, we have $GM = \mathbf{1}_{\mathcal{A}}$.

A reformulation of the above definition is the existence of the following exact sequence (called a *presentation* of $\mathcal{A}$):

$$\mathbb{Z}^k \xrightarrow{M} \mathbb{Z}^r \xrightarrow{G} \mathcal{A} \longrightarrow 1 \ .$$

Definition 4.1.2. *Let $\mathcal{A}$ be a group. We say that (A, D_A) is a* Smith normal form *for $\mathcal{A}$ if (A, D_A) is a system of generators and relations for $\mathcal{A}$, if D_A is a diagonal matrix in Smith normal form (in other words, the diagonal entries a_i are nonnegative and satisfy $a_{i+1} \mid a_i$ for $i < r$), and if no diagonal entry is equal to 1 (if $\mathcal{A}$ is infinite of rank n, this implies that $a_i = 0$ for $1 \leq i \leq n$ and $a_i > 0$ for $n < i \leq r$).*

The elementary divisor theorem tells us that there exists a Smith normal form and that the matrix D_A is unique. However, the generators A are not unique.

The following algorithm, although immediate, will be of constant use.

Algorithm 4.1.3 (SNF for Finite Groups). Let (G, M) be a system of generators and relations for a finite group $\mathcal{A}$. This algorithm computes a Smith normal form (A, D_A) for $\mathcal{A}$. It also outputs a matrix U_a that will be essential for discrete logarithm computations.

1. [Apply HNF] Let H be the Hermite normal form of the matrix M obtained by applying an HNF algorithm. If H is not a square matrix (equivalently, if M is not of maximal rank), output an error message saying either that M cannot be a complete system of relations or that $\mathcal{A}$ is an infinite group, and terminate the algorithm.
2. [Apply SNF] Using a Smith normal form algorithm, compute unimodular matrices U and V and a diagonal matrix D in Smith normal form such that $UHV = D$. Set $A' \leftarrow GU^{-1}$.
3. [Remove trivial components] Let n be the largest i such that $D_{i,i} \neq 1$ (0 if none exist). Let D_A be the matrix obtained from D by keeping only the first n rows and columns, let A be the row vector obtained by keeping only the first n entries of A', and let U_a be the (not necessarily square) matrix obtained by keeping only the first n rows of U. Output (A, D_A), output U_a, and terminate the algorithm.

This algorithm's validity is clear. Note the important relation $AU_a = G$. □

The reason for keeping the matrix U_a is also clear: if an element α of $\mathcal{A}$ is known on the generators G as $\alpha = GX$, then on the new generators A we have $\alpha = A(U_a X)$, so the matrix U_a allows us to go from one system of generators to another.

The following two definitions are rather imprecise but useful.

Definition 4.1.4. *Let $\mathcal{A}$ be a group. We say that we have* effectively computed *the group $\mathcal{A}$ if we have done the following.*

(1) *We have computed a system (G, M) of generators and relations for the group $\mathcal{A}$ or, equivalently, by Algorithm 4.1.3, a Smith normal form (A, D_A).*
(2) *We have found an efficient algorithm that, given an element $\alpha \in \mathcal{A}$, finds a column vector X with integer entries such that $\alpha = GX$ (or $\alpha = AX$ if we have the SNF). The column vector X will be called the* discrete logarithm *of α on the given generators.*

When we say that we have computed an Abelian group, or that a group is known, we will always mean that we have effectively computed it in the above sense. Note that this definition is not really a mathematical one since we have not said what we mean by an efficient algorithm.

A similar definition applies to maps.

Definition 4.1.5. *Let $\mathcal{A}$ and $\mathcal{B}$ be two groups and ψ a homomorphism from $\mathcal{A}$ to $\mathcal{B}$. We say that ψ is* effective *or if the following properties are true. If $\mathcal{B}$ has been computed, then if $\alpha \in \mathcal{A}$, we can compute $\psi(\alpha)$ expressed on the generators of $\mathcal{B}$. Similarly, if $\mathcal{A}$ has been computed, then if $\beta \in \mathrm{Im}(\psi)$, we can compute $\alpha \in \mathcal{A}$ such that $\beta = \psi(\alpha)$.*

4.1.2 Algorithmic Representation of Subgroups

A subgroup of a known group can of course be represented abstractly as (A, D_A) as for any other group, but this is often not convenient since it forgets the subgroup structure. There is an alternate, richer representation, based on the following proposition.

Proposition 4.1.6. *Let $\mathcal{B} = (B, D_B)$ be a finite Abelian group given in SNF, where $B = (\beta_i)_{1 \le i \le n}$. There is a natural one-to-one correspondence between subgroups $\mathcal{A}$ of $\mathcal{B}$ and integral matrices H in Hermite normal form satisfying $H^{-1} D_B \in \mathcal{M}_n(\mathbb{Z})$. The correspondence is as follows.*

(1) *The subgroup $\mathcal{A}$ associated to such a matrix H is the group given by generators and relations (not necessarily in SNF), as $\mathcal{A} = (BH, H^{-1} D_B)$.*
(2) *Conversely, if $\mathcal{A}$ is a subgroup of $\mathcal{B}$ and B' is a row vector of generators of $\mathcal{A}$, we can write $B' = BP$ for some integer matrix P. The corresponding matrix H is the Hermite normal form of the matrix $(P | D_B)$.*
(3) *Let H be a matrix in HNF, and let $\mathcal{A}$ be the corresponding subgroup. Then $|\mathcal{A}| = |\mathcal{B}| / \det(H)$ or, equivalently, $|\mathcal{B}/\mathcal{A}| = [\mathcal{B} : \mathcal{A}] = \det(H)$.*

Proof. Let $B = (\beta_i)_{1 \le i \le n}$ and let $D_B = \mathrm{diag}((b_i)_{1 \le i \le n})$, where $\mathrm{diag}((b_i)_i)$ denotes the diagonal matrix whose diagonal entries are the b_i. By definition, the following sequence is exact:

$$1 \longrightarrow \bigoplus_{i=1}^{n} b_i\mathbb{Z} \longrightarrow \mathbb{Z}^n \xrightarrow{\phi} \mathcal{B} \longrightarrow 1 \ ,$$

where

$$\phi(x_1, \ldots, x_n) = \prod_{1 \le i \le n} \beta_i^{x_i} \ .$$

Let $(\varepsilon_i)_{1 \le i \le n}$ be the canonical basis elements of $\mathbb{Z}^n$, and let Λ be the lattice defined by $\Lambda = \bigoplus_i b_i\varepsilon_i$. We thus have a canonical isomorphism $\mathcal{B} \simeq \mathbb{Z}^n/\Lambda$, obtained by sending the ith generator β_i of $\mathcal{B}$ to the class of ε_i.

Subgroups of $\mathbb{Z}^n/\Lambda$ are of the form Λ'/Λ, where Λ' is a lattice such that $\Lambda \subset \Lambda' \subset \mathbb{Z}^n$. Such a lattice Λ' can be uniquely defined by a matrix H in Hermite normal form so that the columns of this matrix express a $\mathbb{Z}$-basis of Λ' on the ε_i. The condition $\Lambda' \subset \mathbb{Z}^n$ means that H has integer entries, and the condition $\Lambda \subset \Lambda'$ means that $H^{-1}D_B$ also has integer entries, since it is the matrix that expresses the given basis of Λ in terms of that of Λ'. In terms of generators, this correspondence translates into the equality $B' = BH$. Furthermore, $B'X = 1_{\mathcal{B}}$ if and only if $BHX = 1_{\mathcal{B}}$, hence $HX = D_BY$, or $X = H^{-1}D_BY$, and so if $\mathcal{A}$ is the subgroup of $\mathcal{B}$ corresponding to Λ'/Λ, it is given in terms of generators and relations by $(BH, H^{-1}D_B)$, proving (1).

For (2), we note that $BD_B = \mathbf{1}_{\mathcal{B}}$, hence if $B'' = B(P|D_B)$, we have simply added some $1_{\mathcal{A}}$'s to the generators of $\mathcal{A}$. Thus, the group can be defined by the generators B'' and the matrix of relations of maximal rank $(P|D_B)$, hence also by (B'', H), where H is the Hermite normal form of this matrix.

For (3), we know that $H^{-1}D_B$ expresses a basis of Λ in terms of a basis of Λ'; hence

$$|\mathcal{A}| = |\Lambda'/\Lambda| = \det(H^{-1}D_B) = |\mathcal{B}| \,/\, \det(H) \ .$$

□

Example. The matrix H corresponding to the subgroup $\{1_{\mathcal{B}}\}$ of $\mathcal{B}$ is $H = D_B$, and the matrix corresponding to the subgroup $\mathcal{B}$ of $\mathcal{B}$ is $H = I_n$.

A matrix H in HNF such that $H^{-1}D \in \mathcal{M}_n(\mathbb{Z})$ will be called a *left divisor* of D. We will implicitly assume that all left divisors are in HNF, since if H is a left divisor of D, then for any unimodular matrix U, HU is also a left divisor of D. The above proposition states that subgroups of $\mathcal{B}$ are in canonical one-to-one correspondence with left divisors of D_B. Hence, it is usually better to represent a subgroup $\mathcal{A}$ of $\mathcal{B}$ by the matrix H.

If we really want a Smith normal form for $\mathcal{A}$, we simply apply Algorithm 4.1.3 to the system of generators and relations $(BH, H^{-1}D_B)$ for the group $\mathcal{A}$.

Conversely, if we are given a subgroup $\mathcal{A}$ by an SNF (A, D_A) together with an injective group homomorphism ψ from $\mathcal{A}$ to $\mathcal{B}$, we can compute the HNF matrix H associated to $\psi(\mathcal{A})$ as follows. Using the discrete logarithm

algorithm in $\mathcal{B}$, we compute an integral matrix P such that $\psi(A) = BP$, and H is simply the Hermite normal form of the matrix $(P|D_B)$. This is a restatement of Proposition 4.1.6 (2).

Finding a discrete logarithm algorithm for a subgroup is done as follows. If $\mathcal{A}$ is represented as a subgroup of $\mathcal{B}$ by a matrix H, then to compute the discrete logarithm of $\beta \in \mathcal{A}$, we first apply the discrete logarithm algorithm in the full group $\mathcal{B}$, thus obtaining an integer vector X such that $\beta = BX$. Hence $\beta = BH(H^{-1}X)$, so $H^{-1}X$ is the discrete logarithm on the generators BH (it is an integer vector if and only if $\beta \in \mathcal{A}$). We can of course left-multiply by the matrix U_a output by Algorithm 4.1.3 to give the discrete logarithm on the SNF (A, D_A) if we have explicitly computed it.

In rest of this section, we are going to give a number of algorithms for computing with Abelian groups, such as computing kernels, inverse images, images, quotients, extensions, and so forth. In each case, we will choose the most suitable representation for the result, either as an abstract group given in SNF or as a subgroup by an HNF matrix H which is a left divisor of an SNF matrix D as above. Going back and forth between these representations is done as we have just explained. All the algorithms are easy but technical, hence the reader is advised at first to skim through the rest of this section, and to read it carefully only for an actual computer implementation.

4.1.3 Computing Quotients

Let

$$\mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1$$

be an exact sequence of Abelian groups. In this section, we assume that $\mathcal{A}$ and $\mathcal{B}$ are known (in the sense of Definition 4.1.4) and that we want to compute $\mathcal{C}$. We assume also that the maps ψ and ϕ are effective. We do not necessarily assume that ψ is injective. Let (A, D_A) (resp., (B, D_B)) be a Smith normal form of $\mathcal{A}$ (resp., $\mathcal{B}$) (it is only necessary for these to be generators and relations, but usually they will be in SNF).

Since ϕ is surjective, it is clear that if we set $B' = \phi(B)$, B' is a system of generators of $\mathcal{C}$. We must find all the relations between them. Let V be such a relation, expressed as a column vector. Then

$$B'V = 1_{\mathcal{C}} \iff \phi(BV) = 1_{\mathcal{C}} \iff BV \in \mathrm{Im}(\psi) \iff BV = \psi(A)X$$

for a certain integer vector X.

Since the group $\mathcal{B}$ is known, we know how to compute algorithmically a matrix P such that $\psi(A) = BP$. Hence

$$\begin{aligned} B'V = 1_{\mathcal{C}} &\iff BV = BPX &&\iff B(V - PX) = 1_{\mathcal{B}} \\ &\iff V - PX \in \mathrm{Im}(D_B) &&\iff V \in \mathrm{Im}(P|D_B)\ . \end{aligned}$$

It follows that $(\phi(B), (P|D_B))$ is a system of generators and relations for $\mathcal{C}$, and we finish using Algorithm 4.1.3. Formally, this gives the following.

Algorithm 4.1.7 (Quotient of Groups). Given two groups $\mathcal{A} = (A, D_A)$ and $\mathcal{B} = (B, D_B)$ in SNF and an exact sequence $\mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1$, this algorithm computes the SNF of the group $\mathcal{C}$.

1. [Compute P] Using the discrete logarithm algorithm in $\mathcal{B}$, compute a matrix P such that $\psi(A) = BP$.
2. [Compute SNF] Apply Algorithm 4.1.3 to $(\phi(B), (P|D_B))$, output the SNF (C, D_C) of the result and the auxiliary matrix U_a, and terminate the algorithm.

To obtain a corresponding discrete logarithm algorithm, we proceed as follows. Let $\gamma \in \mathcal{C}$. Since ϕ is surjective and is effective, we can find $\beta \in \mathcal{B}$ such that $\gamma = \phi(\beta)$. Since we know how to compute discrete logarithms in $\mathcal{B}$, we can find X such that $\beta = BX$. Hence

$$\gamma = \phi(BX) = \phi(B)X = CU_aX \ ,$$

where U_a is the auxiliary matrix output by Algorithm 4.1.3. It follows that the discrete logarithm of γ on the generators C is given by the vector U_aX.

Remark. If the group $\mathcal{A}$ is given as a subgroup of $\mathcal{B}$ by a left HNF divisor H of D_B, and the map ψ is the natural injection, the algorithm simplifies considerably since the HNF of $(P|D_B)$ is equal to H, since $P = H$ is a left divisor of D_B. Thus we simply apply Algorithm 4.1.3 to the system of generators and relations $(\phi(B), H)$.

4.1.4 Computing Group Extensions

Let $\mathcal{A} = (A, D_A)$ and $\mathcal{C} = (C, D_C)$ be two groups given in SNF, and assume now that we have an exact sequence

$$1 \longrightarrow \mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1 \ .$$

We want to compute the SNF (B, D_B) of the group $\mathcal{B}$.

Let B' be arbitrarily chosen such that $\phi(B') = C$. If $\beta \in \mathcal{B}$ then for some vector Y, we have $\phi(\beta) = CY = \phi(B')Y = \phi(B'Y)$, hence $\beta - B'Y \in \mathrm{Ker}(\phi)$, so $\beta - B'Y = \psi(AX)$ for some integer vector X. Thus $\beta = \psi(A)X + B'Y = (\psi(A)|B')R$, where $R = \left(\frac{X}{Y}\right)$. It follows that $(\psi(A)|B')$ forms a generating set for $\mathcal{B}$ (this is, of course, trivial, but we prefer to do everything in matrix terms).

Let us find the relations between these generators. If $R = \left(\frac{X}{Y}\right)$ is such a relation, we have $\psi(A)X + B'Y = 1_{\mathcal{B}}$. If we apply ϕ to this relation, we obtain $\phi(B')Y = CY = 1_{\mathcal{C}}$, hence $Y \in \mathrm{Im}D_C$, so that $Y = D_CY_1$ for some integral vector Y_1. Thus we have $\psi(A)X + B'D_CY_1 = 1_{\mathcal{B}}$.

Set $B'' = B'D_C$. Then $\phi(B'') = \phi(B')D_C = CD_C = \mathbf{1}_{\mathcal{C}}$, hence the entries of B'' are in $\mathrm{Ker}(\phi) = \mathrm{Im}(\psi)$, and since we have a discrete logarithm algorithm in $\mathcal{A}$, we can find a matrix P such that $B'' = \psi(AP) = \psi(A)P$.

So finally, the equation for our relation is

$$\begin{aligned}\psi(A)X + \psi(A)PY_1 = 1_{\mathcal{B}} &\iff \psi(A)(X + PY_1) = 1_{\mathcal{B}} \\ &\iff A(X + PY_1) = 1_{\mathcal{A}} \iff X + PY_1 \in \mathrm{Im}D_A \\ &\iff X + PY_1 = D_A T\end{aligned}$$

for some integer vector T (note that here we have used the injectivity of ψ).

In other words, $R = \binom{X}{Y}$ is a relation if and only if we have

$$R = \begin{pmatrix} D_A & -P \\ 0 & D_C \end{pmatrix} \begin{pmatrix} T \\ Y_1 \end{pmatrix}$$

for some integer vectors T and Y_1. We then obtain the SNF as before by applying Algorithm 4.1.3. Formally, this gives the following.

Algorithm 4.1.8 (Group Extensions). Given two groups $\mathcal{A} = (A, D_A)$ and $\mathcal{C} = (C, D_C)$ in SNF, and an exact sequence $1 \longrightarrow \mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1$ with ψ and ϕ effective, this algorithm computes the SNF (B, D_B) of the group $\mathcal{B}$.

1. [Compute generators] Compute B' such that $\phi(B') = C$ (which can be done since ϕ is effective), and compute $\psi(A)$.
2. [Compute P] Set $B'' \leftarrow B'D_C$, and let A'' be such that $B'' = \psi(A'')$. (B'' is in the image of ψ and A'' can be found since ψ is effective.) Using the discrete logarithm algorithm in $\mathcal{A}$, compute an integral matrix P such that $A'' = AP$.
3. [Terminate] Set $G \leftarrow (\psi(A)|B')$ and $M \leftarrow \left(\begin{smallmatrix} D_A & -P \\ 0 & D_C \end{smallmatrix}\right)$. Apply Algorithm 4.1.3 to the system of generators and relations (G, M), output the SNF (B, D_B) of $\mathcal{B}$ and the auxiliary matrix U_a, and terminate the algorithm.

It is easy to obtain a corresponding discrete logarithm algorithm. Let $\beta \in \mathcal{B}$. Using the discrete logarithm algorithm in $\mathcal{C}$, we can find Y such that $\phi(\beta) = CY = \phi(B')Y$, hence $\phi(\beta - B'Y) = 1_{\mathcal{C}}$, so $\beta - B'Y \in \mathrm{Im}(\psi)$. Using the discrete logarithm algorithm in $\mathcal{A}$, we obtain $\beta - B'Y = \psi(A)X$ for some X, so $\beta = (\psi(A)|B')\binom{X}{Y}$. Finally, this gives $\beta = BU_a\binom{X}{Y}$; hence $U_a\binom{X}{Y}$ is our desired discrete logarithm.

Remark. From the above discussion, it is clear that the matrix $-P$ measures the obstruction to the fact that the exact sequence is split. More precisely, if $-P = 0$, the sequence splits; conversely, if the sequence splits, then one can find generators such that $-P = 0$ (see Exercise 2).

4.1.5 Right Four-Term Exact Sequences

In view of our application to ray class group computations, we will also use right four-term exact sequences (we will see left four-term exact sequences in Section 4.1.7). More precisely, assume that we have an exact sequence of the form

$$\mathcal{E} \xrightarrow{\rho} \mathcal{Z} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1 \ .$$

We assume that we know the groups $\mathcal{E} = (E, D_E)$, $\mathcal{Z} = (Z, D_Z)$, and $\mathcal{C} = (C, D_C)$, and we want to compute $\mathcal{B}$. In the application we have in mind, $\mathcal{E}$ is in general infinite, in which case the diagonal entry of D_E is equal to 0 for each infinite cyclic component. This could be treated as a three-term exact sequence by introducing the quotient group $\mathcal{Z}/\rho(\mathcal{E})$, but it is more elegant and just as easy to treat it directly as a right four-term exact sequence.

We proceed essentially as in Section 4.1.4. Let B' be such that $\phi(B') = C$, and let $\beta \in \mathcal{B}$. We have $\phi(\beta) = CY$ for some Y, hence $\phi(\beta - B'Y) = 1_{\mathcal{C}}$, hence $\beta - B'Y \in \psi(\mathcal{Z})$. It follows that $\beta - B'Y = \psi(Z)X$ for some X; in other words, $(\psi(Z)|B')$ is a generating set for $\mathcal{B}$.

Let us find the relations between these generators. If $R = \left(\frac{X}{Y}\right)$ is such a relation, we have $\psi(Z)X + B'Y = 1_{\mathcal{B}}$. If we apply ϕ to this relation, we obtain $\phi(B')Y = CY = 1_{\mathcal{C}}$. Hence $Y \in \mathrm{Im} D_C$, so that $Y = D_C Y_1$ for some integral vector Y_1. Thus we have $\psi(Z)X + B'D_C Y_1 = 1_{\mathcal{B}}$.

Set $B'' = B'D_C$. Then $\phi(B'') = \phi(B')D_C = CD_C = \mathbf{1}_{\mathcal{C}}$. Thus all the entries of B'' are in $\mathrm{Ker}(\phi) = \mathrm{Im}(\psi)$. Since ψ is assumed to be effective, we can find Z' such that $\psi(Z') = B''$. Since we have a discrete logarithm algorithm in $\mathcal{Z}$, we can find a matrix P such that $Z' = ZP$.

Thus, $R = \left(\frac{X}{Y}\right)$ is a relation if and only if $\psi(Z)X + \psi(Z)PY_1 = 1_{\mathcal{B}}$, or in other words $Z(X + PY_1) \in \mathrm{Ker}(\psi) = \mathrm{Im}(\rho)$. Thus there exists a vector T such that $Z(X + PY_1) = \rho(E)T$. Using again the discrete logarithm algorithm in $\mathcal{Z}$, we can find a matrix Q such that $\rho(E) = ZQ$. Hence we get $Z(X + PY_1 - QT) = 1_{\mathcal{Z}}$ or, equivalently, $X + PY_1 - QT = D_Z T'$ for still another integer vector T'.

In other words, $R = \left(\frac{X}{Y}\right)$ is a relation if and only if we have

$$R = \begin{pmatrix} Q & D_Z & -P \\ 0 & 0 & D_C \end{pmatrix} \begin{pmatrix} T \\ T' \\ Y_1 \end{pmatrix}$$

for some integer vectors T, T', and Y_1. We then obtain the SNF as before by applying Algorithm 4.1.3. Formally, this gives the following.

Algorithm 4.1.9 (Right Four-Term Exact Sequences). Given three Abelian groups $\mathcal{E} = (E, D_E)$, $\mathcal{Z} = (Z, D_Z)$, and $\mathcal{C} = (C, D_C)$ in SNF and an exact sequence $\mathcal{E} \xrightarrow{\rho} \mathcal{Z} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1$ with ρ, ψ, and ϕ effective, this algorithm computes the SNF (B, D_B) of the group $\mathcal{B}$.

1. [Compute generators] Compute B' such that $\phi(B') = C$ (which can be done since ϕ is effective), and compute $\psi(Z)$.
2. [Compute P] Set $B'' \leftarrow B'D_C$, and let Z' be such that $B'' = \psi(Z')$ (B'' is in the image of ψ and Z' can be found since ψ is effective). Using the discrete logarithm algorithm in $\mathcal{Z}$, compute a matrix P such that $Z' = ZP$.

3. [Compute Q] Using the discrete logarithm algorithm in $\mathcal{Z}$, compute a matrix Q such that $\rho(E) = ZQ$.

4. [Terminate] Set $G \leftarrow (\psi(Z)|B')$ and $M \leftarrow \begin{pmatrix} Q & D_Z & -P \\ 0 & 0 & D_C \end{pmatrix}$. Apply Algorithm 4.1.3 to the system of generators and relations (G, M), output the SNF (B, D_B) of $\mathcal{B}$ and the auxiliary matrix U_a, and terminate the algorithm.

In our application to ray class group computations, the group $\mathcal{E}$ will be the group of units of a number field, hence finitely generated but not finite in general. As can be seen from step 3, however, only a finite set of generators is needed. Apart from this group, all of the other groups we will use are finite.

It is again easy to obtain a corresponding discrete logarithm algorithm. Let $\beta \in \mathcal{B}$. Using the discrete logarithm algorithm in $\mathcal{C}$, we can find Y such that $\phi(\beta) = CY = \phi(B')Y$, hence $\phi(\beta - B'Y) = 1_{\mathcal{C}}$ so $\beta - B'Y \in \mathrm{Im}(\psi)$. Since $\mathcal{Z}$ has been computed and ψ is effective, using the discrete logarithm algorithm in $\mathcal{Z}$ we obtain $\beta - B'Y = \psi(Z)X$ for some X, so $\beta = (\psi(Z)|B')\binom{X}{Y}$. Finally, this gives $\beta = BU_a\binom{X}{Y}$; hence $U_a\binom{X}{Y}$ is our desired discrete logarithm.

4.1.6 Computing Images, Inverse Images, and Kernels

Let $\mathcal{B} = (B, D_B)$ and $\mathcal{C} = (C, D_C)$ be two known Abelian groups, let ϕ be an effective group homomorphism from $\mathcal{B}$ to $\mathcal{C}$, and let $\mathcal{A}$ be a subgroup of $\mathcal{B}$ given by an HNF matrix H_B that is a left divisor of D_B as explained in Proposition 4.1.6. We can easily compute the image of ϕ using the following algorithm.

Algorithm 4.1.10 (Image of a Subgroup). Let $\mathcal{B} = (B, D_B)$ and $\mathcal{C} = (C, D_C)$ be two known Abelian groups in SNF, let ϕ be an effective group homomorphism from $\mathcal{B}$ to $\mathcal{C}$, and let $\mathcal{A}$ be a subgroup of $\mathcal{B}$ given by a left divisor H_B of D_B. This algorithm computes the image $\phi(\mathcal{A})$ as a subgroup of $\mathcal{C}$; in other words, it outputs a left divisor H_C of D_C that represents the subgroup $\phi(\mathcal{A})$ according to Proposition 4.1.6.

1. [Compute matrix P] Using the discrete logarithm algorithm in $\mathcal{C}$, compute an integer matrix P such that $\phi(B) = CP$.

2. [Terminate] Let $M \leftarrow (PH_B|D_C)$ be the horizontal concatenation of PH_B and D_C. Let H_C be the HNF of the matrix M (which is a left divisor of D_C). Output H_C and terminate the algorithm.

Proof. By definition of H_B, $B' = BH_B$ is a system of generators for $\mathcal{A}$, and $\phi(B') = \phi(B)H_B = C(PH_B)$. Hence, by Proposition 4.1.6 (2), the desired matrix H_C is the HNF of the matrix $(PH_B|D_C)$. □

Once again, let $\mathcal{B} = (B, D_B)$ and $\mathcal{C} = (C, D_C)$ be two known Abelian groups, let ϕ be an effective group homomorphism from $\mathcal{B}$ to $\mathcal{C}$, but now let

$\mathcal{A}$ be a subgroup of $\mathcal{C}$, given by an HNF matrix H_C that is a left divisor of D_C. We want to compute $\phi^{-1}(\mathcal{A})$ as a subgroup of $\mathcal{B}$. This is done as follows.

Algorithm 4.1.11 (Inverse Image of a Subgroup). Let $\mathcal{B} = (B, D_B)$ and $\mathcal{C} = (C, D_C)$ be two known Abelian groups in SNF, let ϕ be an effective group homomorphism from $\mathcal{B}$ to $\mathcal{C}$, and let $\mathcal{A}$ be a subgroup of $\mathcal{C}$ given by a left divisor H_C of D_C. This algorithm computes the inverse image $\phi^{-1}(\mathcal{A})$ as a subgroup of $\mathcal{B}$; in other words, it outputs a left divisor H_B of D_B that represents the subgroup $\phi^{-1}(\mathcal{A})$ according to Proposition 4.1.6.

1. [Compute P] Using the discrete logarithm algorithm in $\mathcal{C}$, compute an integral matrix P such that $\phi(B) = CP$.
2. [Compute U_1] Apply an HNF algorithm to the matrix $(P|H_C)$, and let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(P|H_C)U = (0|H)$. We can discard the matrices U_2, U_3, U_4, and H.
3. [Terminate] Let H_B be the HNF of the matrix $(U_1|D_B)$. Output H_B and terminate the algorithm.

Proof. Let X be an integer vector representing an element of $\mathcal{B}$ on the generators B. We have

$$\begin{aligned} BX \in \phi^{-1}(\mathcal{A}) &\iff \phi(B)X \in \mathcal{A} \iff CPX = CH_CY \\ &\iff C(PX - H_CY) = 1_{\mathcal{C}} \end{aligned}$$

for some integer vector Y, so $BX \in \phi^{-1}(\mathcal{A}) \iff PX - H_CY = D_CZ$ for some integer vector Z. We know, however, that H_C is a left divisor of D_C, so $D_C = H_CH'_C$ for some integer matrix H'_C. Hence X represents an element of $\phi^{-1}(\mathcal{A})$ on B if and only if there exist integer vectors Y and Z such that $PX - H_C(Y + H'_CZ) = 0$, hence if and only if there exists an integer vector T such that $PX + H_CT = 0$. (Indeed, $T = -(Y + H'_CZ)$ exists, but conversely if T is given, we can choose $Z = 0$ and $Y = -T$.)

We now use [Coh0, Proposition 2.4.9], which tells us that if

$$(P|H_C)\begin{pmatrix} U_1 & U_2 \\ U_3 & U_4 \end{pmatrix} = (0|H)$$

is the HNF decomposition of the matrix $(P|H_C)$, a $\mathbb{Z}$-basis of the kernel of $(P|H_C)$ is given by the columns of the matrix $\left(\begin{smallmatrix} U_1 \\ U_3 \end{smallmatrix}\right)$. In other words, $PX + H_CT = 0$ if and only if there exists a column vector X_1 such that

$$\begin{pmatrix} X \\ T \end{pmatrix} = \begin{pmatrix} U_1 \\ U_3 \end{pmatrix} X_1 \ .$$

Hence X represents an element of $\phi^{-1}(\mathcal{A})$ if and only if it is in the image of U_1, hence a generating system of $\phi^{-1}(\mathcal{A})$ is given by BU_1, and we conclude by Proposition 4.1.6 (2). □

Remarks

(1) To compute the kernel $\phi^{-1}(\{1_C\})$ of the map ϕ, we apply the above algorithm to the matrix $H_C = D_C$, which is the matrix representing $\{1_C\}$ in the subgroup representation.
(2) It is easy to write an algorithm for computing the *cokernel* $C/\mathrm{Im}(\phi)$ of ϕ; see Exercise 4.

Finally, note the following lemma, which gives an important property of the matrix U_1 used in the above algorithm.

Lemma 4.1.12. *Assume that we have a matrix equality of the form*

$$(P|H_1)U = (0|H) \ ,$$

where U is invertible and H_1 is a square matrix with nonzero determinant. Let r be the number of columns of P or, equivalently, the number of 0 columns on the right-hand side. Then the upper-left $r \times r$ submatrix U_1 of U has nonzero determinant equal to $\pm\det(H_1)/\det(H)$, where the sign is equal to the determinant of U.

Proof. Write $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$. One easily checks the block matrix identity

$$\begin{pmatrix} H_1 & 0 \\ -P & H \end{pmatrix} \begin{pmatrix} U_1 & U_2 \\ 0 & I \end{pmatrix} = \begin{pmatrix} H_1 & 0 \\ 0 & H_1 \end{pmatrix} \begin{pmatrix} U_1 & U_2 \\ U_3 & U_4 \end{pmatrix} \ .$$

Since $\det(H_1) \neq 0$, it follows that $\det(U_1)\det(H) = \pm\det(H_1)$, where the sign is equal to $\det(U)$. This proves the lemma. □

Note that it is easy to write the inverse of U_1 in $\mathrm{GL}_r(\mathbb{Q})$ in terms of the block matrix decomposition of U^{-1}; see Exercise 5.

4.1.7 Left Four-Term Exact Sequences

In Chapter 7, we will also use *left* four-term exact sequences. More precisely, assume that we have an exact sequence of the form

$$1 \longrightarrow \mathcal{E} \xrightarrow{\rho} \mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \ .$$

We assume that we know the groups $\mathcal{E} = (E, D_E)$, $\mathcal{B} = (B, D_B)$, and $\mathcal{C} = (C, D_C)$, and we want to compute $\mathcal{A}$. In the application we have in mind, $\mathcal{E}$, $\mathcal{A}$, and $\mathcal{B}$ will in general be infinite, and in that case the diagonal entries of the corresponding SNFs are equal to 0 for each infinite cyclic component, but apart from this everything that we have done remains valid.

The above left four-term exact sequence could as usual be treated as a concatenation of shorter exact sequences, but we prefer to treat it directly as

a left four-term exact sequence. We will combine the ideas of Sections 4.1.6 and 4.1.4.

We first compute the kernel of ϕ, which is equal to the image of ψ, by using Section 4.1.6. Following Algorithm 4.1.11 (with $\mathcal{A} = \{1_C\}$), we first compute an integral matrix P such that $\phi(B) = CP$. Using an HNF algorithm, we compute a unimodular matrix $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ such that $(P|D_C)U = (0|H)$ is in HNF. We let H_B be the HNF of the matrix $(U_1|D_B)$, so that H_B is a left divisor of D_B such that BH_B is a generating system for $\mathrm{Ker}(\phi) = \mathrm{Im}(\psi)$.

We now follow Section 4.1.4. Let A' be such that $\psi(A') = BH_B$, which is possible since the entries of BH_B are in $\mathrm{Im}(\psi)$. Then as in Section 4.1.4, $(\rho(E)|A')$ forms a generating set for $\mathcal{A}$. Let us find the relations between these generators. If $R = \left(\begin{smallmatrix} X \\ Y \end{smallmatrix}\right)$ is such a relation, we have $\rho(E)X + A'Y = 1_{\mathcal{A}}$. Applying ψ to this relation, we obtain $\psi(A')Y = BH_BY = 1_{\mathcal{B}}$, hence $H_BY = D_BY_1$ for a certain integer vector Y_1. Since H_B is a left divisor of D_B, this gives $Y = H_B^{-1}D_BY_1$. Thus, we have $\rho(E)X + A'H_B^{-1}D_BY_1 = 1_{\mathcal{A}}$.

Set $A'' = A'H_B^{-1}D_B$. Then $\psi(A'') = \psi(A')H_B^{-1}D_B = BD_B = \mathbf{1}_{\mathcal{B}}$. Thus the entries of A'' are in $\mathrm{Ker}(\psi) = \mathrm{Im}(\rho)$, and since we have a discrete logarithm algorithm in $\mathcal{E}$, we can find a matrix Q such that $A'' = \rho(EQ) = \rho(E)Q$. Thus, the equation for our relation is

$$\begin{aligned} \rho(E)X + \rho(E)QY_1 = 1_{\mathcal{A}} &\iff \rho(E)(X + QY_1) = 1_{\mathcal{A}} \\ &\iff E(X + QY_1) = 1_{\mathcal{E}} \\ &\iff X + QY_1 \in \mathrm{Im}D_E \iff X + QY_1 = D_ET \end{aligned}$$

for some integer vector T (note that here we have used the injectivity of ρ).

In other words, $R = \left(\begin{smallmatrix} X \\ Y \end{smallmatrix}\right)$ is a relation if and only if we have

$$R = \begin{pmatrix} D_E & -Q \\ 0 & H_B^{-1}D_B \end{pmatrix} \begin{pmatrix} T \\ Y_1 \end{pmatrix}$$

for some integer vectors T and Y_1. We then obtain the SNF as before by applying Algorithm 4.1.3. Formally, this gives the following.

Algorithm 4.1.13 (Left Four-Term Exact Sequences). Given three Abelian groups $\mathcal{E} = (E, D_E)$, $\mathcal{B} = (B, D_B)$, and $\mathcal{C} = (C, D_C)$ in SNF and an exact sequence $1 \longrightarrow \mathcal{E} \xrightarrow{\rho} \mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C}$ with ρ, ψ, and ϕ effective, this algorithm computes the SNF (A, D_A) of the group $\mathcal{A}$.

1. [Compute P] Using the discrete logarithm algorithm in $\mathcal{C}$, compute an integral matrix P such that $\phi(B) = CP$.
2. [Compute $\mathrm{Ker}(\phi)$] Apply an HNF algorithm to the matrix $(P|D_C)$, let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(P|D_C)U = (0|H)$, and finally let H_B be the HNF of the matrix $(U_1|D_B)$. We can discard all the matrices computed up to now except H_B.
3. [Compute generators] Compute A' such that $\psi(A') = BH_B$ (which can be done since ψ is effective), and compute $\rho(E)$.

4. [Compute Q] Set $A'' \leftarrow A'H_B^{-1}D_B$ and let E'' be such that $A'' = \rho(E'')$ (A'' is in the image of ρ and E'' can be found since ρ is effective). Using the discrete logarithm algorithm in $\mathcal{E}$, compute an integral matrix Q such that $E'' = EQ$.

5. [Terminate] Set $G \leftarrow (\rho(E)|A')$ and $M \leftarrow \begin{pmatrix} D_E & -Q \\ 0 & H_B^{-1}D_B \end{pmatrix}$. Apply Algorithm 4.1.3 to the system of generators and relations (G, M), output the SNF (A, D_A) of $\mathcal{A}$ and the auxiliary matrix U_a, and terminate the algorithm.

As usual, it is easy to obtain a corresponding discrete logarithm algorithm. Let $\alpha \in \mathcal{A}$. Using the discrete logarithm algorithm in $\mathcal{B}$, we can find Y such that $\psi(\alpha) = BY$. Since $\psi(\alpha) \in \mathrm{Ker}(\phi)$, the vector $Z = H_B^{-1}Y$ has integral entries (see Exercise 6). Thus, $\psi(\alpha - A'Z) = BY - BH_BZ = 0$, so $\alpha - A'Z \in \mathrm{Ker}(\psi) = \mathrm{Im}(\rho)$, and hence we can find an integral vector T such that $\alpha - A'Z = \rho(E)T$. Hence $\alpha = \rho(E)T + A'Z = G\binom{T}{Z}$, and the discrete logarithm of α with respect to the generators A is equal to $U_a\binom{T}{Z}$.

4.1.8 Operations on Subgroups

It is easy to modify the preceding algorithms so that they perform operations on subgroups, represented as explained in Section 4.1.2. Thus, let $\mathcal{B} = (B, D_B)$ be a fixed Abelian group in SNF, and let $\mathcal{A}_1$ and $\mathcal{A}_2$ be subgroups of $\mathcal{B}$ given by HNF left divisors H_1 and H_2 of D_B. We want to compute their intersection and their sum, which by definition is the subgroup generated by $\mathcal{A}_1$ and $\mathcal{A}_2$. We leave the easy proof of the following algorithm to the reader (Exercise 7).

Algorithm 4.1.14 (Intersection and Sum of Subgroups). Given an Abelian group $\mathcal{B} = (B, D_B)$ in SNF and two subgroups $\mathcal{A}_1$ and $\mathcal{A}_2$ given by HNF left divisors H_1 and H_2 of D_B, this algorithm computes the HNF left divisors of D_B giving the intersection $\mathcal{A}_1 \cap \mathcal{A}_2$ and the sum $\mathcal{A}_1 + \mathcal{A}_2$ of $\mathcal{A}_1$ and $\mathcal{A}_2$.

1. [Compute HNF] Let $(H_1|H_2)\begin{pmatrix} U_1 & U_2 \\ U_3 & U_4 \end{pmatrix} = (0|H)$ be the HNF decomposition of the matrix $(H_1|H_2)$.

2. [Terminate] Let H_3 be the HNF of $(H_1U_1|D_B)$ (or, equivalently, the HNF of $(H_2U_3|D_B)$). Output H_3 as the HNF of the intersection and H as the HNF of the sum, and terminate the algorithm.

If we want the intersection or the sum of more than two subgroups, we can either apply the above algorithm recursively or directly use the HNF of the concatenation of all the matrices. The first method is clearly preferable since it is better to compute $k-1$ times the HNF of an $n \times 2n$ matrix than the HNF of a single $n \times kn$ matrix.

There is, however, another natural problem, which we will encounter below. As above, let $\mathcal{A}_1$ and $\mathcal{A}_2$ be subgroups of $\mathcal{B}$ given by HNF divisors H_1 and H_2. We would like to compute the intersection of these two subgroups

as a subgroup of $\mathcal{A}_2$, in other words as a left divisor of the SNF of $\mathcal{A}_2$ and not of $\mathcal{B}$. This is easily done using the following algorithm, whose easy proof is again left to the reader (Exercise 8).

Algorithm 4.1.15 (Intersection of Subgroups in a Subgroup). Given an Abelian group $\mathcal{B} = (B, D_B)$ in SNF and two subgroups $\mathcal{A}_1$ and $\mathcal{A}_2$ given by HNF left divisors H_1 and H_2 of D_B, this algorithm computes the SNF (A_2, D_{A_2}) of $\mathcal{A}_2$ and the left HNF divisor of D_{A_2}, giving $\mathcal{A}_1 \cap \mathcal{A}_2$ as a subgroup of $\mathcal{A}_2$.

1. [Compute SNF of $\mathcal{A}_2$] Using Algorithm 4.1.3 (except that one should skip step 1, which is not necessary) applied to the system of generators and relations $(BH_2, H_2^{-1}D_B)$, compute the SNF (A_2, D_{A_2}) of the group $\mathcal{A}_2$ and the matrix U_a, and output this SNF.
2. [Compute HNF of intersection] Let $(H_1|H_2)\begin{pmatrix} U_1 & U_2 \\ U_3 & U_4 \end{pmatrix} = (0|H)$ be the HNF decomposition of the matrix $(H_1|H_2)$.
3. [Terminate] Output the HNF of the matrix $(U_a U_3 | D_{A_2})$ as left HNF divisor of D_{A_2} representing $\mathcal{A}_1 \cap \mathcal{A}_2$, and terminate the algorithm.

4.1.9 p-Sylow Subgroups of Finite Abelian Groups

Let $\mathcal{C} = (C, D_C)$ be a group, and let p be a prime number. We would like to compute the p-Sylow subgroup $\mathcal{C}_p$ of $\mathcal{C}$. Recall that by definition, this is the subgroup of $\mathcal{C}$ consisting of all elements $g \in \mathcal{C}$ whose order is a power of p.

We will use the following convenient notation, which should be standard in number theory. If $m \in \mathbb{Z}$, $m \neq 0$, we will denote by (p^∞, m) the limit as $k \to \infty$ of (p^k, m). Of course, this sequence stabilizes for k large enough, so the limit exists; more precisely, $(p^\infty, m) = p^{v_p(m)}$, where as usual $v_p(m)$ is the p-adic valuation of m. The following proposition gives the answer to our question.

Proposition 4.1.16. *Let $\mathcal{C} = (C, D_C)$ be a group given in SNF, with $C = (\gamma_i)_{1 \le i \le n}$ and $D_C = \mathrm{diag}((c_i)_{1 \le i \le n})$, and let p be a prime number. Let r_c be the largest index $i \le n$ such that $p \mid c_i$ ($r_c = 0$ if none exist). Then $\mathcal{C}_p$ is given in SNF by $\mathcal{C}_p = (C_p, D_{C,p})$, where*

$$C_p = (\gamma_i^{c_i/(p^\infty, c_i)})_{1 \le i \le r_c} \quad \text{and} \quad D_{C,p} = \mathrm{diag}((p^\infty, c_i)_{1 \le i \le r_c}) \ .$$

Proof. Let $g \in \mathcal{C}_p$. There exists $a \ge 0$ such that $g^{p^a} = 1$. Let $g = \prod_{1 \le i \le n} \gamma_i^{x_i}$. Thus, $c_i \mid p^a x_i$, hence $(c_i/(p^a, c_i)) \mid x_i$, which implies that $(c_i/(p^\infty, c_i)) \mid x_i$. Hence, if we set $\gamma_{i,p} = \gamma_i^{c_i/(p^\infty, c_i)}$, the $\gamma_{i,p}$ are generators of $\mathcal{C}_p$, and we can restrict to $i \le r_c$, since otherwise the $\gamma_{i,p}$ are equal to 1. It is clear that the matrix of relations between the $\gamma_{i,p}$ is given by $D_{C,p} = \mathrm{diag}((p^\infty, c_i))$, and since this is already in SNF, this proves the proposition. □

Since this proposition gives explicitly the SNF of $\mathcal{C}_p$, it is not necessary to give a formal algorithm. If we want to consider $\mathcal{C}_p$ as a subgroup of $\mathcal{C}$,

the corresponding HNF left divisor of D_C is evidently the diagonal matrix $\mathrm{diag}((c_i/(p^\infty, c_i))_{1\le i\le n})$.

Consider now a more theoretical problem, which may be useful in certain cases. Assume that we have an exact sequence of Abelian groups. What happens when we take p-Sylow subgroups? The answer to this question is as follows. Taking p-Sylow subgroup is a *left exact* functor in the category of Abelian groups, and it is even an *exact* functor in the subcategory of finite Abelian groups. This means that we have the following proposition.

Proposition 4.1.17. (1) *Let* $1 \longrightarrow \mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C}$ *be an exact sequence of Abelian groups, which are exceptionally not assumed to be finite. Then* $1 \longrightarrow \mathcal{A}_p \xrightarrow{\psi_p} \mathcal{B}_p \xrightarrow{\phi_p} \mathcal{C}_p$ *is also an exact sequence, where the maps are simply the restrictions of the corresponding maps.*

(2) *Let* $\cdots \longrightarrow \mathcal{A} \longrightarrow \mathcal{B} \longrightarrow \mathcal{C} \longrightarrow \cdots$ *be an exact sequence of* finite *Abelian groups of any length. Then* $\cdots \longrightarrow \mathcal{A}_p \longrightarrow \mathcal{B}_p \longrightarrow \mathcal{C}_p \longrightarrow \cdots$ *is again an exact sequence.*

Proof. For (1), we first note that if ψ is a group homomorphism from $\mathcal{A}$ to $\mathcal{B}$, then clearly $\psi(\mathcal{A}_p) \subset \mathcal{B}_p$, so the restricted maps are well-defined. Exactness at $\mathcal{A}$ is also clear since the restriction of an injective map is injective. In addition, the identity $\phi \circ \psi = 0$ is preserved by restriction. Thus, we must simply show that $\mathrm{Ker}(\phi_p) \subset \mathrm{Im}(\psi_p)$. Let $x \in \mathrm{Ker}(\phi_p)$. This means first that $\phi(x) = 1$ in $\mathcal{C}$, hence by the exactness of the initial sequence, that $x = \psi(y)$ for some $y \in \mathcal{A}$. It also means that $x^{p^a} = 1$ for some $a \ge 0$. But then $\psi(y^{p^a}) = 1$, hence $y^{p^a} = 1$ since ψ is injective, and so $y \in \mathcal{A}_p$, and $x \in \mathrm{Im}(\psi_p)$ as desired.

Since any exact sequence is made up of short exact sequences of the type $1 \longrightarrow \mathcal{A} \xrightarrow{\psi} \mathcal{B} \xrightarrow{\phi} \mathcal{C} \longrightarrow 1$, it is enough to prove (2) for short exact sequences of this type, the general result following by induction. By (1), we already know that the sequence of p-Sylow subgroups is exact at $\mathcal{A}_p$ and at $\mathcal{B}_p$. We must show that it is exact at $\mathcal{C}_p$ or, equivalently, that ϕ_p is surjective. For this, we use in an essential way the fact that the groups are finite by using a counting argument. Let $|\mathcal{A}| = p^a k$, $|\mathcal{B}| = p^b m$, and $|\mathcal{C}| = p^c n$, where $p \nmid kmn$. By the exactness of the initial exact sequence, we have $p^b m = p^a k p^c n$, hence in particular $b = a + c$. By the structure theorem for finite Abelian groups, we have $|\mathcal{A}_p| = p^a$, $|\mathcal{B}_p| = p^b$, and $|\mathcal{C}_p| = p^c$, and hence $|\mathcal{C}_p| = |\mathcal{B}_p| \,/\, |\mathcal{A}_p|$. But since we already know exactness at $\mathcal{A}_p$ and $\mathcal{B}_p$, we have $|\phi_p(\mathcal{B}_p)| = |\mathcal{B}_p| \,/\, |\mathcal{A}_p|$, hence $|\phi_p(\mathcal{B}_p)| = |\mathcal{C}_p|$, thus showing that ϕ_p is surjective, as claimed. □

Remark. It is easy to see that (2) is false when the groups are not necessarily assumed to be finite. For example, consider the following exact sequence

$$0 \longrightarrow \mathbb{Z} \xrightarrow{[p]} \mathbb{Z} \xrightarrow{s} \mathbb{Z}/p\mathbb{Z} \longrightarrow 0 \ ,$$

where $[p]$ denotes multiplication by p and s is the canonical surjection. The sequence of p-Sylow subgroups is $0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \mathbb{Z}/p\mathbb{Z} \longrightarrow 0$, which clearly is not exact.

4.1.10 Enumeration of Subgroups

In Chapter 3, we saw that Abelian extensions correspond to equivalence classes of congruence subgroups or, equivalently, to equivalence classes of subgroups of ray class groups $Cl_{\mathfrak{m}}(K)$, which are finite Abelian groups. Thus, it is important to be able to *enumerate* these groups. We consider this problem here.

Let $\mathcal{C} = (C, D_C)$ be a fixed Abelian group given in SNF. By Proposition 4.1.6, enumerating subgroups of $\mathcal{C}$ is equivalent to enumerating HNF matrices H that are left divisors of D_C (two matrices M and M' that differ only by right multiplication by a unimodular matrix U defining the same subgroup).

The question of finding these divisors in an efficient manner is not immediate (see Theorem 4.1.18 below). In the context of class groups, however, it is reasonable to assume that $\mathcal{C}$ will often be cyclic or close to cyclic (see [Coh0, Section 5.10]). Hence, we can proceed as follows. Let n be the number of cyclic components of $\mathcal{C}$ as above, and let $D_C = \mathrm{diag}(c_1, \ldots, c_n)$.

If $n = 1$, H divides D_C if and only if $H = (e_1)$, where $e_1 \mid c_1$ and $e_1 \geq 1$; hence, we simply look at all (positive) divisors of c_1.

If $n = 2$, then an immediate computation shows that $H = \left(\begin{smallmatrix} e_1 & f_1 \\ 0 & e_2 \end{smallmatrix}\right)$ divides D_C if and only if for $i = 1$ and $i = 2$, e_i is a positive divisor of c_i, and $f_1 = ke_1/\gcd(e_1, c_2/e_2)$ with $0 \leq k < \gcd(e_1, c_2/e_2)$ (see Exercise 9).

If $n \geq 3$, we can try all possible HNF matrices $H = (e_{i,j})$ with $e_{i,i} \mid c_i$ and

$$e_{i,i+1} \equiv 0 \pmod{e_{i,i}/\gcd(e_{i,i}, c_{i+1}/e_{i+1,i+1})} \ ,$$

which are easily seen to be necessary conditions (see Exercise 9).

However, this is wasteful, since we need to examine many more HNF matrices than there are subgroups. Thus, we need a method that enables us to construct the HNF matrices that correspond to subgroups only (in other words, the left divisors of D_C).

We first make an important reduction. As in the preceding section, let $\mathcal{C}_p$ be the p-Sylow subgroup of $\mathcal{C}$, generated by the $\gamma_i^{c_i/(p^\infty, c_i)}$. I claim that it suffices to enumerate the subgroups of $\mathcal{C}_p$ for each p. Indeed, let $\mathcal{B}$ be a subgroup of $\mathcal{C}$ given by an HNF matrix H. The p-Sylow subgroup $\mathcal{B}_p$ of $\mathcal{B}$ is, of course, equal to $\mathcal{B} \cap \mathcal{C}_p$. We want to consider it as a subgroup of $\mathcal{C}_p$, and as such it can be computed as a left divisor H_p of the matrix $D_{C,p}$ given by Proposition 4.1.16, by using Algorithm 4.1.15 (note that in that algorithm U_a is the identity matrix in our case). Conversely, if for each p we are given a left divisor of H_p of $D_{C,p}$ corresponding to a subgroup $\mathcal{B}_p$ of $\mathcal{C}_p$, then it is clear that $D_C D_{C,p}^{-1} H_p$ is a left divisor of D_C corresponding to $\mathcal{B}_p$ considered as a

subgroup of $\mathcal{C}$, and hence we can reconstruct the subgroup $\mathcal{B}$ by summing these subgroups using Algorithm 4.1.14.

Although the above may sound like useless nitpicking, it is essential for a correct implementation.

Once this reduction is made, we may assume that our group $\mathcal{C}$ is a p-group, in other words that its order is a power of p. In this case, the complete answer to our problem has been given by G. Birkhoff (see [Bir], [But]).

I give the theorem as stated by L. Butler (slightly modified for our purposes) and refer to [Bir] and [But] for details and proof.

Theorem 4.1.18. *Let $\mathcal{C} = (C, D_C)$ be an Abelian p-group in SNF, and write $D_C = \mathrm{diag}((p^{x_i})_{1\le i\le s})$. Consider all the matrices M obtained as follows.*

(1) *We choose an integer t such that $0 \le t \le s$ and a family of integers $(y_i)_{1\le i\le t}$ such that $y_{i+1} \le y_i$ for $i < t$ and such that $y_i \le x_i$. We set by convention $y_i = 0$ for $t < i \le s$.*

(2) *We choose a permutation σ of $[1,s]$ such that for all $i \le t$, $y_i \le x_{\sigma(i)}$, and for all $i < s$ such that $y_i = y_{i+1}$, then $\sigma(i) > \sigma(i+1)$. Set $\tau = \sigma^{-1}$.*

(3) *We choose integers $c_{i,j}$ for $\tau(i) > j$, $1 \le i \le s$, $1 \le j \le t$, satisfying the following:*

$$\begin{aligned}
&(a)\quad i < \sigma(j) && \Longrightarrow\ 1 \le c_{i,j} \le p^{y_j - y_{\tau(i)}}\ ;\\
&(b)\quad i > \sigma(j) \text{ and } x_i < y_j && \Longrightarrow\ 1 \le c_{i,j} \le p^{x_i - y_{\tau(i)}}\ ;\\
&(c)\quad i > \sigma(j) \text{ and } x_i \ge y_j && \Longrightarrow\ 1 \le c_{i,j} \le p^{y_j - y_{\tau(i)} - 1}\ .
\end{aligned}$$

(4) *We define the $s \times t$ matrix $M = (m_{i,j})$ by setting*

$$m_{i,j} = \begin{cases} p^{x_i} & \text{if } \tau(i) < j \\ p^{x_i - y_j} & \text{if } \tau(i) = j \\ c_{i,j}p^{x_i - y_j} & \text{if } \tau(i) > j \text{ in case } (a) \\ c_{i,j} & \text{if } \tau(i) > j \text{ in case } (b) \\ c_{i,j}p^{x_i - y_j + 1} & \text{if } \tau(i) > j \text{ in case } (c). \end{cases}$$

To each subgroup $\mathcal{A}$ of $\mathcal{C}$ is associated a unique such matrix M, where the SNF of $\mathcal{A}$ is $(CM, \mathrm{diag}((p^{y_i})_{1\le i\le t}))$ and conversely each such matrix M gives rise to a subgroup whose corresponding left HNF divisor of D_C is the HNF of the matrix $(M|D_C)$.

Using this theorem and the algorithmic reductions to p-groups that we have made above, we can now easily write a complete algorithm for the enumeration of subgroups of a finite Abelian group, but we will not do this formally (see Exercise 10). Note that it may be more efficient to first choose the permutation and then the y_i.

Let us give an example. Assume that we want to describe all subgroups of $\mathcal{C} = (\mathbb{Z}/p^2\mathbb{Z}) \times (\mathbb{Z}/p\mathbb{Z})$ using Birkhoff's theorem. We find the following matrices M and the corresponding left HNF divisors obtained as the HNF of $(M|D_C)$, where $D_C = \mathrm{diag}(p^2, p)$.

(1) For $t = 0$, M is the 2×0 matrix and

$$H = D_C = \begin{pmatrix} p^2 & 0 \\ 0 & p \end{pmatrix} .$$

(2) For $t = 1$, $y_1 = 1$, and σ the identity,

$$M = \begin{pmatrix} p \\ p \end{pmatrix}, \quad H = \begin{pmatrix} p & 0 \\ 0 & p \end{pmatrix} .$$

(3) For $t = 1$, $y_1 = 1$, and σ the transposition,

$$M = \begin{pmatrix} c_{1,1}p \\ p \end{pmatrix}, \quad H = \begin{pmatrix} p^2 & c_{1,1}p \\ 0 & 1 \end{pmatrix} ,$$

where $c_{1,1}$ takes every integer value such that $1 \leq c_{1,1} \leq p$.

(4) For $t = 1$, $y_1 = 2$, σ is necessarily the identity,

$$M = \begin{pmatrix} 1 \\ c_{2,1} \end{pmatrix}, \quad H = \begin{pmatrix} p & c_{2,1}^{-1} \\ 0 & 1 \end{pmatrix}$$

if $p \nmid c_{2,1}$,

$$H = \begin{pmatrix} 1 & 0 \\ 0 & p \end{pmatrix}$$

if $p \mid c_{2,1}$. Here, $c_{2,1}$ takes every integer value such that $1 \leq c_{2,1} \leq p$, and when $p \nmid c_{2,1}$, $c_{2,1}^{-1}$ is the inverse of $c_{2,1}$ modulo p such that $1 \leq c_{2,1}^{-1} < p$.

(5) For $t = 2$, $y_1 = y_2 = 1$, σ is necessarily the transposition,

$$M = \begin{pmatrix} p & p \\ 1 & p \end{pmatrix}, \quad H = \begin{pmatrix} p & 0 \\ 0 & 1 \end{pmatrix} .$$

(6) For $t = 2$, $y_1 = 2$, $y_2 = 1$, σ is necessarily the identity,

$$M = \begin{pmatrix} 1 & p^2 \\ 1 & 1 \end{pmatrix}, \quad H = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} .$$

This gives a total of 1 subgroup of order 1, $p + 1$ subgroups of order p, $p + 1$ subgroups of order p^2, and 1 subgroup of order p^3.

The situation simplifies considerably if we want to enumerate not all subgroups, but only subgroups of given *index*. (By class field theory, in the ray class field case, computing congruence subgroups of given index is equivalent to computing Abelian extensions of given degree.) In particular, if the degree is prime, we have the following proposition.

Proposition 4.1.19. *Let $C = (C, D_C)$ be an Abelian group given in SNF, with $D_C = \mathrm{diag}((c_i)_i)$, and let ℓ be a prime number. Let r_c be the largest index i such that $\ell \mid c_i$ ($r_c = 0$ if none exist), so that r_c is the ℓ-rank of C.*

The subgroups of C of index ℓ correspond under Proposition 4.1.6 to matrices $H = (h_{i,j})$ such that there exists a row index k (necessarily unique) satisfying the following properties.

(1) *We have* $k \le r_c$.
(2) *For* $i \neq k$, *then* $h_{i,j} = 0$ *for* $j \neq i$ *and* $h_{i,i} = 1$.
(3) *We have* $h_{k,k} = \ell$, $h_{k,j} = 0$ *if* $j < k$ *or* $j > r_c$, *and* $0 \le h_{k,j} < \ell$ *if* $k < j \le r_c$.

In particular, there are $(\ell^{r_c}-1)/(\ell-1)$ *subgroups of index* ℓ (*this is, of course, a well-known and easy result*).

Proof. The proof of this proposition is easy and is left to the reader (Exercise 11). □

This proposition leads to the following algorithm.

Algorithm 4.1.20 (Subgroups of Index ℓ). Given an Abelian group $\mathcal{C} = (C, D_C)$ in HNF with $D_C = \mathrm{diag}((c_i)_{1\le i\le n})$ and a prime number ℓ, this algorithm computes the list $\mathcal{C}$ of all subgroups of $\mathcal{C}$ of index ℓ as HNF left divisors of D_C.

1. [Initializations] Let $\mathcal{C} \leftarrow \varnothing$, let r_c be the largest index i (0 if none exist) such that $\ell \mid c_i$, and set $k \leftarrow 0$.
2. [Loop on k] Set $k \leftarrow k+1$. If $k > r_c$, output $\mathcal{C}$ and terminate the algorithm. Otherwise, set $A \leftarrow -1$.
3. [Loop on A] Let $A \leftarrow A+1$. If $A \ge \ell^{r_c-k-1}$, go to step 2. Otherwise, let $H \leftarrow I_n$ be the identity matrix of order n, set $H_{k,k} \leftarrow \ell$, set $j \leftarrow k$, and $a \leftarrow A$.
4. [Loop on j] Set $j \leftarrow j+1$. If $j > r_c$, set $\mathcal{C} \leftarrow \mathcal{C} \cup \{H\}$ and go to step 3. Otherwise, let $a = \ell q + r$ be the Euclidean division of a by ℓ with $0 \le r < \ell$, set $H_{k,j} \leftarrow r$, $a \leftarrow q$, and go to step 4.

4.1.11 Application to the Solution of Linear Equations and Congruences

It is easy to apply the above techniques to the solution of a system of linear equations in integers or to a system of linear congruences. For the first problem, we can use the following algorithm (which should be in [Coh0]).

Algorithm 4.1.21 (Solving Linear Systems in Integers). Given an $m \times n$ matrix P with integer entries and an m-component integral column vector B, this algorithm either says that the system of linear equations $PX = B$ has no integral solution, or gives the general solution as a particular solution together with the general solution of the homogeneous system.

1. [Compute HNF] Using an HNF algorithm, compute a unimodular $n \times n$ matrix U and a (not necessarily square) HNF matrix H such that $PU = (0|H)$, let k be the number of columns equal to 0 in the right-hand side, and write $U = (U_1|U_2)$, where U_1 is an $n \times k$ and U_2 is an $(n-k) \times n$ matrix.

2. [Compute inverse image] Using [Coh0, Algorithm 2.3.4], check whether there exists an inverse image Z_2 of Y by H (if it exists it will be unique). If it does not exist or if it does not have integral entries, the system has no solution, so terminate the algorithm.
3. [Solve system] Output $X_0 \leftarrow U_2 Z_2$ as a particular solution of our linear system, the columns of the matrix U_1 as a $\mathbb{Z}$-basis of the homogeneous system, and terminate the algorithm.

Proof. The easy proof is left to the reader (Exercise 13). □

Consider now the similar problem with congruences. Let $P = (p_{i,j})$ be an $m \times n$ matrix with integer entries, let $(d_1, \ldots, d_m)$ be a set of positive integers, and let $B = (b_1, \ldots, b_m)^t$ be an integral column vector. We want to solve the system of m linear congruences in the n unknowns x_j

$$\sum_{1 \le j \le n} p_{i,j} x_j \equiv b_i \pmod{d_i} \quad \text{for} \quad 1 \le i \le m \ .$$

We must first give a meaning to the problem. Let

$$\mathcal{C} = \bigoplus_{1 \le i \le m} (\mathbb{Z}/d_i\mathbb{Z}) \ ,$$

and, if V is an integer column vector with m components, denote by $\overline{V}$ the image of V in $\mathcal{C}$ by the natural surjection from $\mathbb{Z}^m$ to $\mathcal{C}$. The matrix P defines a natural map from $\mathbb{Z}^n$ to $\mathcal{C}$ which sends X to $\overline{PX}$. Since $\mathcal{C}$ is a finite group, the kernel of this map is a lattice in $\mathbb{Z}^n$ that can therefore be represented as an HNF matrix H. If X_0 is a particular solution of our system (if it exists), the set of solutions to our system of congruences is then equal to $X_0 + HZ$ for any integer vector Z.

This solution is not completely satisfactory, however. Since we are dealing only with finite groups, we really want a finite solution set. Let d be the lowest common multiple (LCM) of the d_i, in other words the exponent of the group $\mathcal{C}$. Then clearly we can ask for solution vectors modulo d; in other words, we introduce $\mathcal{B} = (\mathbb{Z}/d\mathbb{Z})^n$ and consider P as a map from $\mathcal{B}$ to $\mathcal{C}$. The kernel of this map can now be computed by Algorithm 4.1.11 as a subgroup of $(\mathbb{Z}/d\mathbb{Z})^n$, and we can then compute its SNF in its own right, giving the solution set of the homogeneous system as a group (A, D_A), where the generators A are elements of $(\mathbb{Z}/d\mathbb{Z})^n$ and $D_A = \mathrm{diag}(a_1, \ldots, a_r)$ is a diagonal matrix in SNF such that $a_i \mid d$ for all i.

To transform this into an algorithm, let $D = \mathrm{diag}(d_1, \ldots, d_m)$ be the diagonal matrix of the d_i (which is not in SNF in general). Using the techniques of Section 4.1.6, we proceed as follows. Let U be an $(m+n) \times (m+n)$ unimodular matrix such that $(P|D)U = (0|H)$ with H in HNF, and write $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$. If $X = (x_1, \ldots, x_n)^t$ is a column vector representing a solution to our system of congruences, there exists another column vector

Y such that $(P|D)\binom{X}{Y} = B$ or, equivalently, $(P|D)U\binom{Z_1}{Z_2} = B$ if we set $\binom{Z_1}{Z_2} = U^{-1}\binom{X}{Y}$. Since $(P|D)U = (0|H)$, we obtain $HZ_2 = B$. Thus, our system has a solution if and only if $H^{-1}B$ is an integral vector. The general solution to our system is thus $\binom{Z_1}{Z_2} = \binom{Z_1}{H^{-1}B}$ for an arbitrary integral vector Z_1; hence $\binom{X}{Y} = U\binom{Z_1}{H^{-1}B}$, so

$$X = U_1 Z_1 + U_2 H^{-1} B$$

is the general solution in $\mathbb{Z}^n$ of our system of congruences.

The vector $U_2 H^{-1} B$ represents a particular solution to our system, while $U_1 Z_1$ is the general solution of the homogeneous system. To obtain the solution as a subgroup of $\mathcal{B} = (\mathbb{Z}/d\mathbb{Z})^n$, as in the final step of Algorithm 4.1.11 we compute the HNF H_B of $(U_1|dI_n)$.

Putting all this together gives the following algorithm.

Algorithm 4.1.22 (Linear System of Congruences). Let $\sum_{1\le j\le n} p_{i,j}x_j \equiv b_i \pmod{d_i}$ for $1 \le i \le m$ be a system of m linear congruences in the n unknowns x_j, and let d be the LCM of the d_i. This algorithm either says that the system has no solution, or gives the general solution as a particular solution together with the general solution of the homogeneous system, considered in $(\mathbb{Z}/d\mathbb{Z})^n$. We will denote by P the $m \times n$ matrix of the $p_{i,j}$, by D the diagonal matrix of the d_i, by B the column vector of the b_i, and represent solutions to our system by column vectors X with n components.

1. [Compute HNF of $(P|D)$] Apply an HNF algorithm to the matrix $(P|D)$, and let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(P|D)U = (0|H)$. We can discard the matrices U_3 and U_4.
2. [Test if solution] Let $Z_2 \leftarrow H^{-1}B$. If Z_2 is not an integral vector, the system has no solution, so terminate the algorithm. Otherwise, set $X_0 \leftarrow U_2 Z_2$ (this is a particular solution to our system).
3. [Terminate] Let H_B be the HNF of the matrix $(U_1|dI_n)$. Output H_B and terminate the algorithm (the general solution to our system in $\mathcal{B} = (\mathbb{Z}/d\mathbb{Z})^n$ will be $X_0 + H_B Z$ for an arbitrary vector $Z \in \mathcal{B}$).

Note in particular that the number of solutions of our system modulo d is either 0 (if $H^{-1}B$ is not integral) or equal to $d^n/\det(H_B)$.

If we want the solution of the homogeneous system as a group in its own right, we apply Algorithm 4.1.3 to the system of generators and relations (EH_B, dH_B^{-1}), where E is the canonical basis of $(\mathbb{Z}/d\mathbb{Z})^n$.

Finally, consider the problem of a combined system of linear congruences and linear equations. This simply corresponds to the choice of some d_i equal to 0 in the congruences. We cannot directly use Algorithm 4.1.22 since the matrix $D = \mathrm{diag}(d_i)$ is not of maximal rank. We call such a system a *mixed linear system*.

There are two ways to solve the problem. The first one is to start by solving the linear system using Algorithm 4.1.21, finding (if it exists) a particular solution plus the general solution of the homogeneous system. We then plug this into the system of congruences, giving a new system of congruences in new variables, which we can then solve using Algorithm 4.1.22. We leave the details to the reader (Exercise 14).

The second method is direct and gives the following algorithm, which is only a slight modification of Algorithm 4.1.22, and we leave its proof to the reader (Exercise 15).

Algorithm 4.1.23 (Mixed Linear System). Let $\sum_{1\le j\le n} p_{i,j}x_j = b_i$ for $1 \le i \le m_1$ and $\sum_{1\le j\le n} p_{i,j}x_j \equiv b_i \pmod{d_i}$ for $m_1 < i \le m = m_1 + m_2$ be a system of m_1 linear equations and m_2 linear congruences in the n unknowns x_i, and let d be the LCM of the d_i. This algorithm either says that the system has no solution, or gives the general solution as a particular solution together with the general solution of the homogeneous system. We will denote by P the $m \times n$ matrix of the $p_{i,j}$, by D the diagonal matrix of the d_i, by B the m-component column vector of the b_i, and represent solutions to our system by column vectors X with n components. We assume that there do exist linear equations, in other words, that there exists (i,j) with $1 \le i \le m_1$ and $1 \le j \le n$ such that $p_{i,j} \neq 0$ (otherwise, use Algorithm 4.1.22).

1. [Compute HNF of $(P|D_0)$] Let $D_0 \leftarrow \binom{0}{D}$ be the $m \times m_2$ matrix obtained by vertically concatenating an $m_1 \times m_2$ zero matrix with the diagonal matrix D. Apply an HNF algorithm to the matrix $(P|D_0)$, and let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(P|D_0)U = (0|H)$. We can discard the matrices U_3 and U_4. Note that the matrix H will not necessarily be square, but its columns are independent.
2. [Test if solution] Using [Coh0, Algorithm 2.3.4], check whether there exists an inverse image Z_2 of B by H (if it exists, it will be unique). If it does not exist or if Z_2 is not an integral vector, the system has no solution, so terminate the algorithm. Otherwise, set $X_0 \leftarrow U_2Z_2$ (this is a particular solution to our system).
3. [Compute integer kernel] Let $P_1 = (p_{i,j})_{1\le i\le m_1, 1\le j\le n}$ be the matrix of the linear system (obtained by extracting the first m_1 rows of P). Using an integer kernel algorithm (for example, [Coh0, Algorithm 2.4.10]), compute a matrix J whose columns give a $\mathbb{Z}$-basis for the integer kernel of P_1.
4. [Terminate] Let H_B be the HNF of the matrix $(U_1|dJ)$. Output H_B and terminate the algorithm (the general solution to our system in $\mathbb{Z}^n$ will be $X_0 + H_BZ$ for an arbitrary integer vector Z).

4.2 Computing the Structure of $(\mathbb{Z}_K/\mathfrak{m})^*$

Let K be a number field and $\mathfrak{m}$ a modulus of K. In this section, we explain how to compute the group $(\mathbb{Z}_K/\mathfrak{m})^*$ in the sense of Definition 4.1.4, using

the tools we developed in the preceding section. The theoretical answer to this question is solved, in principle, in [Nak2]. This is, however, not suited to algorithmic purposes and, in addition, is much more complicated than the solution we present below.

We give two answers to this question. The first answer gives a theoretical and practical answer valid in many, but not all, cases (Section 4.2.2). The second answer is slightly more complex, but gives a complete algorithmic answer to the problem (Section 4.2.5). These solutions are complementary.

4.2.1 Standard Reductions of the Problem

Let $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$ with $\mathfrak{m}_0 = \prod_{\mathfrak{p}} \mathfrak{p}^{v_\mathfrak{p}}$. An element of $(\mathbb{Z}_K/\mathfrak{m})^*$ will be written as a pair $(\overline{\alpha}, w)$, where $\alpha \in \mathbb{Z}_K$ is coprime to $\mathfrak{m}_0$, and $w \in (\mathbb{Z}/2\mathbb{Z})^{\mathfrak{m}_\infty}$. Note that, although the natural map from the elements of $\mathbb{Z}_K$ coprime to $\mathfrak{m}$ into $(\mathbb{Z}_K/\mathfrak{m})^*$ is *surjective* (as a consequence of the strong approximation theorem), it is *not* a good idea to represent elements of $(\mathbb{Z}_K/\mathfrak{m})^*$ as $(\overline{\alpha}, s(\alpha))$, where $s(\alpha) \in (\mathbb{Z}/2\mathbb{Z})^{\mathfrak{m}_\infty}$ is the vector of signs of α at all the places of $\mathfrak{m}_\infty$. The main reason for this will be seen in Section 4.3.2.

For each $\sigma \in \mathfrak{m}_\infty$, let e_σ denote the corresponding canonical basis element of $(\mathbb{Z}/2\mathbb{Z})^{\mathfrak{m}_\infty}$ (all its coordinates are equal to 0 except at σ, where the coordinate is equal to 1). By definition, we have

$$(\mathbb{Z}_K/\mathfrak{m})^* = (\mathbb{Z}_K/\mathfrak{m}_0)^* \oplus \bigoplus_{\sigma \in \mathfrak{m}_\infty} (\mathbb{Z}/2\mathbb{Z})e_\sigma \ ,$$

so we are reduced to computing $(\mathbb{Z}_K/\mathfrak{m}_0)^*$. The reader may wonder why I go to such pains in writing what is, after all, a trivial isomorphism, but I recall that isomorphisms must be proscribed in algorithmic practice.

We now have a similar problem. We know theoretically that

$$(\mathbb{Z}_K/\mathfrak{m}_0)^* \simeq \prod_{\mathfrak{p}} (\mathbb{Z}_K/\mathfrak{p}^{v_\mathfrak{p}})^* \ ,$$

but this is not usable in algorithmic practice, since we must absolutely have an equality and not an isomorphism. This is obtained by using the following lemma.

Lemma 4.2.1. *Let $\mathfrak{a}$ and $\mathfrak{c}$ be two coprime integral ideals of K, and set $\mathfrak{b} = \mathfrak{a}\mathfrak{c}$.*

(1) *We can find in polynomial time elements a and c such that $a \in \mathfrak{a}$, $c \in \mathfrak{c}$, and $a + c = 1$.*
(2) *We have a split exact sequence*

$$1 \longrightarrow (\mathbb{Z}_K/\mathfrak{a})^* \xrightarrow{\psi} (\mathbb{Z}_K/\mathfrak{b})^* \xrightarrow{\phi} (\mathbb{Z}_K/\mathfrak{c})^* \longrightarrow 1 \ ,$$

where $\psi(\overline{\alpha}) = \overline{c\alpha + a}$, $\phi(\overline{\beta}) = \overline{\beta}$, *and a section* σ *of* ϕ *is given by* $\sigma(\overline{\gamma}) = \overline{a\gamma + c}$. (*Here* $\overline{\ }$ *denotes the classes in the respective groups, but using the same notation for each will not lead to any confusion as long as we know in which group we work.*)

(3) *Assume that* $(\mathbb{Z}_K/\mathfrak{a})^* = \bigoplus(\mathbb{Z}/a_i\mathbb{Z})\overline{\alpha_i}$ *and* $(\mathbb{Z}_K/\mathfrak{c})^* = \bigoplus(\mathbb{Z}/c_j\mathbb{Z})\overline{\gamma_j}$. *Then*

$$(\mathbb{Z}_K/\mathfrak{b})^* = \bigoplus(\mathbb{Z}/a_i\mathbb{Z})\overline{(c\alpha_i + a)} \oplus \bigoplus(\mathbb{Z}/c_j\mathbb{Z})\overline{(a\gamma_j + c)} \ .$$

(*Note that this is not quite a representation in SNF, but it can easily be transformed into one.*)

Proof. The proof is a little tedious but straightforward.

(1). This is a restatement of Proposition 1.3.1.

(2) a). The map ψ is well-defined: if $\overline{\alpha} = \overline{\alpha'}$, then $\alpha' - \alpha \in \mathfrak{a}$; hence, $(c\alpha' + a) - (c\alpha + a) = c(\alpha' - \alpha) \in \mathfrak{ac} = \mathfrak{b}$ since $c \in \mathfrak{c}$.

b). The map ψ is a group homomorphism. Indeed, this follows from the fact that a and c are orthogonal idempotents modulo $\mathfrak{b}$; in other words, that $ac \in \mathfrak{b}$ and

$$a^2 - a = -a(1-a) = -ac = -c(1-c) = c^2 - c \in \mathfrak{b} \ .$$

Hence,

$$\psi(\overline{\alpha})\psi(\overline{\alpha'}) = \overline{(c\alpha + a)(c\alpha' + a)} = \overline{c\alpha\alpha' + a} = \psi(\overline{\alpha\alpha'}) \ .$$

c). The map ψ is injective. Indeed,

$$\begin{aligned} \psi(\overline{\alpha}) = \overline{1} &\iff \overline{c\alpha + a} \equiv 1 \pmod{\mathfrak{ac}} \implies c\alpha \equiv 1 \pmod{\mathfrak{a}} \\ &\implies \alpha \equiv 1 \pmod{\mathfrak{a}} \qquad \iff \overline{\alpha} = 1 \end{aligned}$$

since $c \equiv 1 \pmod{\mathfrak{a}}$.

d). The map ϕ is clearly well-defined and is a group homomorphism.

e). By symmetry with a), b), and c), the map σ is well-defined and is an injective group homomorphism. Furthermore, since $a \equiv 1 \pmod{\mathfrak{c}}$, $\phi \circ \sigma$ is the identity map, which implies that σ is a section of ϕ and in particular that ϕ is surjective.

Statement (3) is an immediate consequence of (2). □

Remarks

(1) This lemma can easily be generalized to the case where $\mathfrak{a}$ and $\mathfrak{c}$ are coprime moduli (meaning that $\mathfrak{a}_0 + \mathfrak{c}_0 = \mathbb{Z}_K$ and $\mathfrak{a}_\infty \cap \mathfrak{c}_\infty = \varnothing$). The proof of this is left to the reader (see Exercise 16).

(2) This lemma is simply the Chinese remainder theorem for ideals (or, more generally, for moduli).

(3) In Section 4.1.4 we mentioned that the matrix P introduced there measures the obstruction to an exact sequence being split. Let $\mathcal{A} = (\mathbb{Z}_K/\mathfrak{a})^* = (A, D_A)$, $\mathcal{B} = (\mathbb{Z}_K/\mathfrak{b})^*$, and $\mathcal{C} = (\mathbb{Z}_K/\mathfrak{c})^* = (C, D_C)$. If we follow Algorithm 4.1.8, we must first choose lifts B' of C. Since our sequence is split, we will take $B' = \sigma(\overline{C}) = \overline{aC + c\mathbf{1}_{\mathbb{Z}_K}}$. Since σ is a homomorphism, we have $B'D_C = \sigma(\overline{CD_C}) = \sigma(\mathbf{1}_C) = \mathbf{1}_B$, hence $P = 0$, as claimed (see also Exercise 2).

By induction, it follows from this lemma that to compute the structure of $(\mathbb{Z}_K/\mathfrak{m})^*$ it is enough to compute the structure of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ for prime ideals $\mathfrak{p}$. Hence, we can proceed in one of two ways. Either use Lemma 4.2.1 (3) recursively or, preferably, we can use the following more global algorithm.

Algorithm 4.2.2 (Nonrecursive Chinese for Ideals). Let $\mathfrak{m}_0 = \prod_{\mathfrak{p}} \mathfrak{p}^{v_{\mathfrak{p}}}$ be an integral ideal, and assume that we are given the SNF of $(\mathbb{Z}_K/\mathfrak{p}^{v_{\mathfrak{p}}})^* = (G_{\mathfrak{p}}, D_{\mathfrak{p}})$. This algorithm computes the SNF of $(\mathbb{Z}_K/\mathfrak{m}_0)^*$.

1. [Compute $\alpha_{\mathfrak{p}}$ and $\beta_{\mathfrak{p}}$] For each $\mathfrak{p} \mid \mathfrak{m}_0$, do as follows. Using Algorithm 1.3.2, compute $\alpha_{\mathfrak{p}}$ and $\beta_{\mathfrak{p}}$ such that $\alpha_{\mathfrak{p}} \in \mathfrak{m}_0/\mathfrak{p}^{v_{\mathfrak{p}}}$, $\beta_{\mathfrak{p}} \in \mathfrak{p}^{v_{\mathfrak{p}}}$, and $\alpha_{\mathfrak{p}} + \beta_{\mathfrak{p}} = 1$.
2. [Terminate] Let G be the concatenation of the $\beta_{\mathfrak{p}}\mathbf{1}_{\mathbb{Z}_K} + \alpha_{\mathfrak{p}}G_{\mathfrak{p}}$ and let D be the diagonal concatenation of the SNF matrices $D_{\mathfrak{p}}$. Using Algorithm 4.1.3 on the system of generators and relations (G, D), output the SNF of the group $(\mathbb{Z}_K/\mathfrak{m}_0)^*$ and the auxiliary matrix U_a, and terminate the algorithm.

Proof. If $G_{\mathfrak{p}} = (\gamma_i)$, it is clear that if we set $\gamma_i' = \beta_{\mathfrak{p}} + \alpha_{\mathfrak{p}}\gamma_i$ then $\gamma_i' \equiv \gamma_i \pmod{\mathfrak{p}^{v_{\mathfrak{p}}}}$; hence the γ_i' are also generators of $(\mathbb{Z}_K/\mathfrak{p}^{v_{\mathfrak{p}}})^*$ with the same matrix of relations $D_{\mathfrak{p}}$. In particular, the γ_i' are coprime to $\mathfrak{p}^{v_{\mathfrak{p}}}$, but on the other hand $\gamma_i' \equiv 1 \pmod{\mathfrak{m}_0/\mathfrak{p}^{v_{\mathfrak{p}}}}$, so the γ_i' are also coprime to $\mathfrak{m}_0/\mathfrak{p}^{v_{\mathfrak{p}}}$, hence to $\mathfrak{m}_0$, so if we concatenate all the γ_i' we clearly obtain a generating system for $(\mathbb{Z}_K/\mathfrak{m}_0)^*$ whose matrix of relations is the diagonal concatenation of the $D_{\mathfrak{p}}$. □

Note that, as usual, the matrix U_a allows us to obtain a corresponding discrete logarithm algorithm.

We have thus reduced the problem to the computation of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$. For this, we first introduce a definition.

Definition and Proposition 4.2.3. *Let $\mathfrak{a}$ and $\mathfrak{b}$ be (nonzero) ideals. Assume that $\mathfrak{a} \mid \mathfrak{b} \mid \mathfrak{a}^k$ for some positive integer k. We denote by $(1+\mathfrak{a})/(1+\mathfrak{b})$ the quotient set of $1+\mathfrak{a}$ by the equivalence relation $\mathcal{R}$ defined by $(1+x)\ \mathcal{R}\ (1+y) \iff x \equiv y \pmod{\mathfrak{b}}$. Multiplication in K induces a multiplication in $(1+\mathfrak{a})/(1+\mathfrak{b})$, which makes this set into an Abelian group.*

Proof. It is clear that $\mathcal{R}$ is an equivalence relation. Since $\mathfrak{a}$ is an ideal, $1+\mathfrak{a}$ is stable by multiplication, and since $\mathfrak{b}$ is an ideal, $\mathcal{R}$ is compatible with multiplication. Thus $(1+\mathfrak{a})/(1+\mathfrak{b})$ has a natural commutative multiplication, and the class of 1 is the unit element. We need only to show that any element

has an inverse. But if $x \in \mathfrak{a}$, then by assumption $x^k \in \mathfrak{b}$. It follows that for any $x \in \mathfrak{a}$ we have

$$(1+x)\left(1+\sum_{i=1}^{k-1}(-1)^i x^i\right) = 1+(-1)^{k-1}x^k \ ;$$

hence, if we set $y = \sum_{i=1}^{k-1}(-1)^i x^i$, then $y \in \mathfrak{a}$ and $(1+x)(1+y)-1 \in \mathfrak{b}$, so the class of $1+y$ is an inverse of the class of $1+x$. Thus $(1+\mathfrak{a})/(1+\mathfrak{b})$ is in a natural way an Abelian group.

It is easy to prove that this group is also finite. This will, in fact, follow from the results proven in the rest of this section. □

Proposition 4.2.4. *Let $\mathfrak{p}$ be a prime ideal of degree f, and let $q = p^f = |\mathbb{Z}_K/\mathfrak{p}|$. Set $G = (\mathbb{Z}_K/\mathfrak{p}^k)^*$. Let*

$$W = \{x \in G/\ x^{q-1} = 1\} \quad \textit{and} \quad G_{\mathfrak{p}} = (1+\mathfrak{p})/(1+\mathfrak{p}^k) \ .$$

Then

(1) *$W \simeq (\mathbb{Z}_K/\mathfrak{p})^*$, and in particular W is a cyclic subgroup of order $q-1$ of G. More precisely, if $\overline{g_0}$ is a generator of $(\mathbb{Z}_K/\mathfrak{p})^*$, then $\lceil \log_2(k) \rceil$ iterations of $g \leftarrow g - (g^{q-1}-1)/((q-1)g^{q-2}) \bmod \mathfrak{p}^k$ applied to g_0 gives a generator of W.*

(2) *$G_{\mathfrak{p}}$ is a p-subgroup of G of order q^{k-1}.*

(3) *$G = W \times G_{\mathfrak{p}}$.*

Proof. (1). All nonzero elements of $\mathbb{Z}_K/\mathfrak{p}$ are roots of the polynomial equation $X^{q-1}-1 = 0$; hence this equation has exactly $q-1$ distinct solutions in the field $\mathbb{Z}_K/\mathfrak{p}$. Thus

$$X^{q-1}-1 \equiv \prod_{a \in (\mathbb{Z}_K/\mathfrak{p})^*} (X-a) \pmod{\mathfrak{p}} \ .$$

It follows from Hensel's lemma that this factorization can be lifted to a factorization modulo any power of $\mathfrak{p}$. Thus there exists a group isomorphism between $(\mathbb{Z}_K/\mathfrak{p})^*$ and solutions to $X^{q-1}-1 \equiv 0 \pmod{\mathfrak{p}^k}$; in other words, between $(\mathbb{Z}_K/\mathfrak{p})^*$ and W.

It follows that W is a cyclic group of order $q-1 = p^f-1$ and that a generator of W can be obtained by Hensel lifting a generator of $(\mathbb{Z}_K/\mathfrak{p})^*$. This is done using the Newton–Hensel iteration given in the proposition.

(2). If we send the class of $1+x$ to the class of x, it is clear that, *as a set*, $G_{\mathfrak{p}}$ is isomorphic to $\mathfrak{p}/\mathfrak{p}^k$. In fact, as we will see in more detail below, the whole difficulty of the structure problem for $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ comes from the fact that this is only a set isomorphism, and not always a group isomorphism.

In any case, it follows that

$$|G_{\mathfrak{p}}| = |\mathfrak{p}/\mathfrak{p}^k| = \mathcal{N}(\mathfrak{p}^k)/\mathcal{N}(\mathfrak{p}) = q^{k-1} \ ,$$

so $G_{\mathfrak{p}}$ is a p-subgroup of G of order $q^{k-1} = p^{f(k-1)}$.

(3). Consider the map ϕ from $W \times G_{\mathfrak{p}}$ to G defined by $\phi((x,y)) = x \cdot y$. It is clearly a group homomorphism, and it is an isomorphism since an element of W is characterized by its residue modulo $\mathfrak{p}$, and each nonzero residue is attained. □

It follows from this proposition that to compute the structure of $(\mathbb{Z}_K/\mathfrak{m})^*$ it is sufficient to compute the structure of $G_{\mathfrak{p}}$, which is of course the p-Sylow subgroup of G.

4.2.2 The Use of $\mathfrak{p}$-adic Logarithms

To compute $G_{\mathfrak{p}}$, a natural idea is the use of $\mathfrak{p}$-adic logarithms. This is indeed useful but cannot be applied in complete generality. We study it in detail here.

We first recall some basic notions about $\mathfrak{p}$-adic numbers. We refer to [Ami], [Bac], [Kob], and many other textbooks on the subject.

Definition 4.2.5. *Let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$. A* $\mathfrak{p}$-adic integer *is a sequence $(a_k)_{k\geq 0}$, where $a_k \in \mathbb{Z}_K/\mathfrak{p}^k$ is such that $a_{k+1} \equiv a_k \pmod{\mathfrak{p}^k}$. The set of $\mathfrak{p}$-adic integers is an integral domain denoted $\mathbb{Z}_{K,\mathfrak{p}}$, and its field of fractions, denoted $K_{\mathfrak{p}}$, is called the* $\mathfrak{p}$-adic completion *of the number field K.*

In practice, although we will always work with elements modulo some fixed power $\mathfrak{p}^k$ of $\mathfrak{p}$, it is much more convenient to consider this as the truncation at level k of a $\mathfrak{p}$-adic number.

Definition 4.2.6. *Let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$ and x an element of K. We define the* $\mathfrak{p}$-adic logarithm *of $1+x$ by the expansion*

$$\log_{\mathfrak{p}}(1+x) = \sum_{i=1}^{\infty} (-1)^{i-1} \frac{x^i}{i} \ .$$

We define the $\mathfrak{p}$-adic exponential *of x by the expansion*

$$\exp_{\mathfrak{p}}(x) = \sum_{i=0}^{\infty} \frac{x^i}{i!} \ .$$

The basic properties of these $\mathfrak{p}$-adic functions are as follows.

Proposition 4.2.7. *Let $\mathfrak{p}$ be a prime ideal above a prime number p, and let $e = e(\mathfrak{p}/p) = v_{\mathfrak{p}}(p)$ be its ramification index.*

(1) *The expansion for* $\log_{\mathfrak{p}}(1+x)$ *converges* $\mathfrak{p}$*-adically if and only if* $v_{\mathfrak{p}}(x) \geq 1$.
(2) *The expansion for* $\exp_{\mathfrak{p}}(x)$ *converges* $\mathfrak{p}$*-adically if and only if* $v_{\mathfrak{p}}(x) > e/(p-1)$ *or, equivalently, if and only if* $v_{\mathfrak{p}}(x) \geq 1 + \lfloor e/(p-1) \rfloor$.
(3) *We have*

$$\log_{\mathfrak{p}}((1+x)(1+y)) = \log_{\mathfrak{p}}(1+x) + \log_{\mathfrak{p}}(1+y)$$

whenever this makes sense — more precisely, whenever $v_{\mathfrak{p}}(x) \geq 1$ *and* $v_{\mathfrak{p}}(y) \geq 1$.
(4) *We have* $\exp_{\mathfrak{p}}(x+y) = \exp_{\mathfrak{p}}(x)\exp_{\mathfrak{p}}(y)$ *whenever this makes sense — more precisely, whenever* $v_{\mathfrak{p}}(x) > e/(p-1)$ *and* $v_{\mathfrak{p}}(y) > e/(p-1)$.
(5) *We have* $\log_{\mathfrak{p}}(\exp_{\mathfrak{p}}(x)) = x$ *and* $\exp_{\mathfrak{p}}(\log_{\mathfrak{p}}(1+x)) = 1+x$ *whenever* $v_{\mathfrak{p}}(x) > e/(p-1)$.

Proof. (1). It is easily shown that a series $\sum_i u_i$ converges $\mathfrak{p}$-adically if and only if u_i tends to zero $\mathfrak{p}$-adically; in other words, if and only if the $\mathfrak{p}$-adic valuation $v_{\mathfrak{p}}(u_i)$ tends to infinity as $i \to \infty$.

We have

$$v_{\mathfrak{p}}\left(\frac{x^i}{i}\right) = iv_{\mathfrak{p}}(x) - v_{\mathfrak{p}}(i) = iv_{\mathfrak{p}}(x) - ev_p(i) \ .$$

Thus, if $v_{\mathfrak{p}}(x) \geq 1$, we have $v_{\mathfrak{p}}(x^i/i) \geq i - ev_p(i) \geq i - e\log(i)/\log(p) \to \infty$ as $i \to \infty$; hence the series converges $\mathfrak{p}$-adically. On the other hand, if $v_{\mathfrak{p}}(x) \leq 0$, then $v_{\mathfrak{p}}(x^i/i) \leq -ev_p(i)$, which does not tend to $+\infty$ as $i \to \infty$.

(2). We have

$$v_p(i!) = \sum_{j \geq 1} \left\lfloor \frac{i}{p^j} \right\rfloor \ ,$$

hence $v_p(i!) < i/(p-1)$, so $v_p(i!) \leq (i-1)/(p-1)$ with equality if and only if i is a power of p (see Exercise 17).

Thus

$$v_{\mathfrak{p}}\left(\frac{x^i}{i!}\right) = iv_{\mathfrak{p}}(x) - ev_p(i!) \geq i\left(v_{\mathfrak{p}}(x) - \frac{e}{p-1}\right) \to \infty$$

as $i \to \infty$ when $v_{\mathfrak{p}}(x) > e/(p-1)$ or, equivalently, when $v_{\mathfrak{p}}(x) \geq 1 + \lfloor e/(p-1) \rfloor$.

Conversely, if $v_{\mathfrak{p}}(x) \leq e/(p-1)$, then when i is a power of p we have $v_{\mathfrak{p}}(x^i/i!) = iv_{\mathfrak{p}}(x) - e(i-1)/(p-1) \leq e/(p-1)$, and this does not tend to infinity as $i \to \infty$, so the series does not converge in this case.

(3), (4), (5). The identities themselves are purely formal and are equivalent to standard combinatorial identities on binomial coefficients, which in turn can be proved via the properties of the usual (complex) logarithm and exponential functions. We must, however, also find their domain of validity. For (3), the result is clear since the functions $\log_{\mathfrak{p}}(1+x)$ and $\log_{\mathfrak{p}}(1+y)$ must be defined; hence $v_{\mathfrak{p}}(x) \geq 1$ and $v_{\mathfrak{p}}(y) \geq 1$, but then $v_{\mathfrak{p}}(x+y+xy) \geq 1$ also. The proof of (4) is similar.

Let us prove (5). For $\log_{\mathfrak{p}}(\exp_{\mathfrak{p}}(x))$ to be defined, we must have at least $v_{\mathfrak{p}}(x) > e/(p-1)$. Conversely, assume that this is satisfied. I claim that $v_{\mathfrak{p}}(\exp_{\mathfrak{p}}(x)-1) \geq 1$, and so the logarithm will be defined.

Indeed, for all $i > 0$, we have

$$v_{\mathfrak{p}}\left(\frac{x^i}{i!}\right) = iv_{\mathfrak{p}}(x) - e\sum_{j\geq 1}\left\lfloor \frac{i}{p^j}\right\rfloor > iv_{\mathfrak{p}}(x) - \frac{ie}{p-1} > 0 \ ,$$

proving our claim.

Conversely, if $v_{\mathfrak{p}}(x) > e/(p-1)$, then for all i we have

$$v_{\mathfrak{p}}\left(\frac{x^i}{i}\right) - v_{\mathfrak{p}}(x) = (i-1)v_{\mathfrak{p}}(x) - ev_p(i) \ ,$$

so if $i = p^a m$ with $p \nmid m$, we have

$$v_{\mathfrak{p}}\left(\frac{x^i}{i}\right) - v_{\mathfrak{p}}(x) = (p^a m - 1)v_{\mathfrak{p}}(x) - ea \ .$$

If we set $f(a) = (p^a-1)v_{\mathfrak{p}}(x) - ea$, we have, for all $a \geq 1$,

$$f(a) - f(a-1) = p^{a-1}(p-1)v_{\mathfrak{p}}(x) - e \geq (p-1)v_{\mathfrak{p}}(x) - e > 0 \ ;$$

hence for all $a \geq 1$ we have $f(a) > f(0) = 0$. From this it follows that $v_{\mathfrak{p}}(x^i/i) - v_{\mathfrak{p}}(x) > 0$ for $a > 0$, and for $a = 0$, $v_{\mathfrak{p}}(x^i/i) - v_{\mathfrak{p}}(x) = (i-1)v_{\mathfrak{p}}(x) > 0$ when $i > 1$. So for $i > 1$, we have $v_{\mathfrak{p}}(x^i/i) > v_{\mathfrak{p}}(x)$; hence $v_{\mathfrak{p}}(\log_{\mathfrak{p}}(1+x)) = v_{\mathfrak{p}}(x)$. It follows that the exponential is defined, and the identity follows. □

Corollary 4.2.8. *Let $\mathfrak{p}$ be a prime ideal above a prime number p, let $e = e(\mathfrak{p}/p) = v_{\mathfrak{p}}(p)$ be its ramification index, and set $k_0 = 1 + \lfloor e/(p-1)\rfloor$. For any integers a and b such that $b > a \geq k_0$, the functions $\log_{\mathfrak{p}}$ and $\exp_{\mathfrak{p}}$ induce inverse isomorphisms between the multiplicative group $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ and the additive group $\mathfrak{p}^a/\mathfrak{p}^b$. In particular, if $e < p-1$ and $k \geq 2$, they induce inverse isomorphisms between $G_{\mathfrak{p}} = (1+\mathfrak{p})/(1+\mathfrak{p}^k)$ and $\mathfrak{p}/\mathfrak{p}^k$.*

Proof. Set $v = v_{\mathfrak{p}}(x)$. We have seen above that if $v > e/(p-1)$, then $v_{\mathfrak{p}}(\log_{\mathfrak{p}}(1+x)) = v_{\mathfrak{p}}(x) = v$. Hence $\log_{\mathfrak{p}}$ sends $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ to $\mathfrak{p}^a/\mathfrak{p}^b$, and it is a group homomorphism because of the additive property of the logarithm.

On the other hand, since $v > e/(p-1)$, the function $\exp_{\mathfrak{p}}(x)$ converges for $x \in \mathfrak{p}^a$. Furthermore, since $v_p(i!) \leq (i-1)/(p-1)$, when $v = v_{\mathfrak{p}}(x) \geq k_0$ we have

$$v_{\mathfrak{p}}\left(\frac{x^i}{i!}\right) = iv - ev_p(i!) \geq v + (i-1)\left(v - \frac{e}{p-1}\right) \ .$$

Therefore, if $i > 1$, we have $v_\mathfrak{p}(x^i/i!) > v$; hence $v_\mathfrak{p}(\exp_\mathfrak{p}(x) - 1) = v$. It follows that $\exp_\mathfrak{p}$ sends $\mathfrak{p}^a/\mathfrak{p}^b$ to $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$, and it is the inverse map of $\log_\mathfrak{p}$ by the proposition, proving the corollary. □

If $\mathfrak{p}$ is an unramified prime ideal above a prime $p \geq 3$, then $k_0 = 1$, and so we have an explicit isomorphism $G_\mathfrak{p} \simeq \mathfrak{p}/\mathfrak{p}^k$. Hence, apart from a small finite number of prime ideals, this corollary reduces a relatively difficult multiplicative problem to a much easier additive one, as we will now see.

We first have the following easy lemma.

Lemma 4.2.9. *Let $\mathfrak{a}$ and $\mathfrak{b}$ be (nonzero) integral ideals of $\mathbb{Z}_K$. The additive group $\mathfrak{b}/\mathfrak{a}\mathfrak{b}$ is isomorphic to the additive group $\mathbb{Z}_K/\mathfrak{a}$.*

Proof. By the approximation theorem for Dedekind domains, there exists $\alpha \in \mathbb{Z}_K$ such that $v_\mathfrak{p}(\alpha) = v_\mathfrak{p}(\mathfrak{b})$ for all $\mathfrak{p}$ dividing $\mathfrak{a}$ and $v_\mathfrak{p}(\alpha) \geq v_\mathfrak{p}(\mathfrak{b})$ for all $\mathfrak{p}$ dividing $\mathfrak{b}$. In particular, $\alpha \in \mathfrak{b}$. Thus the map $x \mapsto \alpha x$ induces a well-defined additive group homomorphism from $\mathbb{Z}_K/\mathfrak{a}$ to $\mathfrak{b}/\mathfrak{a}\mathfrak{b}$. Since

$$\overline{\alpha x} = \overline{0} \iff \alpha x \in \mathfrak{a}\mathfrak{b} \iff \forall \mathfrak{p} \quad v_\mathfrak{p}(\alpha) + v_\mathfrak{p}(x) \geq v_\mathfrak{p}(\mathfrak{a}) + v_\mathfrak{p}(\mathfrak{b}) \ ,$$

it follows from our choice of α that, for all $\mathfrak{p}$ dividing $\mathfrak{a}$, we have $v_\mathfrak{p}(x) \geq v_\mathfrak{p}(\mathfrak{a})$, and hence $x \in \mathfrak{a}$ so $\overline{x} = \overline{0}$. Thus our map is an injective group homomorphism. Since the norm is multiplicative in $\mathbb{Z}_K$, we have

$$|\mathfrak{b}/\mathfrak{a}\mathfrak{b}| = \mathcal{N}(\mathfrak{a}\mathfrak{b})/\mathcal{N}(\mathfrak{b}) = \mathcal{N}(\mathfrak{a}) = |\mathbb{Z}_K/\mathfrak{a}| \ ,$$

and hence our map is also surjective, proving the lemma. □

Coming back to our original problem, by Corollary 4.2.8 we know that if $b > a > e/(p-1)$, the multiplicative group $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ is isomorphic to $\mathfrak{p}^a/\mathfrak{p}^b$ and hence, by the above lemma, to $\mathbb{Z}_K/\mathfrak{p}^{b-a}$.

The structure of these additive groups can be completely described as follows.

Theorem 4.2.10. *Let $\mathfrak{p}$ a prime ideal above p, with ramification index $e = e(\mathfrak{p}/p)$ and residual degree $f = f(\mathfrak{p}/p)$, and let $k \geq 1$ be an integer. Write*

$$k + e - 1 = eq + r \quad \textit{with} \quad 0 \leq r < e \ .$$

Then

$$\left(\mathbb{Z}_K/\mathfrak{p}^k\right) \simeq \left(\mathbb{Z}/p^q\mathbb{Z}\right)^{(r+1)f} \times \left(\mathbb{Z}/p^{q-1}\mathbb{Z}\right)^{(e-r-1)f} \ .$$

Proof. We have $\left|\mathbb{Z}_K/\mathfrak{p}^k\right| = \mathcal{N}(\mathfrak{p}^k) = p^{kf}$, hence $\mathbb{Z}_K/\mathfrak{p}^k$ is a p-group of cardinality p^{kf}, so we can write

$$\mathbb{Z}_K/\mathfrak{p}^k \simeq \prod_{i \geq 1} (\mathbb{Z}/p^i\mathbb{Z})^{a_i}, \quad \text{with} \quad \sum_{i \geq 1} i a_i = kf \ .$$

Since $p^k \in \mathfrak{p}^k$, we must have $a_i = 0$ for $i > k$. Let us assume that we have computed $a_k, a_{k-1}, \ldots, a_{j+1}$ (initially with $j = k$). We want to compute a_j.

Note that

$$p^{j-1}(\mathbb{Z}_K/\mathfrak{p}^k) \simeq \prod_{i \geq j} (\mathbb{Z}/p^{i-j+1}\mathbb{Z})^{a_i} \ ;$$

hence

$$\left|p^{j-1}(\mathbb{Z}_K/\mathfrak{p}^k)\right| = p^s \quad \text{with} \quad s = \sum_{i \geq j} (i-j+1)a_i \ .$$

On the other hand, we have

$$p^{j-1}(\mathbb{Z}_K/\mathfrak{p}^k) = (p^{j-1}\mathbb{Z}_K + \mathfrak{p}^k)/\mathfrak{p}^k \ .$$

The ideal $\mathfrak{b} = p^{j-1}\mathbb{Z}_K + \mathfrak{p}^k$ is an integral ideal that contains $\mathfrak{p}^k$ and hence is a power of $\mathfrak{p}$. Furthermore,

$$v_{\mathfrak{p}}(\mathfrak{b}) = \min(v_{\mathfrak{p}}(p^{j-1}), v_{\mathfrak{p}}(\mathfrak{p}^k)) = \min(e(j-1), k) \ ,$$

so $\mathfrak{b} = \mathfrak{p}^{\min(e(j-1),k)}$.

Since the ideal norm is multiplicative and $(\mathbb{Z}_K/\mathfrak{p}^a)/(\mathfrak{p}^b/\mathfrak{p}^a) \simeq \mathbb{Z}_K/\mathfrak{p}^b$, we have $\left|\mathfrak{p}^b/\mathfrak{p}^a\right| = \mathcal{N}(\mathfrak{p})^{a-b}$, from which it finally follows that

$$\left|p^{j-1}(\mathbb{Z}_K/\mathfrak{p}^k)\right| = p^{s'} \text{ with } s' = (k - \min(e(j-1), k))f = \max(k - e(j-1), 0)f \ .$$

Comparing the two expressions, we obtain the recursion formula

$$\sum_{i \geq j} (i-j+1)a_i = \max(k - e(j-1), 0)f \ . \tag{1}$$

Since $a_i = 0$ for $i > k$, it follows by induction that $a_j = 0$ for $e(j-1) \geq k$; in other words, $a_j = 0$ for $j > \lceil k/e \rceil$ (this is clear anyhow since $p^{\lceil k/e \rceil} \in \mathfrak{p}^k$). Let $k + e - 1 = eq + r$ with $0 \leq r < e$ be the Euclidean division of $k + e - 1$ by e, so that $q = \lfloor (k+e-1)/e \rfloor = \lceil k/e \rceil$. Since $a_i = 0$ for $i > q$, applying the above recursion with $j = q$ gives us

$$a_q = (k - e(q-1))f = (r+1)f \ .$$

Applying the recursion with $j = q - 1$ gives us

$$a_{q-1} = (k - e(q-2))f - 2a_q = (e - 1 - r)f \ .$$

Finally, since

$$qa_q + (q-1)a_{q-1} = (qr + q + qe - q - qr - e + 1 + r)f = (eq - e + r + 1)f = kf \ ,$$

we must have $a_i = 0$ for $i < q - 1$, proving the theorem. □

As the following corollary shows, we can now obtain the multiplicative structure of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ in most cases.

Corollary 4.2.11. *Let $\mathfrak{p}$ be a prime ideal above p, with ramification index $e = e(\mathfrak{p}/p)$ and residual degree $f = f(\mathfrak{p}/p)$, and let $k \geq 2$ be an integer. Write*

$$k + e - 2 = eq + r \quad \text{with} \quad 0 \leq r < e \ .$$

Assume that $p \geq \min(e+2, k)$. Then

$$\left(\mathbb{Z}_K/\mathfrak{p}^k\right)^* \simeq \left(\mathbb{Z}/(p^f-1)\mathbb{Z}\right) \times \left(\mathbb{Z}/p^q\mathbb{Z}\right)^{(r+1)f} \times \left(\mathbb{Z}/p^{q-1}\mathbb{Z}\right)^{(e-r-1)f} \ .$$

Proof. Assume first that $k \geq e+2$. Then $p \geq e+2$ or, in other words, $e < (p-1)$. We can thus apply Corollary 4.2.8, and Lemma 4.2.9, Theorem 4.2.10, together with Proposition 4.2.4, imply the result in this case.

Assume now that $k \leq e+1$. Then $e \leq k+e-2 \leq 2e-1$; hence $q = 1$ and $r = k-2$. Thus, we must prove that

$$(1+\mathfrak{p})/(1+\mathfrak{p}^k) \simeq (\mathbb{Z}/p\mathbb{Z})^{(k-1)f} \ ,$$

and since these groups have the same cardinality, we must simply show that $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$ is killed by p. Since

$$(1+x)^p = 1 + x^p + \sum_{1 \leq i \leq p-1} \binom{p}{i} x^i \ ,$$

when $x \in \mathfrak{p}$ and $1 \leq i \leq p-1$, we have

$$v_{\mathfrak{p}}\left(\binom{p}{i} x^i\right) = e + i v_{\mathfrak{p}}(x) \geq e+1 \geq k \ ,$$

and $v_{\mathfrak{p}}(x^p) = p v_{\mathfrak{p}}(x) \geq p \geq k$ by assumption. Hence, if $x \in \mathfrak{p}$, we have $(1+x)^p \equiv 1 \pmod{\mathfrak{p}^k}$, and so $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$ is killed by p, as claimed, proving the corollary. □

This gives the solution for the structure problem in all but a finite number of cases, but it seems hopeless to have a general nonalgorithmic answer to the problem which is valid in every case. Even in [Nak2], the given answer is algorithmic, although not very usable.

To illustrate this complexity, we give the following supplementary proposition, which covers some more cases (the theorem can be extended at will if desired; see Exercise 19).

Proposition 4.2.12. *Let $\mathfrak{p}$ be a prime ideal above p, of ramification index $e = e(\mathfrak{p}/p)$ and degree f, and let $k \geq 1$ be an integer. We have*

$$(\mathbb{Z}_K/\mathfrak{p}^k)^* \simeq \left(\mathbb{Z}/(p^f-1)\mathbb{Z}\right) \times G_{\mathfrak{p}} \ , \ \text{where} \quad G_{\mathfrak{p}} = \prod_{i=1}^{k-1} \left(\mathbb{Z}/p^i\mathbb{Z}\right)^{a_i}$$

for certain nonnegative integers a_i. For $2 \leq k \leq 4$, they are given by the following table.

(1) *If* $k = 2$, *then* $(a_1) = (f)$.
(2) *If* $k = 3$, *then* (a_1, a_2) *is given by*

$$\begin{array}{ll} (0, f) & \text{if } p \geq 3 \text{ and } e = 1; \\ (2f, 0) & \text{if } p \geq 3 \text{ and } e \geq 2; \\ (2, f-1) & \text{if } p = 2 \text{ and } e = 1; \\ (0, f) & \text{if } p = 2 \text{ and } e \geq 2. \end{array}$$

(3) *If* $k = 4$, *then* (a_1, a_2, a_3) *is given by*

$$\begin{array}{ll} (0, 0, f) & \text{if } p \geq 5 \text{ and } e = 1; \\ (f, f, 0) & \text{if } p \geq 5 \text{ and } e = 2; \\ (3f, 0, 0) & \text{if } p \geq 5 \text{ and } e \geq 3; \\ (0, 0, f) & \text{if } p = 3 \text{ and } e = 1; \\ (f + 2a, f - a, 0) & \text{if } p = 3 \text{ and } e = 2; \\ (f, f, 0) & \text{if } p = 3 \text{ and } e \geq 3; \\ (1, 1, f-1) & \text{if } p = 2 \text{ and } e = 1; \\ (f, f, 0) & \text{if } p = 2 \text{ and } e \geq 2. \end{array}$$

In the above, $a = 1$ *if there exists* $x \in \mathbb{Z}_K$ *such that* $x^2 \equiv -3 \pmod{\mathfrak{p}^3}$, *and* $a = 0$ *otherwise.*

(Note that in this proposition, we have, as usual, mixed multiplicative and additive notation.)

Proof. We must first prove that $G_{\mathfrak{p}}$ is killed by p^{k-1}; in other words, that $(1+x)^{p^{k-1}} \equiv 1 \pmod{\mathfrak{p}^k}$ for all $x \in \mathfrak{p}$. We prove this by induction on k. The statement is trivially true for $k = 1$, so assume that it is true for k. Thus $(1+x)^{p^{k-1}} = 1 + y$ with $y \in \mathfrak{p}^k$. Hence

$$(1+x)^{p^k} = (1+y)^p = 1 + \sum_{1 \leq j \leq p-1} \binom{p}{j} y^j + y^p \ .$$

Since $p \mid \binom{p}{j}$ for $1 \leq j \leq p-1$, we have

$$v_{\mathfrak{p}}\left(\binom{p}{j} y^j\right) = e + j v_{\mathfrak{p}}(y) \geq 1 + k \ .$$

On the other hand,

$$v_{\mathfrak{p}}(y^p) = p v_{\mathfrak{p}}(y) \geq pk \geq k + 1$$

since $p \geq 2$, from which our assertion follows by induction. Note that one can prove a much more precise statement than this (see Exercise 20).

Corollary 4.2.11 gives us directly a number of special cases. Specifically, it gives the cases $k = 2$, $k = 3$ and $p \geq 3$; $k = 4$ and $p \geq 5$; and $k = 4$, $p = 3$, and $e = 1$.

Let us look at the remaining cases. The easiest way is probably to use the following lemma, similar to the proof of Theorem 4.2.10.

Lemma 4.2.13. *With the notation of the above proposition, let p^{k_j} be the cardinality of the kernel of the map $x \mapsto x^{p^j}$ from $G_\mathfrak{p}$ into itself. The exponents a_i are given by $a_i = 0$ for $i \geq k$ and the following backwards recursion:*

$$a_j = (k-1)f - \sum_{i=j+1}^{k-1} (i-j+1)a_i - k_{j-1}.$$

The reader is invited to compare this with recursion (1).

Proof. Since $G_\mathfrak{p} = (1+\mathfrak{p})/(1+\mathfrak{p}^k)$ is killed by p^{k-1}, we can write $G_\mathfrak{p} \simeq \prod_{1 \leq i \leq k-1} (\mathbb{Z}/p^i\mathbb{Z})^{a_i}$. Let K_j be the kernel of the map $x \mapsto x^{p^j}$ from $G_\mathfrak{p}$ into itself, and let $p^{k_j} = |K_j|$ be its cardinality (it will be a power of p since K_j is a subgroup of the p-group $G_\mathfrak{p}$). Then

$$K_j \simeq \prod_{1 \leq i \leq j} (\mathbb{Z}/p^i\mathbb{Z})^{a_i} \prod_{j+1 \leq i \leq k-1} (p^{i-j}\mathbb{Z}/p^i\mathbb{Z})^{a_i} ,$$

from which it follows that

$$k_j = \sum_{1 \leq i \leq j} i a_i + j \sum_{j+1 \leq i \leq k-1} a_i .$$

Since $G_\mathfrak{p} = p^{(k-1)f}$, we have $\sum_{1 \leq i \leq k-1} i a_i = (k-1)f$, and so $k_j = (k-1)f - \sum_{j+1 \leq i \leq k-1} (i-j)a_i$. Changing j into $j-1$ gives the backwards recursion of the lemma. □

Resuming the proof of the corollary, we look at the cases not covered by Corollary 4.2.11.

Assume first that $k = 3$, $p = 2$. We have $K_1 = \{\overline{1+x} \in G_\mathfrak{p},\ (1+x)^2 \equiv 1 \pmod{\mathfrak{p}^3}\}$, so

$$K_1 = \{\overline{1+x} \in G_\mathfrak{p},\ \mathfrak{p}^3 \mid x(x+2)\} = \{\overline{1+x} \in G_\mathfrak{p},\ \mathfrak{p}^2 \mid x \text{ or } \mathfrak{p}^2 \mid x+2\} .$$

Hence, if $e \geq 2$, these two conditions are equivalent, so $K_1 = (1+\mathfrak{p}^2)/(1+\mathfrak{p}^3)$ and $2^{k_1} = \mathcal{N}(\mathfrak{p}) = 2^f$, while if $e = 1$, $K_1 = \pm(1+\mathfrak{p}^2)/(1+\mathfrak{p}^3)$, and $2^{k_1} = 2\mathcal{N}(\mathfrak{p}) = 2^{f+1}$. From the backwards recursion, it follows that $a_2 = f$, $a_1 = 0$ when $e \geq 2$, while $a_2 = f-1$, $a_1 = 2$ when $e = 1$ (note that $k_0 = 0$). A similar reasoning left to the reader gives the formulas for $k = 4$, $p = 2$.

Assume now that $k = 4$, $p = 3$, $e \geq 2$ (the case $e = 1$ follows from Corollary 4.2.11). By definition,

$$K_2 = \{\overline{1+x} \in G_{\mathfrak{p}},\ (1+x)^9 \equiv 1 \pmod{\mathfrak{p}^4}\} \ .$$

Since $e \geq 2$, when $x \in \mathfrak{p}$, we have

$$(1+x)^9 \equiv 1 + 9x + 36x^2 + 84x^3 \equiv 1 \pmod{\mathfrak{p}^4} \ ,$$

and so $K_2 = G_{\mathfrak{p}}$ and hence $3^{k_2} = 3^{3f}$ and $a_3 = 0$.

Similarly,

$$K_1 = \{\overline{1+x} \in G_{\mathfrak{p}},\ (1+x)^3 \equiv 1 \pmod{\mathfrak{p}^4}\} \ .$$

Since $e \geq 2$, when $x \in \mathfrak{p}$, we have $(1+x)^3 = 1 + 3x + 3x^2 + x^3 \equiv 1 + 3x + x^3 \pmod{\mathfrak{p}^4}$, and so

$$K_1 = \{\overline{1+x} \in G_{\mathfrak{p}},\ \mathfrak{p}^4 \mid x(3+x^2)\} \ .$$

If $e \geq 3$, this is equivalent to $\mathfrak{p}^2 \mid x$; hence $K_1 = (1+\mathfrak{p}^2)/(1+\mathfrak{p}^4)$, so $3^{k_1} = 3^{2f}$, and the recursion formula gives $a_2 = f$ and $a_1 = f$.

If $e = 2$, then either $\mathfrak{p}^2 \mid x$ or $x^2 \equiv -3 \pmod{\mathfrak{p}^3}$, these two conditions being exclusive. If this last congruence has no solution (if $a = 0$), then we again have $3^{k_1} = 3^{2f}$ and $a_2 = a_1 = f$. If the congruence has a solution x_0, we have $v_{\mathfrak{p}}(x_0) = 1$, and since $x(3+x^2) \equiv x(x-x_0)(x+x_0) \pmod{\mathfrak{p}^3}$, it follows that

$$K_1 = \{\overline{1+x} \in G_{\mathfrak{p}},\ x \equiv 0, x_0, -x_0 \pmod{\mathfrak{p}^2}\} \ ,$$

and so $3^{k_1} = 3 \cdot 3^{2f} = 3^{2f+1}$. The recursion formula gives $a_2 = f - 1$ and $a_1 = f + 2$, as desired. □

Note that in the above cases, we have given only the abstract structure of the groups $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ and not a complete algorithmic description in the sense of Definition 4.1.4, but this can also easily be done if desired (see Algorithm 4.2.15 below and the discussion that precedes it).

We see that the use of $\mathfrak{p}$-adic logarithms gives a satisfactory answer to our structure problem in most cases (see Exercise 21 for still another possibility of the same nature). However, it is not complete, and we must therefore look for another idea to be able to treat the general problem. We shall see that this idea indeed leads to a complete algorithmic and satisfactory solution to the problem, but not to a theoretical formula of the same nature as the one given by Proposition 4.2.12.

4.2.3 Computing $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ by Induction

We now explain how to algorithmically compute the groups $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ in all cases. The method is based on an induction procedure using the following proposition.

Proposition 4.2.14. (1) *Let $a \leq b \leq c$ be integers. We have the exact sequence*

$$1 \longrightarrow (1+\mathfrak{p}^b)/(1+\mathfrak{p}^c) \longrightarrow (1+\mathfrak{p}^a)/(1+\mathfrak{p}^c) \longrightarrow (1+\mathfrak{p}^a)/(1+\mathfrak{p}^b) \longrightarrow 1 \ .$$

(2) *Assume that $b \leq 2a$. Then the map from the multiplicative group $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ to the additive group $\mathfrak{p}^a/\mathfrak{p}^b$, which sends the class of $1+x$ modulo $1+\mathfrak{p}^b$ to the class of x modulo $\mathfrak{p}^b$, is well-defined and is a group isomorphism.*

Proof. The existence of the exact sequence is trivial. For (2), the definition of $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ (Definition 4.2.3) shows that the map $\overline{1+x} \mapsto \overline{x}$ is a bijection. However, it is not a group homomorphism in general (otherwise, $G_{\mathfrak{p}}$ would always be isomorphic to $\mathfrak{p}/\mathfrak{p}^k$, and we have seen in Proposition 4.2.12 that this is not always the case). If, however, $b \leq 2a$ and x and y belong to $\mathfrak{p}^a$, we have $\mathfrak{p}^b \mid \mathfrak{p}^{2a} \mid xy$, and hence $(1+x)(1+y) = 1+x+y+xy \equiv 1+x+y \pmod{\mathfrak{p}^b}$, and so the map is a group homomorphism. □

Assume that we can algorithmically compute $\mathfrak{p}^a/\mathfrak{p}^b$ for $a \leq b \leq 2a$. Using the explicit isomorphism above, we thus compute $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$. Then using Proposition 4.2.14 (1) and Algorithm 4.1.8, we inductively compute $(1+\mathfrak{p})/(1+\mathfrak{p}^2)$, $(1+\mathfrak{p})/(1+\mathfrak{p}^4), \ldots, (1+\mathfrak{p})/(1+\mathfrak{p}^{2^m})$, $G_{\mathfrak{p}} = (1+\mathfrak{p})/(1+\mathfrak{p}^k)$, where $m = \lfloor \log_2(k-1) \rfloor$.

Thanks to Lemma 4.2.9 and Theorem 4.2.10, we know the structure of $\mathfrak{p}^a/\mathfrak{p}^b$ as an abstract Abelian group. Although everything is explicit, it is not very convenient to deduce from the proof of Theorem 4.2.10 a system of generators and relations for $\mathfrak{p}^a/\mathfrak{p}^b$.

To compute $\mathfrak{p}^a/\mathfrak{p}^b$ algorithmically, the simplest is perhaps to proceed as follows. Let $\mathfrak{p} = p\mathbb{Z}_K + \pi\mathbb{Z}_K$ be a two-element representation of $\mathfrak{p}$, where we may assume π chosen so that $v_{\mathfrak{p}}(\pi) = 1$ (if this is not the case, then $v_{\mathfrak{p}}(p) = 1$, so $\mathfrak{p}$ is unramified and we replace π by $\pi + p$).

Then for all m, if $q = \lceil m/e \rceil = \lfloor (m+e-1)/e \rfloor$ as above, by Proposition 2.3.15 (or directly), we have $\mathfrak{p}^m = p^q\mathbb{Z}_K + \pi^m\mathbb{Z}_K$.

From this, it is easy to compute the Hermite normal form of $\mathfrak{p}^m$ on some fixed integral basis of $\mathbb{Z}_K$: construct the $n \times 2n$ matrix obtained by concatenating p^q times the identity matrix with the $n \times n$ matrix giving the endomorphism multiplication by π^m on the integral basis, and then apply a Hermite normal form algorithm to obtain the desired HNF.

Let A and B be the Hermite normal form of $\mathfrak{p}^a$ and $\mathfrak{p}^b$, respectively. Since $\mathfrak{p}^b \subset \mathfrak{p}^a$, the matrix $A^{-1}B$, which expresses the HNF basis of $\mathfrak{p}^b$ on the HNF basis of $\mathfrak{p}^a$, has integer entries. If we apply the Smith normal form algorithm to this matrix, we will find unimodular matrices U and V such that $UA^{-1}BV = D_C$ is a diagonal matrix in Smith normal form. If $D_C = \mathrm{diag}((c_i)_i)$ and we set $C = AU^{-1}$, then the columns of C give the coordinates on the chosen integral basis of elements $\gamma_i \in \mathfrak{p}^a$, and we have

$\mathfrak{p}^a/\mathfrak{p}^b = \bigoplus(\mathbb{Z}/c_i\mathbb{Z})\overline{\gamma_i}$, where $\overline{\gamma}$ denotes the class of γ modulo $\mathfrak{p}^b$. If, in addition, $b \leq 2a$, it follows from Proposition 4.2.14 that

$$(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b) = \bigoplus(\mathbb{Z}/c_i\mathbb{Z})\overline{(1+\gamma_i)} \ .$$

Note that the above is simply a rephrasing of the method explained in Section 4.1.3.

We can now give formal algorithms for computing $(\mathbb{Z}_K/\mathfrak{p}^k)^*$. We begin with a basic subalgorithm corresponding to Proposition 4.2.14.

Algorithm 4.2.15 (Computation of $\mathfrak{p}^a/\mathfrak{p}^b$ and $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$). Let K be a number field, let $\mathfrak{p}$ be a prime ideal given by a two-element representation, and let a and b be two positive integers such that $b > a$. This algorithm computes integers $c_{a,i}$ and elements $\gamma_{a,i} \in \mathfrak{p}^a$ such that $\mathfrak{p}^a/\mathfrak{p}^b = \bigoplus(\mathbb{Z}/c_{a,i}\mathbb{Z})\overline{\gamma_{a,i}}$ and $c_{a,i+1} \mid c_{a,i}$. Hence, if in addition $b \leq 2a$, $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b) = \bigoplus(\mathbb{Z}/c_{a,i}\mathbb{Z})\overline{(1+\gamma_{a,i})}$. Furthermore, it outputs an additional matrix U_a, which will be needed for discrete logarithm computations.

1. [Compute HNF matrices] By using the method explained above, compute the Hermite normal forms A and B of $\mathfrak{p}^a$ and $\mathfrak{p}^b$, respectively.
2. [Apply Smith] Apply the Smith normal form algorithm to the integral matrix $A^{-1}B$, thus obtaining unimodular matrices U and V such that $UA^{-1}BV = D_C$ is a diagonal matrix in Smith normal form.
3. [Terminate] Let $D_C = \mathrm{diag}((c_{a,i})_i)$. For each i, let $\gamma_{a,i}$ be the element of $\mathbb{Z}_K$ (in fact of $\mathfrak{p}^a$) whose coefficients on the given integral basis are the entries of the ith column of the matrix AU^{-1}. Output $\mathfrak{p}^a/\mathfrak{p}^b = \bigoplus(\mathbb{Z}/c_{a,i}\mathbb{Z})\overline{\gamma_{a,i}}$, and if $b \leq 2a$, $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b) = \bigoplus(\mathbb{Z}/c_{a,i}\mathbb{Z})\overline{(1+\gamma_{a,i})}$. For future use set $U_a \leftarrow UA^{-1}$, output the matrix U_a, and terminate the algorithm.

We could clean up the trivial components as we did at the end of Algorithm 4.1.3. Since this is essentially going to be used only as a subalgorithm of the complete algorithm for computing $(\mathbb{Z}_K/\mathfrak{m})^*$, we will clean up at the very end.

The corresponding discrete logarithm algorithm is essentially trivial. Indeed, since $a \leq 2b$, Proposition 4.2.14 tells us that

$$\prod(1+\gamma_{a,i})^{x_i} \equiv 1 + \sum x_i\gamma_{a,i} \pmod{1+\mathfrak{p}^b} \ .$$

Hence, if $\overline{\beta} \in (1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$, we want to solve $\sum x_i\gamma_{a,i} = \beta - 1$, or in matrix terms on the integral basis, $AU^{-1}X = B - 1_K$, where B is the column vector representing β on the integral basis, and 1_K is the column vector representing 1 (equal to $(1,0,\ldots,0)^t$ since we chose an integral basis starting with 1). It follows that $X = UA^{-1}(B - 1_K) = U_a(B - 1_K)$ is the desired discrete logarithm, and this is the reason we have kept the matrix U_a.

The second basic subalgorithm we need is the following.

Algorithm 4.2.16 (Discrete Logarithm in $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$). Let K be a number field, $\mathfrak{p}$ a prime ideal and k an integer, which we can assume to be greater than or equal to 2; otherwise the problem is trivial. For each $a \geq 1$ such that $2^a \leq k$, we assume that we have computed the $c_{a,i}$, $\gamma_{a,i}$, and U_a corresponding to $b = \min(2a, k)$ by Algorithm 4.2.15. Finally, let $\beta \in (1+\mathfrak{p})$, where β is given by a column vector B on the integral basis. This algorithm computes the discrete logarithm of $\overline{\beta}$ in $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$ with respect to the $\overline{1+\gamma_{a,i}}$; more precisely, it computes integers $y_{a,i}$ such that $\overline{\beta} = \prod_{a,i} \overline{(1+\gamma_{a,i})}^{y_{a,i}}$ in $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$. (The $\gamma_{a,i}$ and $c_{a,i}$ do not give a Smith basis of $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$, but this is not necessary. In addition, a will always be a power of 2.) As above, we let 1_K denote the column vector representing 1.

1. [Initialize] Set $a \leftarrow 1$.
2. [Main step] Set $Z \leftarrow U_a(B - 1_K)$. If $Z = (z_i)$, for each i set $y_{a,i} \leftarrow -((-z_i) \bmod c_{a,i})$, where we choose the smallest nonnegative residue of $-z_i$. Finally, set $\beta \leftarrow \beta \prod_i (1+\gamma_{a,i}^{-y_{a,i}})$. Note that this product should be reduced modulo $\mathfrak{p}^k$ (see Section 4.3.2) and that the exponents are nonnegative.
3. [Loop and terminate] Set $a \leftarrow 2a$. If $a < k$, let B be the column vector whose entries are the coefficients of β on the integral basis, and go to step 2. Otherwise, output the $y_{a,i}$ and terminate the algorithm.

There is a little trick in the main step of this algorithm. We could have simply set $y_{a,i} \leftarrow z_i \bmod c_i$. We would then have to set $\beta \leftarrow \beta / \prod_i (1+\gamma_{a,i}^{y_{a,i}}) \bmod \mathfrak{p}^k$, and although division modulo an ideal is not too difficult, it is slower than multiplication, hence we prefer to use the above trick (see Section 4.3.2).

We are now ready to give the algorithm for computing $(\mathbb{Z}_K/\mathfrak{p}^k)^*$.

Algorithm 4.2.17 (Computation of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$). Let K be a number field, let $\mathfrak{p}$ be a prime ideal of degree f above p given by a two-element representation, and let k be a positive integer. This algorithm computes integers d_i and elements δ_i of $\mathbb{Z}_K$ such that $(\mathbb{Z}_K/\mathfrak{p}^k)^* = \bigoplus (\mathbb{Z}/d_i\mathbb{Z})\overline{\delta_i}$ with $d_{i+1} \mid d_i$. It also outputs a number of other quantities that will be needed in other algorithms.

1. [Initialize] If $k = 1$, go to step 4. Otherwise, set $a \leftarrow 1$ and $b \leftarrow 2$.
2. [Compute $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$] Using Algorithm 4.2.15, compute the quantities $c_{a,i}$, $1+\gamma_{a,i}$, and U_a giving the structure of $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$. Call n_a the number of cyclic components $c_{a,i}$ (or $\gamma_{a,i}$). For future use, output all these quantities.
3. [Loop] Set $a \leftarrow 2a$. If $a < k$, set $b \leftarrow \min(2b, k)$ and go to step 2.
4. [Prime to p part] Set $q \leftarrow p^f$. By choosing elements at random in $\mathbb{Z}_K \setminus \mathfrak{p}$ (essentially by using Algorithm 1.3.13), find $g_0 \in \mathbb{Z}_K$ such that $g_0 \bmod \mathfrak{p}$ is of order exactly $q-1$ in $(\mathbb{Z}_K/\mathfrak{p})^*$ (so that the class of g_0 is a generator of $(\mathbb{Z}_K/\mathfrak{p})^*$).

5. [Start computation of big matrix] Set $c_{0,1} \leftarrow q-1$, $\gamma_{0,1} \leftarrow g_0 - 1$, $n_0 \leftarrow 1$. In the next step, we will compute a square $h \times h$ matrix H, where $h = \sum_{a \geq 0} n_a$. It is very convenient to index the rows and columns of H by the pairs (a,i) as for the generators, and we will consider these pairs to be lexicographically ordered, so that $(a,i) \leq (b,j)$ if and only if $a < b$ or $a = b$ and $i \leq j$.
6. [Compute big matrix H] For each pair (a,i), do the following. Set $H_{a,i} \leftarrow V$, where $V = (v_{b,j})$ is the column vector computed as follows. Set $v_{b,j} \leftarrow 0$ for $(b,j) < (a,i)$, set $v_{a,i} \leftarrow c_{a,i}$. Let $\beta \leftarrow (1+\gamma_{a,i})^{c_{a,i}}$. Using Algorithm 4.2.16, compute the discrete logarithm $(y_{b,j})$ of β (we will have $y_{b,j} = 0$ for $(b,j) \leq (a,i)$). Finally, for $(b,j) > (a,i)$, set $v_{b,j} \leftarrow -y_{b,j}$.
7. [Terminate] Let G be the row vector of the $\overline{1+\gamma_{a,i}}$ (here $\bar{}$ is modulo $\mathfrak{p}^k$), and let H be the matrix whose columns are the $H_{a,i}$. Apply Algorithm 4.1.3 to the system of generators and relations (G,H), output the SNF (P,D_P), the auxiliary matrix $U_{\mathfrak{p}} = U_a$ obtained in that algorithm, and terminate the algorithm.

The corresponding discrete logarithm algorithm is the following.

Algorithm 4.2.18 (Discrete Logarithm in $(\mathbb{Z}_K/\mathfrak{p}^k)^*$). In addition to the data given in Algorithm 4.2.17, we are given an element $\beta \in \mathbb{Z}_K$ coprime to $\mathfrak{p}^k$ (or, equivalently, to $\mathfrak{p}$). This algorithm computes the discrete logarithm of β with respect to the δ_i output by Algorithm 4.2.17.

1. [Compute discrete log modulo $\mathfrak{p}$] Using, for example, Shanks's baby-step, giant-step method or a more sophisticated method, compute the discrete logarithm $y_{0,1}$ of $\overline{\beta}$ with respect to $\overline{g_0}$ in $(\mathbb{Z}_K/\mathfrak{p})^*$ (this may be the most time-consuming part of the algorithm if $q = |\mathbb{Z}_K/\mathfrak{p}|$ is large). Then set $\beta \leftarrow \beta/g_0^{y_{0,1}} \bmod \mathfrak{p}^k$.
2. [Use Algorithm 4.2.16] (Here $\beta \in (1+\mathfrak{p})$.) Compute the discrete logarithm $(y_{a,i})$ of β in $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$ in the sense of Algorithm 4.2.16, and let $Y = (y_{a,i})$ be the column vector of the $y_{a,i}$ (always in lexicographic order, and including $y_{0,1}$).
3. [Terminate] Using the matrix $U_{\mathfrak{p}}$ output in Algorithm 4.2.17 (whose columns are indexed by the pairs (a,i), but whose rows are indexed normally), compute $X \leftarrow U_{\mathfrak{p}}Y$, output X, and terminate the algorithm.

Remark. Using the result of Exercise 21, we have at our disposal at least three methods for computing the structure of $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$.

(1) The use of $\mathfrak{p}$-adic logarithms. This method gives the result in one step if $e < p-1$; otherwise one needs to use other methods to compute the structure of $(1+\mathfrak{p})/(1+\mathfrak{p}^{k_0})$ with $k_0 = 1 + \lfloor e/(p-1) \rfloor$.
(2) The use of the map $1+x \mapsto x$ as we have done above. This method needs to be applied recursively since it is applicable only for $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ when $b \leq 2a$, and the number of iterations is roughly $\log k / \log 2$.

(3) The use of the Artin–Hasse logarithm explained in Exercise 21, in other words the map $1+x \mapsto \sum_{1\le i<p}(-1)^{i-1}x^i/i$ modulo $\mathfrak{p}^b$. This method also needs to be applied recursively (unless $k \le p$) since it is applicable only for $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ when $b \le pa$, and the number of iterations is roughly $\log k/\log p$. Thus, this method always needs fewer iterations than the previous method, at the expense of the computation of a more complicated function. It is not clear which method is the fastest.

The computation of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ by the recursive method explained above has the advantage of working in all cases, but it is rather heavy and can lead to quite large generators. In the next section we will see how to reduce the size of these generators. We can, however, usually improve the above algorithm by using a *combination* of $\mathfrak{p}$-adic logarithms and exponentials, with the recursive method. Indeed, by Corollary 4.2.8, we know that if $k_0 = 1 + \lfloor e/(p-1)\rfloor$, then the $\mathfrak{p}$-adic logarithm and exponential give isomorphisms between $(1+\mathfrak{p}^{k_0})/(1+\mathfrak{p}^k)$ and $\mathfrak{p}^{k_0}/\mathfrak{p}^k$. Thus, we can use Algorithm 4.2.17 to compute explicitly $(1+\mathfrak{p})/(1+\mathfrak{p}^{k_0})$, $\mathfrak{p}$-adic techniques to compute $(1+\mathfrak{p}^{k_0})/(1+\mathfrak{p}^k)$, and Proposition 4.2.14 and Algorithm 4.1.8 to put both structures together. The details are left to the reader, but a serious implementation should use this approach (see Exercise 22).

One important special case of this that deserves mention is when $e = p-1$, which is of frequent use in explicit constructions of Kummer theory (see Chapter 5). In this case, we have the following proposition.

Proposition 4.2.19. *Assume $\mathfrak{p}$ is a prime ideal above p of ramification index e and residual degree f, and assume that $e = p-1$. Let ω_i be such that $(\mathbb{Z}_K/\mathfrak{p}) = \bigoplus_{i\in D_p}(\mathbb{Z}/p\mathbb{Z})\overline{\omega_i}$ with the notation of Proposition 2.4.6 and Corollary 2.4.7, and let π be a uniformizer of $\mathfrak{p}$ (in other words, $\pi \in \mathfrak{p}\setminus\mathfrak{p}^2$). Finally, let $\overline{\gamma_j}$ as output by Algorithm 4.2.15 be such that $\mathfrak{p}^2/\mathfrak{p}^p = \bigoplus(\mathbb{Z}/p^{c_j}\mathbb{Z})\overline{\gamma_j}$ with $c_j \ge 1$ (after removing the trivial components).*

Then for all j we have $c_j = 1$, $(1+\mathfrak{p})/(1+\mathfrak{p}^p)$ is a $\mathbb{Z}/p\mathbb{Z}$ vector space of dimension $ef = (p-1)f$, and a basis for this vector space is given by the classes of $1+\pi\omega_i$ for $i \in D_p$, together with the classes of the $\exp_{\mathfrak{p}}(\gamma_j)$. In other words,

$$(1+\mathfrak{p})/(1+\mathfrak{p}^p) = \bigoplus_{i\in D_p}(\mathbb{Z}/p\mathbb{Z})\overline{(1+\pi\omega_i)} \bigoplus_{1\le j\le(p-2)f}(\mathbb{Z}/p\mathbb{Z})\overline{\exp_{\mathfrak{p}}(\gamma_j)}\ .$$

Proof. Since π is a uniformizer of $\mathfrak{p}$, we have $\mathfrak{p}/\mathfrak{p}^2 = \bigoplus_{i\in D_p}(\mathbb{Z}/p\mathbb{Z})\overline{\pi\omega_i}$; hence

$$(1+\mathfrak{p})/(1+\mathfrak{p}^2) = \bigoplus_{i\in D_p}(\mathbb{Z}/p\mathbb{Z})\overline{(1+\pi\omega_i)}\ .$$

Note that these are equalities, and not only isomorphisms.

On the other hand, $\mathfrak{p}^2/\mathfrak{p}^p$ is clearly killed by p, so $c_j = 1$ for all j such that $c_j > 0$. Since $e = p-1$, by Corollary 4.2.8, we also have

$$(1+\mathfrak{p}^2)/(1+\mathfrak{p}^p) = \bigoplus_{1\le j\le (p-2)f} (\mathbb{Z}/p\mathbb{Z})\overline{\exp_{\mathfrak{p}}(\gamma_j)} \ .$$

Finally, I claim that the exact sequence

$$1 \longrightarrow (1+\mathfrak{p}^2)/(1+\mathfrak{p}^p) \longrightarrow (1+\mathfrak{p})/(1+\mathfrak{p}^p) \longrightarrow (1+\mathfrak{p})/(1+\mathfrak{p}^2) \longrightarrow 1$$

is *split*, which will prove the proposition. To prove this, instead of giving a direct proof (which is easy; see Exercise 23), we will use Algorithm 4.1.8.

Using the same symbol $\overline{\ }$ to mean the class modulo different subgroups, we have $A = \left(\overline{\exp_{\mathfrak{p}}(\gamma_j)}\right)_j$, $D_A = pI_{f(e-1)}$, $C = \left(\overline{1+\pi\omega_i}\right)_{i\in D_p}$, and $D_C = pI_f$ (where I_n always denotes the $n \times n$ identity matrix). Thus, following the algorithm, we take $B' = \left(\overline{1+\pi\omega_i}\right)_{i\in D_p}$ and $\psi(A) = \left(\overline{\exp_{\mathfrak{p}}(\gamma_j)}\right)$.

Set $\alpha = 1+\pi\omega_i$. By the binomial theorem, we have

$$\alpha^p = 1 + \sum_{1\le k\le p-1} \binom{p}{k} \pi^k \omega_i^k + \pi^p \omega_i^p \ .$$

Now $v_{\mathfrak{p}}(\pi^p\omega_i^p) \ge pv_{\mathfrak{p}}(\pi) = p$, while for $1 \le k \le p-1$,

$$v_{\mathfrak{p}}\left(\binom{p}{k}\pi^k\omega_i^k\right) \ge v_{\mathfrak{p}}(p) + k = e + k \ge p$$

since $e = p-1$. It follows that $B'' = B'D_C$ is made only of unit elements of $(1+\mathfrak{p})/(1+\mathfrak{p}^p)$; hence we can take A'' also made of unit elements, so we can take $P = 0$. Thus,

$$G = (\psi(A)|B') = \left(\left(\overline{\exp_{\mathfrak{p}}(\gamma_j)}\right)_j \middle| \left(\overline{1+\pi\omega_i}\right)_{i\in D_p}\right)$$

and $M = pI_{ef}$, which is already in SNF, so we have proved both our claim and the proposition. □

Remark. If we use the Artin–Hasse exponential $\exp_a$ (see Exercise 21), we have also

$$(1+\mathfrak{p})/(1+\mathfrak{p}^p) = \bigoplus_{1\le j\le (p-1)f} (\mathbb{Z}/p\mathbb{Z})\overline{\exp_a(\delta_j)} \ ,$$

where the δ_j are the generators of the additive group $\mathfrak{p}/\mathfrak{p}^p$. This may seem simpler than the formula given by the above proposition, but it is not clear if it is really any faster since the computation of the $\exp_a(\delta_j)$ is longer than that of the $\pi\omega_i$.

4.2.4 Representation of Elements of $(\mathbb{Z}_K/\mathfrak{m})^*$

We are now ready to give a complete algorithm for computing $(\mathbb{Z}_K/\mathfrak{m})^*$ and the corresponding discrete logarithm algorithm. Before doing this, we must

first explain how to represent elements of $(\mathbb{Z}_K/\mathfrak{m})^*$. The immediate idea that comes to mind is to represent them as classes of elements of $\mathbb{Z}_K$ modulo the equivalence relation defining $(\mathbb{Z}_K/\mathfrak{m})^*$. In fact, this idea is almost forced upon us by the notation used.

This has two closely related flaws. First, the surjectivity of the map going from the subset of $\mathbb{Z}_K$ of elements coprime to $\mathfrak{m}$ to $(\mathbb{Z}_K/\mathfrak{m})^*$ is not completely trivial, since it is a consequence of the strong approximation theorem in Dedekind domains. Second, the elements of $\mathbb{Z}_K$ we will have to choose to represent elements of $(\mathbb{Z}_K/\mathfrak{m})^*$ will have to be quite "large" since they must have specific signatures.

There is, however, a better representation. If $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$, we represent an element in $(\mathbb{Z}_K/\mathfrak{m})^*$ as a pair $(\overline{\alpha}, v)$, where $\overline{\alpha} \in (\mathbb{Z}_K/\mathfrak{m}_0)^*$ and $v \in \mathbb{F}_2^{\mathfrak{m}_\infty}$ considered as a *column* vector. This is simply the definition of the group $(\mathbb{Z}_K/\mathfrak{m})^*$, but the whole point is that it is much simpler to handle these pairs than directly elements of $\mathbb{Z}_K$ together with their signatures. Note that even when $\mathfrak{m}$ is an ideal — in other words, when $\mathfrak{m}_\infty = \varnothing$ — we still consider pairs $(\overline{\alpha}, v)$, where v is the unique vector in 0-dimensional space over $\mathbb{F}_2$.

If $(\overline{\alpha}, v) \in (\mathbb{Z}_K/\mathfrak{m})^*$, we will say that $\overline{\alpha}$ is the *finite part* and v the *infinite part*, or the *Archimedean part*. The group law in $(\mathbb{Z}_K/\mathfrak{m})^*$ corresponds to multiplying the finite parts and adding the infinite parts.

In all the algorithms that we will present, the above representation is sufficient and simpler than the one-element representation. In some cases, however, it may be desirable to obtain such a one-element representation. For this, the following naive algorithm works well.

Algorithm 4.2.20 (One-Element Representation in $(\mathbb{Z}_K/\mathfrak{m})^*$). Let $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$ be a modulus and $(\overline{\alpha}, v)$ a pair representing an element of $(\mathbb{Z}_K/\mathfrak{m})^*$, with $v = (v_j)_{j\in\mathfrak{m}_\infty}$. Call s the sign homomorphism from K^* to $\mathbb{F}_2^{\mathfrak{m}_\infty}$. This algorithm computes an element $\beta \in \mathbb{Z}_K$ such that $\beta \equiv \alpha \pmod{{}^*\mathfrak{m}_0}$ and $s(\beta) = v$.

1. [Initialize] If it has not already been done, compute a $\mathbb{Z}$-basis $\gamma_1, \ldots, \gamma_n$ of the ideal $\mathfrak{m}_0$ and set $k \leftarrow |\mathfrak{m}_\infty|$.
2. [Find elements] By considering small linear combinations of the γ_i, find k elements $\beta_1, \ldots, \beta_k$ in $\mathfrak{m}_0$ such that the matrix A over $\mathbb{F}_2$ whose columns are the $s(1+\beta_j)$ is invertible.
3. [Multiply] Set $w \leftarrow A^{-1}v$, and let $w = (w_j)_{1\le j\le k}$. Set $\beta \leftarrow \alpha$, and for each j such that $w_j \neq 0$, set $\beta \leftarrow \beta(1+\beta_j)$. Output β and terminate the algorithm.

Evidently, if several conversions of this sort must be done, steps 1 and 2 should be done once and for all. The final β may be large, and it is desirable to reduce it. This cannot be done too rashly, however, since we must preserve the signature of β. We will discuss this in Section 4.3.2.

Warning. As we have already mentioned, the map from $(\mathbb{Z}_K/\mathfrak{m})^*$ to $Cl_{\mathfrak{m}}(K)$ used in Proposition 3.2.3 is *not* the map coming from the algorithmically natural representation $(\overline{\alpha}, v)$ but the map coming from the above one-element representation.

4.2.5 Computing $(\mathbb{Z}_K/\mathfrak{m})^*$

Using Algorithm 4.2.2, we can now put everything together and obtain the algorithmic description of the group $(\mathbb{Z}_K/\mathfrak{m})^*$, in the sense of Definition 4.1.4. Thanks to the representation explained above, the algorithms are very easy to implement (they would be much more painful if we used the one-element representation).

Call s the signature homomorphism from K^* to $\mathbb{F}_2^{\mathfrak{m}_\infty}$ defined by

$$s(\alpha) = (\mathrm{sign}(v(\alpha)))_{v\in\mathfrak{m}_\infty} .$$

Denote by $\mathbf{0}$ the zero vector in $\mathbb{F}_2^{\mathfrak{m}_\infty}$. We will apply our exact sequence techniques to the split exact sequence

$$0 \longrightarrow \mathbb{F}_2^{\mathfrak{m}_\infty} \longrightarrow (\mathbb{Z}_K/\mathfrak{m})^* \longrightarrow (\mathbb{Z}_K/\mathfrak{m}_0)^* \longrightarrow 1 .$$

For $1 \le j \le |\mathfrak{m}_\infty|$, set $\varepsilon_j = \left(\overline{1}, e_j\right) \in (\mathbb{Z}_K/\mathfrak{m})^*$, where e_j denotes the jth canonical basis element of $\mathbb{F}_2^{\mathfrak{m}_\infty}$. The ε_j form a generating set for $\mathbb{F}_2^{\mathfrak{m}_\infty}$, and the matrix of relations between them is clearly equal to twice the identity matrix of order $k = |\mathfrak{m}_\infty|$. If (C, D_C) are generators and relations for $(\mathbb{Z}_K/\mathfrak{m}_0)^*$ with $C = \left(\overline{\gamma_i}\right)$, we lift the $\overline{\gamma_i}$ to $\gamma_i' = \left(\overline{\gamma_i}, \mathbf{0}\right)$ (this is the reason the sequence is split). The γ_i' together with the ε_j form a generating set for $(\mathbb{Z}_K/\mathfrak{m})^*$ whose relation matrix is equal to $\left(\begin{smallmatrix} D_C & 0 \\ 0 & 2I_k \end{smallmatrix}\right)$ and we conclude as usual with a Smith normal form computation.

Combining all this with the methods of Sections 4.1.3 and 4.1.4, we obtain the following algorithm for computing $(\mathbb{Z}_K/\mathfrak{m})^*$.

Algorithm 4.2.21 (Computation of $(\mathbb{Z}_K/\mathfrak{m})^*$). Given a modulus $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$, this algorithm computes the SNF $(\mathbb{Z}_K/\mathfrak{m})^* = (Z, D_Z)$. The entries ζ_i of Z will be represented as pairs $\left(\overline{\gamma_i}, v_i\right)$, where $\overline{\gamma_i}$ denotes the class modulo $\mathfrak{m}_0$ of an element coprime to $\mathfrak{m}_0$ and $v_i \in \mathbb{F}_2^{\mathfrak{m}_\infty}$. The algorithm also outputs some additional information necessary for computing discrete logarithms.

1. [Factor $\mathfrak{m}_0$] Using Algorithm 2.3.22, find distinct prime ideals $\mathfrak{p}$ and exponents $v_\mathfrak{p}$ such that $\mathfrak{m}_0 = \prod_\mathfrak{p} \mathfrak{p}^{v_\mathfrak{p}}$.
2. [Compute the $(\mathbb{Z}_K/\mathfrak{p}^{v_\mathfrak{p}})^*$] For each $\mathfrak{p}$ dividing $\mathfrak{m}_0$, apply Algorithm 4.2.17, thus finding integers $d_{\mathfrak{p},i}$ and elements $\delta_{\mathfrak{p},i}$ of $\mathbb{Z}_K$ coprime to $\mathfrak{p}$ such that $(\mathbb{Z}_K/\mathfrak{p}^{v_\mathfrak{p}})^* = \bigoplus_i (\mathbb{Z}/d_{\mathfrak{p},i}\mathbb{Z})\overline{\delta_{\mathfrak{p},i}}$, and let $n_\mathfrak{p}$ be the number of cyclic components in this sum. For future use output the auxiliary matrix $U_\mathfrak{p}$ also given by this algorithm.
3. [Modify generators] For each $\mathfrak{p}$ dividing $\mathfrak{m}_0$ do the following. Using Lemma 4.2.1 applied to $\mathfrak{a} = \mathfrak{p}^{v_\mathfrak{p}}$ and $\mathfrak{c} = \mathfrak{m}_0/\mathfrak{a}$, compute elements $a_\mathfrak{p} \in \mathfrak{a}$ and $c_\mathfrak{p} \in \mathfrak{c}$ such that $a_\mathfrak{p} + c_\mathfrak{p} = 1$. Then for all i, set $\varepsilon_{\mathfrak{p},i} \leftarrow \left(\overline{c_\mathfrak{p}\delta_{\mathfrak{p},i} + a_\mathfrak{p}}, \mathbf{0}\right) \in (\mathbb{Z}_K/\mathfrak{m})^*$ (these generators are coprime to $\mathfrak{m}_0$, and they are still congruent to the initial $\delta_{\mathfrak{p},i}$ modulo $\mathfrak{p}^{v_\mathfrak{p}}$).
4. [Deal with $\mathfrak{m}_\infty$] Set $n_\infty \leftarrow |\mathfrak{m}_\infty|$, and for $i = 1$ to $i = n_\infty$ set $d_{\infty,i} \leftarrow 2$ and $\varepsilon_{\infty,i} \leftarrow \left(\overline{1}, e_j\right)$, where e_j is the jth canonical basis vector of $\mathbb{F}_2^{\mathfrak{m}_\infty}$.

5. [Compute big matrix M] Let S be the set formed by all prime ideals dividing $\mathfrak{m}_0$ and the symbol ∞. In this step, we will create a square $h \times h$ matrix M, where $h = \sum_{\mathfrak{p} \in S} n_{\mathfrak{p}}$. It is convenient to index the rows and columns of M by the pairs $(\mathfrak{p}, i)$ as for the generators. Then set $M_{(\mathfrak{p},i),(\mathfrak{q},j)} \leftarrow 0$ if $(\mathfrak{q}, j) \neq (\mathfrak{p}, i)$, $M_{(\mathfrak{p},i),(\mathfrak{p},i)} \leftarrow d_{\mathfrak{p},i}$ otherwise.
6. [Terminate] Let G be the row vector of the $\varepsilon_{\mathfrak{p},i}$, and let M be the matrix whose columns are the $M_{\mathfrak{p},i}$. Apply Algorithm 4.1.3 to the system of generators and relations (G, M), output the SNF (Z, D_Z) and the auxiliary matrix U_a obtained in that algorithm, and terminate the algorithm.

It will be useful to compute discrete logarithms for elements of K^* coprime to $\mathfrak{m}$ which are not necessarily in $\mathbb{Z}_K$ (see Definition 3.2.1). For this, we need the following subalgorithm.

Algorithm 4.2.22 (Coprime Representative Computation). Given a nonzero integral ideal $\mathfrak{a}$ and an element β of K^* coprime to $\mathfrak{a}$, this algorithm computes elements α and γ of $\mathbb{Z}_K$ coprime to $\mathfrak{a}$ such that $\beta = \alpha/\gamma$. We assume β given by its coordinates on an integral basis of K.

1. [Trivial case] Let d be the lowest common multiple of the denominators of the coordinates of β, and set $\mathfrak{b} \leftarrow d\mathbb{Z}_K + \mathfrak{a}$. If $\mathfrak{b} = \mathbb{Z}_K$, set $\gamma \leftarrow d$, set $\alpha \leftarrow d\beta$, and terminate the algorithm.
2. [Compute exponent] Let $\mathfrak{b} = \prod_{\mathfrak{p}} \mathfrak{p}^{v_{\mathfrak{p}}}$ be the prime ideal decomposition of $\mathfrak{b}$. Denote by $e(\mathfrak{p})$ the ramification index of $\mathfrak{p}$. Compute

$$k \leftarrow \sup_{\mathfrak{p} \mid \mathfrak{b}} \lfloor v_p(d) e(\mathfrak{p}) / v_{\mathfrak{p}} \rfloor + 1 \ ,$$

   where $v_p(d)$ denotes the ordinary exponent of p in d and p is the prime number below $\mathfrak{p}$.
3. [Compute $\mathfrak{d}^{-1}$] Using standard ideal operations, compute the ideal $\mathfrak{d} \leftarrow d\mathbb{Z}_K + \mathfrak{b}^k$ and the inverse ideal $\mathfrak{d}^{-1}$.
4. [Terminate] (Here $d\mathfrak{d}^{-1}$ and $\mathfrak{b}^k\mathfrak{d}^{-1}$ are coprime integral ideals.) Using Algorithm 1.3.2, compute a and c such that $a \in d\mathfrak{d}^{-1}$, $c \in \mathfrak{b}^k\mathfrak{d}^{-1}$, and $a + c = 1$. Output $\alpha \leftarrow a\beta$ and $\gamma \leftarrow a$, and terminate the algorithm.

We leave to the reader the (easy) proof of this algorithm's validity (Exercise 24). Note that steps 3 and 4 are applications of Theorem 1.3.3. □

In the important special case where $\mathfrak{a} = m\mathbb{Z}_K$ is a principal ideal generated by an element of $\mathbb{Z}$, we have the following proposition whose proof can of course be immediately made into an algorithm (compare also with Lemma 1.2.31).

Proposition 4.2.23. *Let $\beta \in K^*$ be such that $v_{\mathfrak{p}}(\beta) \geq 0$ for all $\mathfrak{p} \mid m\mathbb{Z}_K$ (this is the case, in particular, if β is coprime to $m\mathbb{Z}_K$). Then the least common multiple of the denominators occurring in the representation of β*

on an integral basis is coprime to m. In other words, there exist $d \in \mathbb{Z}$ and $\alpha \in \mathbb{Z}_K$ such that $\beta = \alpha/d$ and $(d, m) = 1$.

Proof. Write $\beta = \alpha_0/d_0$ for $\alpha_0 \in \mathbb{Z}_K$ and $d_0 \in \mathbb{Z}$, for the moment arbitrary. Let

$$g = (d_0, m^\infty) = \prod_{\mathfrak{p}|d_0,\ \mathfrak{p}|m} \mathfrak{p}^{v_\mathfrak{p}(d_0)} .$$

By definition, we have $(d_0/g, m) = 1$. On the other hand, let $\mathfrak{p}$ be a prime ideal and p be the prime number below $\mathfrak{p}$. Then either $p \nmid g$, in which case $v_\mathfrak{p}(\alpha_0/g) = v_\mathfrak{p}(\alpha_0) \geq 0$, or $p \mid g$, in which case we have $p \mid d_0$, $p \mid m$, and $v_p(g) = v_p(d_0)$; hence $v_\mathfrak{p}(g) = v_\mathfrak{p}(d_0) = v_p(g)e(\mathfrak{p}/p)$, so

$$v_\mathfrak{p}\left(\frac{\alpha_0}{g}\right) = v_\mathfrak{p}(\alpha_0) - v_\mathfrak{p}(d_0) = v_\mathfrak{p}\left(\frac{\alpha_0}{d_0}\right) = v_\mathfrak{p}(\beta) \geq 0 .$$

It follows that $\alpha_0/g \in \mathbb{Z}_K$, so $\beta = (\alpha_0/g)/(d_0/g)$ is a suitable representation, proving the proposition. □

The discrete logarithm algorithm in $(\mathbb{Z}_K/\mathfrak{m})^*$ applied to elements of K^* coprime to $\mathfrak{m}$ is as follows.

Algorithm 4.2.24 (Discrete Logarithm in $(\mathbb{Z}_K/\mathfrak{m})^*$). Given a modulus $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$, the structure $(\mathbb{Z}_K/\mathfrak{m})^* = \bigoplus_i (\mathbb{Z}/z_i\mathbb{Z})\zeta_i$ found by Algorithm 4.2.21, and an element β of K^* coprime to $\mathfrak{m}$, this algorithm computes the discrete logarithm of β with respect to the ζ_i. We assume β given by its coordinates on a fixed integral basis. We let $\mathfrak{m}_0 = \prod_\mathfrak{p} \mathfrak{p}^{v_\mathfrak{p}}$ and, as in Algorithm 4.2.21, S is the set of prime ideals dividing $\mathfrak{m}_0$ union the symbol ∞.

1. [Check if integral] If $\beta \in \mathbb{Z}_K$ (in other words, if the coordinates of β on the integral basis are all integers), go to step 2. Otherwise, using Algorithm 4.2.22, compute α and γ in $\mathbb{Z}_K$ coprime to $\mathfrak{m}$ such that $\beta = \alpha/\gamma$. Let L_α (resp., L_γ) be the discrete logarithm of α (resp., γ) obtained by applying the present algorithm. Output $L_\alpha - L_\gamma$ (where each coordinate can be reduced modulo the respective z_i), and terminate the algorithm.
2. [Compute modulo $\mathfrak{p}^{v_\mathfrak{p}}$] (Here $\beta \in \mathbb{Z}_K$.) Using Algorithm 4.2.18, compute the discrete logarithm $(x_{\mathfrak{p},i})$ of β in $(\mathbb{Z}_K/\mathfrak{p}^{v_\mathfrak{p}})^*$. Compute the vector $V = (\overline{v_i})$ of the signature of β, and for $1 \leq i \leq |\mathfrak{m}_\infty|$ set $x_{\infty,i} \leftarrow v_i$, where v_i is any lift of $\overline{v_i}$ to $\mathbb{Z}$ (for example, in $\{0, 1\}$). Finally, let $X = (x_{\mathfrak{p},i})$ be the column vector indexed by the pairs $(\mathfrak{p}, i)$ for $\mathfrak{p} \in S$.
3. [Terminate] Set $W \leftarrow U_a X$, where U_a is the matrix output in step 6 of Algorithm 4.2.21. Reduce each component w_j of W modulo the corresponding z_j (for example, in the interval $[0, z_j - 1]$), output W, and terminate the algorithm.

This terminates the algorithmic computation of the group $(\mathbb{Z}_K/\mathfrak{m})^*$.

4.3 Computing Ray Class Groups

4.3.1 The Basic Ray Class Group Algorithm

Let $\mathfrak{m}$ be a modulus. Recall the exact sequence coming from Proposition 3.2.3:

$$U(K) \xrightarrow{\rho} (\mathbb{Z}_K/\mathfrak{m})^* \xrightarrow{\psi} Cl_{\mathfrak{m}}(K) \xrightarrow{\phi} Cl(K) \longrightarrow 1 \ .$$

To apply the techniques that we have developed, we need to verify a number of things. First, the groups $Cl(K)$ and $U(K)$ must be known in the sense of Definition 4.1.4. This can be done using either the techniques of [Poh-Zas] or those of [Coh0]. Note that [Coh0] assumes the GRH, but in fact in practical situations it is rather easy to remove the GRH condition by *certifying* the result unconditionally. We refer to [Zan] and [Dia-Oli] for details. Note also that we need a discrete logarithm algorithm in $Cl(K)$ (in $U(K)$ the problem is ordinary linear algebra; see Algorithm 5.3.10). The solution to this is also given in [Coh0], where, in fact, even more information is obtained as part of the *principal ideal algorithm*: if an ideal is principal, the algorithm also gives a generator. More precisely, if the g_i are ideals such that the $\overline{g_i}$ are the given generators of $Cl(K)$, then if g is an ideal of K, we can find (v_i) such that $\overline{g} = \prod_i \overline{g_i}^{v_i}$, but the same algorithm gives also $\alpha \in K$ such that $g = \alpha \prod_i g_i^{v_i}$. We will also do this in the context of ray class groups.

The group $(\mathbb{Z}_K/\mathfrak{m})^*$ has been dealt with extensively in Section 4.2.

Finally, we must show that the maps are effective, in the sense of Definition 4.1.5. This is not completely trivial. First, consider the map ψ from $(\mathbb{Z}_K/\mathfrak{m})^*$ to $Cl_{\mathfrak{m}}(K)$. Since $Cl_{\mathfrak{m}}(K)$ is not yet known, to say that ψ is effective means that if $\overline{g} \in Cl_{\mathfrak{m}}(K)$ is of the form $\psi(\overline{\alpha})$, we can find $\alpha \in \mathbb{Z}_K$. But then g is an ideal of K coprime to $\mathfrak{m}$ that is a principal ideal in the ordinary sense, hence by using the principal ideal algorithm mentioned above, we can algorithmically find α such that $g = \alpha\mathbb{Z}_K$. Since g is coprime to $\mathfrak{m}$, α will also be coprime to $\mathfrak{m}$. Using Algorithm 4.2.22, we can find β and γ such that $\alpha = \beta/\gamma$, and β and γ are integral and coprime to $\mathfrak{m}$, hence we can take $\overline{\alpha} = \overline{\beta}/\overline{\gamma}$ in $(\mathbb{Z}_K/\mathfrak{m})^*$.

Consider now the map ϕ from $Cl_{\mathfrak{m}}(K)$ to $Cl(K)$. Since $Cl_{\mathfrak{m}}(K)$ is not yet known, to say that ϕ is effective means that if $\overline{g} \in Cl_{\mathfrak{m}}(K)$, we can compute $\phi(\overline{g}) \in Cl(K)$, which is of course trivial, but it also means that if $\overline{g} \in Cl(K) = \mathrm{Im}(\phi)$, we can find an ideal g' coprime to $\mathfrak{m}$ such that $\phi(\overline{g'}) = \overline{g}$. This follows from Algorithm 1.3.14.

We can now put everything together. We consider the above ray class group exact sequence as a right four-term exact sequence and apply the results of Section 4.1.5, giving the following algorithm for computing the ray class group $Cl_{\mathfrak{m}}(K)$.

Algorithm 4.3.1 (Computing Ray Class Groups). Let $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$ be a modulus. This algorithm computes ideals h_i coprime to $\mathfrak{m}$ such that the SNF of

$Cl_{\mathfrak{m}}(K)$ is (B, D_B), where $B = \left(\overline{h_i}\right)_i$ and $\overline{h_i}$ denotes the ideal class of h_i in $Cl_{\mathfrak{m}}(K)$. It also outputs some extra information necessary for computing discrete logarithms. We assume that we have already computed $U(K) = (E, D_E)$ with $E = (\varepsilon_i)_{0 \le i \le r}$, $Cl(K) = (C, D_C)$ with $C = (\gamma_i) = \left(\overline{g_i}\right)$ (using [Poh-Zas] or [Coh0]), and $(\mathbb{Z}_K/\mathfrak{m})^* = (Z, D_Z)$ with $Z = (\zeta_i)$ (using Algorithm 4.2.21). We denote by ψ the map from $(\mathbb{Z}_K/\mathfrak{m})^*$ to $Cl_{\mathfrak{m}}(K)$.

1. [Find new g_i] Using Algorithm 1.3.14, for each i compute $\alpha_i' \in K^*$ such that $g_i' = \alpha_i' g_i$ is an integral ideal coprime to $\mathfrak{m}$. Let G' be the row vector of the g_i', and let A' be the row vector of the α_i'. For future use, output the elements α_i'.
2. [Find principal ideals] For each ideal g_i, compute $g_i^{c_i}$ (where $D_C = \mathrm{diag}((c_i)_i)$), and using the principal ideal algorithm (see [Coh0, Algorithm 6.5.10]), find $\alpha_i \in \mathbb{Z}_K$ such that $g_i^{c_i} = \alpha_i \mathbb{Z}_K$.
3. [Compute P] (Here the $\alpha_i'^{c_i}\alpha_i$ are elements of $\mathbb{Z}_K$ coprime to $\mathfrak{m}$.) Using Algorithm 4.2.24, compute the matrix P whose columns are the discrete logarithms of the $\alpha_i'^{c_i}\alpha_i$ with respect to the ζ_i.
4. [Compute Q] Again using Algorithm 4.2.24, compute the matrix Q whose columns are the discrete logarithms of the ε_j with respect to the ζ_i for $0 \le j \le r$.
5. [Terminate] Let $B' \leftarrow (\psi(Z)|G')$ and $M \leftarrow \begin{pmatrix} Q & D_Z & -P \\ 0 & 0 & D_C \end{pmatrix}$. Apply Algorithm 4.1.3 to the system of generators and relations (B', M), and let (B, D_B) be the Smith normal form of $Cl_{\mathfrak{m}}(K)$ thus obtained. If $B = (\beta_i)$, for each i let h_i be an ideal (coprime to $\mathfrak{m}$) belonging to the class β_i. Output the h_i, D_B, the auxiliary matrix U_a output by Algorithm 4.1.3, and terminate the algorithm.

To end this section, we give a corresponding discrete logarithm algorithm in $Cl_{\mathfrak{m}}(K)$. As in the case of $Cl(K)$ itself, we will actually solve a stronger problem and write a principal ideal algorithm in ray class groups.

Algorithm 4.3.2 (Principal Ideal Algorithm in Ray Class Groups). Let $\mathfrak{m}$ be a modulus and let $(\mathbb{Z}_K/\mathfrak{m})^* = (Z, D_Z)$ and $Cl_{\mathfrak{m}}(K) = (B, D_B)$ be as computed by Algorithms 4.2.21 and 4.3.1, respectively. Write $B = \left(\overline{h_i}\right)_i$, and let H denote the row vector of ideals h_i. Given a fractional ideal $\mathfrak{a}$ coprime to $\mathfrak{m}$, this algorithm computes a column vector V and an element $\beta \in K$ such that $\mathfrak{a} = \beta H V$ and $\beta \equiv 1 \pmod{^*\mathfrak{m}}$.

1. [Work in $Cl(K)$] Applying the principal ideal algorithm in $Cl(K)$, find a column vector W and $\gamma \in K$ such that $\mathfrak{a} = \gamma G W$ (where G is the row vector of the ideals g_i whose classes are the given generators of $Cl(K)$).
2. [Work in $(\mathbb{Z}_K/\mathfrak{m})^*$] Set $\alpha \leftarrow \gamma/A'W$, where the $A' = (\alpha_i')$ is the row vector of elements computed in step 1 of Algorithm 4.3.1 (α will be coprime to $\mathfrak{m}$). Using Algorithm 4.2.24, compute Y such that $\alpha \equiv ZY \pmod{^*\mathfrak{m}}$, and let $\alpha' \leftarrow ZY$ as an element of $\mathbb{Z}_K$.

3. [Terminate] Let U_a be the matrix output by Algorithm 4.3.1. Output $V \leftarrow U_a\left(\frac{Y}{W}\right)$ and $\beta \leftarrow \alpha/\alpha'$, and terminate the algorithm.

Remark. It is essential that the generators Z of $(\mathbb{Z}_K/\mathfrak{m})^*$ used in this algorithm be the same as those used in Algorithm 4.3.1.

This finishes the description of the algorithmic computation of the ray class groups $Cl_{\mathfrak{m}}(K)$. It should be emphasized that although many algorithms and subalgorithms are involved, the basic computations are rather simple and the main bottlenecks will be in two places. The first will be in the computation of discrete logarithms in $(\mathbb{Z}_K/\mathfrak{p})^*$. For this, considering the vast amount of effort spent on the problem, we have nothing more to say.

The second bottleneck will be the size of the generators. Indeed, several times we have to multiply a given set of generators by a unimodular matrix U^{-1}, or multiply generators by elements to make them coprime to certain ideals. All this makes the coefficients of the generators grow in size. Since this can rapidly make the algorithms completely useless in practice, we would like to give a few indications on how to get down to generators of manageable size.

4.3.2 Size Reduction of Elements and Ideals

The main place where size reduction is necessary is in Algorithm 4.1.3, that is, in the SNF algorithm for Abelian groups. Recall that in this algorithm, a system of generators and relations (G, M) is given, and after reducing M to its HNF H, which is generally a harmless process, we use the SNF algorithm to compute unimodular matrices U and V such that $UHV = D$ (and afterwards we remove the trivial components). The main difficulty comes from the fact that the new generators are given essentially by GU^{-1}, and these may be large objects if U^{-1} has large entries.

There are several complementary ways to improve this situation, and all should be applied.

(1) The matrix U^{-1} is not unique in general; hence, it is worthwhile to find a matrix U^{-1} that is as small as possible. This can be done using the techniques of [Hav-Maj2]. In many cases this just cannot be done, however, and all possible matrices U^{-1} have large entries.

(2) Another idea is to observe that $GM = \mathbf{1}$ in the Abelian group; hence, if we add to the columns of U^{-1} any $\mathbb{Z}$-linear combination of the columns of M (or of H), the resulting generators GU^{-1} are unchanged. The simplest way for doing this reduction is probably as follows. Let X be a column vector that we want to reduce modulo the columns of H. First compute the matrix L obtained by applying the LLL algorithm to the columns of H. Then replace X by $X - L\lfloor L^{-1}X \rceil$, where $\lfloor A \rceil$ denotes the result of rounding each entry of a matrix to the nearest integer. This should now be rather small.

(3) We should try to avoid divisions as much as possible, since they are generally expensive operations. For this, instead of computing a product of the form $\prod_i g_i^{u_i}$ in the naive way, we write

$$\prod_i g_i^{u_i} = \prod_{i,\, u_i>0} g_i^{u_i} / \prod_{i,\, u_i<0} g_i^{-u_i} ,$$

so that we need to perform only one division.

(4) In the (very frequent) case where the group consists of classes of elements of a set modulo some equivalence relation, the elements of the group are usually given by the classes of some representatives, but the latter should be chosen with care. In other words, we should try to reduce modulo the equivalence relation as much as possible.
Let us look in detail at the two cases of importance to us; that of $(\mathbb{Z}_K/\mathfrak{m})^*$ and that of $Cl_{\mathfrak{m}}(K)$.

a) Recall that elements of $(\mathbb{Z}_K/\mathfrak{m})^*$ are represented by pairs $(\overline{\alpha}, v)$ with $\alpha \in \mathbb{Z}_K$ coprime to $\mathfrak{m}_0$ and $v \in \mathbb{F}_2^{\mathfrak{m}_\infty}$. To reduce such a pair, we consider α represented by a column vector X on a fixed integral basis. As in (2) we compute an LLL-reduced basis L of the ideal $\mathfrak{m}_0$, and set $Y \leftarrow X - L\lfloor L^{-1}X \rceil$ (see Algorithm 1.4.13). This will be a reasonably small vector giving an element β congruent to α modulo $\mathfrak{m}_0$. We can then replace $(\overline{\alpha}, v)$ by $(\overline{\beta}, v)$. This is where the two-element representation is the most useful since we do not have to worry about the signature of β.

b) A simple but very important remark is that if m is the smallest positive integer belonging to $\mathfrak{m}_0$ (the upper-left entry in the HNF representation), we can reduce all the coefficients of β modulo m. This can easily be done because we use the two-element representation of elements of $(\mathbb{Z}/\mathfrak{m})^*$; if we had used the one-element representation, it could not have been done so easily.

c) To reduce an ideal representing some ideal class in $Cl_{\mathfrak{m}}(K)$, we proceed as follows. First, exactly as in the case of $(\mathbb{Z}_K/\mathfrak{m})^*$, instead of representing ideal classes as classes of ideals coprime to $\mathfrak{m}_0$ modulo $P_{\mathfrak{m}}$, we will represent them as pairs $(\mathfrak{a}, v)$, where $\mathfrak{a} \in I_{\mathfrak{m}_0}$ is an ideal coprime to $\mathfrak{m}_0$ and $v \in \mathbb{F}_2^{\infty}$ as usual. The equivalence relation $\mathcal{R}$ on these pairs is defined by $(\mathfrak{a}', v') \mathcal{R} (\mathfrak{a}, v)$ if and only if there exists $\beta \equiv 1 \pmod{{}^*\mathfrak{m}_0}$ such that $v' = v + s(\beta)$. As in the case of $(\mathbb{Z}_K/\mathfrak{m})^*$, this representation avoids annoying problems due to signatures.

We will use the following basic algorithm.

Algorithm 4.3.3 (Reduction of an Ideal). Let $\mathfrak{a}$ and $\mathfrak{b}$ be coprime integral ideals. This algorithm computes an element $\gamma \in \mathfrak{a}$ such that $\gamma \equiv 1 \pmod{\mathfrak{b}}$ and $\mathfrak{a}/\gamma$ an LLL-reduced ideal, in the sense of [Coh0, Section 6.5.1].

1. [LLL-reduce] Let α be the first element of an LLL-reduced basis of the ideal product $\mathfrak{ab}$ for the quadratic form $\sum|\sigma_i(\alpha)|^2$ (see step 2 of [Coh0, Algorithm 6.5.5]). If $\mathfrak{b}=\mathbb{Z}_K$, output α and terminate the algorithm.
2. [Use Algorithm 1.3.2] Using Algorithm 1.3.2, compute $a\in\mathfrak{a}$ and $b\in\mathfrak{b}$ such that $a+b=1$.
3. [Terminate] Compute the element $\alpha'\leftarrow a/\alpha$, and let q be the element obtained by rounding to the nearest integer the coefficients of α' on the integral basis. Output $\gamma\leftarrow a-q\alpha$ and terminate the algorithm.

Using this algorithm, we can now write an algorithm for reducing a representative of an ideal class modulo $\mathfrak{m}$.

Algorithm 4.3.4 (Reduction of the Representative of a Ray Ideal Class). Given a modulus $\mathfrak{m}=\mathfrak{m}_0\mathfrak{m}_\infty$ and an element of $Cl_{\mathfrak{m}}(K)$ represented by a pair $(\mathfrak{a},v)$ as above, this algorithm computes another representative $(\mathfrak{a}',v')$ of the same class in $Cl_{\mathfrak{m}}(K)$ such that $\mathfrak{a}'$ is an "almost-reduced" integral ideal.

1. [Use Algorithm 4.3.3] Using Algorithm 4.3.3 applied to $\mathfrak{a}$ and $\mathfrak{b}=\mathfrak{m}_0$, compute $\gamma\in\mathfrak{a}$ such that $\gamma\equiv 1\pmod{\mathfrak{m}_0}$ and $\mathfrak{a}/\gamma$ is an LLL-reduced ideal.
2. [Use Algorithm 4.3.3 again] Again using Algorithm 4.3.3, but this time applied to $\gamma/\mathfrak{a}$ and $\mathfrak{m}_0$ (which are integral coprime ideals), compute $\delta\in\gamma/\mathfrak{a}$ such that $\delta\equiv 1\pmod{\mathfrak{m}_0}$ and $(\gamma/\mathfrak{a})/\delta$ is an LLL-reduced ideal. Set $\alpha\leftarrow\delta/\gamma$.
4. [Terminate] Set $\mathfrak{a}'\leftarrow\alpha\mathfrak{a}$ and $v'\leftarrow v+s(\alpha)$, output the pair $(\mathfrak{a}',v')$, and terminate the algorithm.

The proof of these algorithms' validity is trivial and is left to the reader (Exercise 25). □

4.4 Computations in Class Field Theory

Thanks to the above algorithms, we have complete control on the ray class groups $Cl_{\mathfrak{m}}(K)$. Let us look at what remains to be done to put the main results of class field theory in algorithmic form.

4.4.1 Computations on Congruence Subgroups

First of all, we must enumerate congruence subgroups C modulo $\mathfrak{m}$ or, equivalently, subgroups $\overline{C}$ of $Cl_{\mathfrak{m}}(K)$. This is, of course, done by enumerating HNF left divisors of the SNF of $Cl_{\mathfrak{m}}(K)$, as explained in Section 4.1.10. Usually, $Cl_{\mathfrak{m}}(K)$ does not have too many cyclic components, so this computation is not difficult in practice. In addition, if we are interested only in subgroups of given index, corresponding to Abelian extensions L/K of given degree, the enumeration is much simpler in general, as can be seen, for example, in Algorithm 4.1.20.

Thus, if the SNF of $Cl_{\mathfrak{m}}(K)$ is equal to (A, D_A) and the HNF left divisor of D_A corresponding to the subgroup $\overline{C}$ is equal to H_A, we will represent the congruence subgroup $(\mathfrak{m}, C)$ by the triplet (A, D_A, H_A). The following algorithm, which is a reformulation in our special case of Algorithm 4.1.10, shows how to go from a modulus to a divisor.

Algorithm 4.4.1 (Computation of $CP_{\mathfrak{n}}$). Let $(\mathfrak{m}, C)$ be a congruence subgroup, and let $\mathfrak{n}$ be a divisor of $\mathfrak{m}$ such that $I_{\mathfrak{m}} \cap P_{\mathfrak{n}} \subset C$, so that by Proposition 3.3.5 we have $(\mathfrak{m}, C) \sim (\mathfrak{n}, CP_{\mathfrak{n}})$. Let $Cl_{\mathfrak{m}}(K) = (A, D_A)$ and $Cl_{\mathfrak{n}}(K) = (B, D_B)$ be the respective SNFs, with $A = (\overline{\mathfrak{a}_i})$, and let H_A the HNF left divisor of D_A representing the subgroup $\overline{C}$ of $Cl_{\mathfrak{m}}(K)$. This algorithm computes the HNF left divisor H_B of D_B representing the subgroup $\overline{CP_{\mathfrak{n}}}$ of $Cl_{\mathfrak{n}}(K)$.

1. [Compute matrix P] Using Algorithm 4.3.2, compute the discrete logarithms of the ideals $\mathfrak{a}_i$ in $Cl_{\mathfrak{n}}(K)$, thus obtaining a matrix P such that $s_{\mathfrak{m},\mathfrak{n}}(A) = BP$, where $s_{\mathfrak{m},\mathfrak{n}}$ is the canonical surjection from $Cl_{\mathfrak{m}}(K)$ to $Cl_{\mathfrak{n}}(K)$.
2. [Compute H_B] Let $M \leftarrow (PH_A | D_B)$. Compute the HNF H_B of M, output H_B, and terminate the algorithm.

We also need to compute the conductor of the congruence subgroup $(\mathfrak{m}, C)$. This is done by applying Corollary 3.3.13 as follows.

Algorithm 4.4.2 (Conductor of a Congruence Subgroup). Let $(\mathfrak{m}, C)$ be a congruence subgroup. This algorithm computes the conductor $\mathfrak{f}$ of $(\mathfrak{m}, C)$. Recall that for any congruence subgroup $(\mathfrak{n}, D)$ we denote by $h_{\mathfrak{n},D}$ the cardinality of the group $Cl_{\mathfrak{n}}/\overline{D}$.

1. [Initialize] Set $\mathfrak{f} \leftarrow \mathfrak{m}$, $D \leftarrow C$, $h \leftarrow h_{\mathfrak{f},D}$.
2. [Loop] For each $\mathfrak{p} \mid \mathfrak{f}$ (finite or infinite), compute $D_{\mathfrak{p}} \leftarrow DP_{\mathfrak{f}/\mathfrak{p}}$ using Algorithm 4.4.1, compute $h_{\mathfrak{p}} \leftarrow h_{\mathfrak{f}/\mathfrak{p},D_{\mathfrak{p}}}$, and test whether $h_{\mathfrak{p}} = h$. If this is true for some $\mathfrak{p}$, set $\mathfrak{f} \leftarrow \mathfrak{f}/\mathfrak{p}$, $D \leftarrow D_{\mathfrak{p}}$, $h \leftarrow h_{\mathfrak{p}}$, and go to step 2.
3. [Terminate] Output $\mathfrak{f}$ (and $D = CP_{\mathfrak{f}}$ if desired) and terminate the algorithm.

Proof. If $h_{\mathfrak{f}/\mathfrak{p},DP_{\mathfrak{f}/\mathfrak{p}}} = h_{\mathfrak{f},D}$, then $\mathfrak{f}$ is not the conductor by Corollary 3.3.13, and by Proposition 3.3.6 we have $(\mathfrak{f}/\mathfrak{p}, DP_{\mathfrak{f}/\mathfrak{p}}) \sim (\mathfrak{f}, D)$, so we can replace $\mathfrak{f}$ by $\mathfrak{f}/\mathfrak{p}$. Conversely, if for all $\mathfrak{p}$ we have $h_{\mathfrak{f}/\mathfrak{p},DP_{\mathfrak{f}/\mathfrak{p}}} < h_{\mathfrak{f},D}$, then Corollary 3.3.13 tells us that $\mathfrak{f}$ is the conductor. □

Remark. If we do not need to compute the conductor but simply need to check whether or not $\mathfrak{m}$ is equal to the conductor, we exit the algorithm as soon as we find some $\mathfrak{p} \mid \mathfrak{m}$ such that $h_{\mathfrak{p}} = h$.

4.4.2 Computations on Abelian Extensions

Consider now the other side of class field theory: in other words, isomorphism classes of Abelian extensions L/K.

We first want to compute the *norm group*, that is, the Artin or Takagi group corresponding to a modulus $\mathfrak{m}$. This is done using Theorem 3.4.4 as follows.

Algorithm 4.4.3 (Computation of the Norm Group). Let L/K be an Abelian extension defined by an irreducible monic polynomial $T \in \mathbb{Z}_K[X]$, and let $\mathfrak{m}$ be a modulus of K known to be a multiple of the conductor of L/K. We assume that the SNF of the ray class group $Cl_{\mathfrak{m}}(K) = (C, D_C)$ has already been computed. This algorithm computes the norm group $T_{\mathfrak{m}}(L/K) = A_{\mathfrak{m}}(L/K)$ as a subgroup of $Cl_{\mathfrak{m}}(K)$; in other words, it outputs an HNF matrix that is the left divisor of D_C corresponding to this subgroup.

1. [Initialize] Set $n \leftarrow [L : K]$, $M \leftarrow D_C$, $\mathfrak{d} \leftarrow \mathrm{disc}(T)$, $p \leftarrow 0$, $g \leftarrow 0$, $i \leftarrow 0$.
2. [Finished?] If $\det(M) = n$, output M and terminate the algorithm.
3. [Next $\mathfrak{p}$] If $i < g$, set $i \leftarrow i + 1$. Otherwise, replace p by the smallest prime number strictly greater than p. Using [Coh0, Algorithm 6.2.9], factor $p\mathbb{Z}_K$ into a power product of prime ideals $(\mathfrak{p}_i)_{1 \le i \le g}$ (the exponents e_i are irrelevant), and set $i \leftarrow 1$. Finally, set $\mathfrak{p} \leftarrow \mathfrak{p}_i$.
4. [Factor $\mathfrak{p}\mathbb{Z}_L$] If $\mathfrak{p} \mid \mathfrak{d}$ or $\mathfrak{p} \mid \mathfrak{m}$, go to step 3. Otherwise, let $\overline{T(X)} = \prod_{1 \le j \le g} \overline{T_j(X)}$ be the factorization of $\overline{T(X)}$ into distinct, monic, irreducible polynomials in $(\mathbb{Z}_K/\mathfrak{p})[X]$. There will be no repeated factors, and all the T_j will have the same degree; call it f.
5. [Compute discrete logarithm] Let L be the discrete logarithm of $\mathfrak{p}$ on the given generators of $Cl_{\mathfrak{m}}(K)$, computed using Algorithm 4.3.2. Set M equal to the Hermite normal form of the horizontal concatenation $(M|fL)$ of M with the one-column matrix (fL), and go to step 2.

Proof. We note that fL as computed in step 5 is the discrete logarithm of $\mathfrak{p}^f$ on the generators C. Hence by Theorem 3.4.4, it corresponds to an element of the norm group $T_{\mathfrak{m}}(L/K)$ expressed on the generators. Thus, the successive matrices M represent successively larger subgroups of $Cl_{\mathfrak{m}}(K)$ (equivalently, $\det(M)$ *decreases*), all contained in the norm group. Since we know that the norm group is generated by the $\mathfrak{p}^f$, we will obtain the norm group after a finite number of steps, characterized by $\det(M) = [L : K] = n$ by Proposition 4.1.6 (3). □

Remark. In step 4, we have removed prime ideals $\mathfrak{p}$ such that $\mathfrak{p} \mid \mathfrak{d}$ and $\mathfrak{p} \mid \mathfrak{m}$. This has two purposes. First, it removes prime ideals dividing $\mathfrak{m}$ and in particular ramified prime ideals, which is necessary for Theorem 3.4.4. But also, $\mathfrak{p}$ will not divide the index $(v_{\mathfrak{p}}(\mathfrak{d}) = v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = 0)$, so we are in the easy case of Algorithm 2.4.13, where we simply need to factor $\overline{T(X)}$ in $(\mathbb{Z}_K/\mathfrak{p})[X]$. In fact, since we need only to compute the common degree f of the irreducible factors of $\overline{T}$, we can simply use the distinct degree factorization algorithm [Coh0, Algorithm 3.4.3], where we replace p by $q = |\mathbb{Z}_K/\mathfrak{p}|$ without actually finding the factors.

It is now easy to compute the conductor of a finite Abelian extension L/K.

Algorithm 4.4.4 (Conductor of an Abelian Extension). Let L/K be an Abelian extension defined by an irreducible monic polynomial $T \in \mathbb{Z}_K[X]$. This algorithm computes the conductor $\mathfrak{f}(L/K)$ and the corresponding norm group $T_{\mathfrak{f}}(L/K) = A_{\mathfrak{f}}(L/K)$.

1. [Compute $\mathfrak{d}(L/K)$] Using Algorithm 2.4.9 and its subalgorithms, compute the relative discriminant ideal $\mathfrak{d}(L/K)$.
2. [Compute ramified real places] Set $\mathfrak{m}_\infty \leftarrow \emptyset$, and for each real embedding σ_i of K, using Sturm's algorithm ([Coh0, Algorithm 4.1.11]) test whether T^{σ_i} has only real roots. For each i for which this is not the case, set $\mathfrak{m}_\infty \leftarrow \mathfrak{m}_\infty \cup \{\sigma_i\}$.
3. [Compute norm group] Set $\mathfrak{m} \leftarrow \mathfrak{d}(L/K)\mathfrak{m}_\infty$. Using Algorithm 4.4.3 above, let $C \leftarrow T_{\mathfrak{m}}(L/K)$.
4. [Compute conductor] Using Algorithm 4.4.2, compute the conductor $(\mathfrak{f}, CP_{\mathfrak{f}})$ of $(\mathfrak{m}, C)$, output $\mathfrak{f}(L/K) \leftarrow \mathfrak{f}$, $T_{\mathfrak{f}}(L/K) \leftarrow CP_{\mathfrak{f}}$, and terminate the algorithm.

Proof. Since $\mathfrak{m}$ as defined in the algorithm is a multiple of the conductor, the algorithm's validity is a simple consequence of the (deep) result asserting that the conductor of a congruence subgroup $(\mathfrak{m}, C)$ is equal to the conductor of the corresponding Abelian extension L/K (Theorem 3.4.6). □

Remark. We could modify the algorithm by taking $\mathrm{disc}(T)$ instead of $\mathfrak{d}(L/K)$, which avoids the round 2 algorithm, at the cost of more class group computations in step 4. Hence, it is not clear whether this gives any improvement.

It is interesting to note that the above algorithms can also be used as an efficient test to determine whether or not an arbitrary extension of number fields L/K is Abelian. For this, we must first modify Algorithm 4.4.3 so that it can still work for an arbitrary extension.

Algorithm 4.4.5 (Norm Group or Non-Abelian Extension). Let L/K be an extension of number fields defined by an irreducible monic polynomial $T \in \mathbb{Z}_K[X]$. Let $\mathfrak{m}$ be a modulus of K known to be a multiple of the conductor of L^{ab}/K, where L^{ab} is the maximal Abelian subextension of L/K. We assume that the SNF of the ray class group $Cl_{\mathfrak{m}}(K) = (C, D_C)$ has already been computed. This algorithm either outputs a failure message indicating under the GRH that L/K is not Abelian or unconditionally computes the norm group $T_{\mathfrak{m}}(L/K) = A_{\mathfrak{m}}(L/K)$ as a subgroup of $Cl_{\mathfrak{m}}(K)$. In other words, it outputs an HNF matrix that is a left divisor of D_C corresponding to this subgroup.

1. [Initialize] Set $n \leftarrow [L : K]$, $M \leftarrow D_C$, $\mathfrak{d} \leftarrow \mathrm{disc}(T)$, $p \leftarrow 0$, $g \leftarrow 0$, $i \leftarrow 0$ and $B \leftarrow (4\log(|d(L)|) + 2.5[L : \mathbb{Q}] + 5)^2$.

2. [Finished?] If $p > B$, do as follows. If $\det(M) \neq n$, output a failure message (L/K is not an Abelian extension), while if $\det(M) = n$, output M. In either case, terminate the algorithm.
3. [Next $\mathfrak{p}$] If $i < g$, set $i \leftarrow i + 1$. Otherwise, replace p by the smallest prime number strictly greater than p. Using [Coh0, Algorithm 6.2.9], factor $p\mathbb{Z}_K$ into a power product of prime ideals $(\mathfrak{p}_i)_{1 \leq i \leq g}$ (the exponents e_i are irrelevant), and set $i \leftarrow 1$. Finally, set $\mathfrak{p} \leftarrow \mathfrak{p}_i$.
4. [Factor $\mathfrak{p}\mathbb{Z}_L$] If $\mathfrak{p} \mid \mathfrak{d}$ or $\mathfrak{p} \mid \mathfrak{m}$, go to step 3. Otherwise, let $\overline{T(X)} = \prod_{1 \leq j \leq g} \overline{T_j(X)}$ be the factorization of $\overline{T(X)}$ into distinct, monic, irreducible polynomials in $(\mathbb{Z}_K/\mathfrak{p})[X]$. There will be no repeated factors. If all the T_j do not have the same degree, then output a failure message (L/K is not a normal extension) and terminate the algorithm. Otherwise, let f be the common degree of the T_j.
5. [Compute discrete logarithm] Let L be the discrete logarithm of $\mathfrak{p}$ on the given generators of $Cl_{\mathfrak{m}}(K)$, computed using Algorithm 4.3.2. Set M equal to the Hermite normal form of the horizontal concatenation $(M|fL)$ of M with the one-column matrix (fL). If $\det(M) < n$, output a failure message (L/K is not an Abelian extension) and terminate the algorithm; otherwise go to step 2.

Proof. As this algorithm is essentially identical to Algorithm 4.4.3, we need only to discuss the cases of failure. If the failure occurs in step 4 or in step 5, then we can unconditionally assert that the extension L/K is not Abelian (and even not normal if the failure is in step 4). If the failure occurs because $p > B$ in step 2, the situation is different.

A result of Bach and Sorenson [Bac-Sor] implies that, under the GRH, the norm group will be generated by prime ideals of norm less than or equal to the bound B computed in step 1. Thus, if the GRH is true, the primes up to the bound B are sufficient, and hence the algorithm is correct as is (and the extension is Abelian if $\det(M) = n$). Thus, the correctness of the algorithm is unconditional for the failure in steps 4 and 5 and is valid only under the GRH for step 2. If we do not want to assume the GRH, we can increase the bound B, but as usual we will have a much larger bound, of the order of $|d(L)|^{1/2}$. □

The modification of Algorithm 4.4.4 is now immediate.

Algorithm 4.4.6 (Is an Extension Abelian?). Let L/K be an extension of number fields defined by an irreducible monic polynomial $T \in \mathbb{Z}_K[X]$. This algorithm determines under the GRH whether or not L/K is an Abelian extension. If it is, it computes the Galois group $G(L/K)$, the conductor $\mathfrak{f}(L/K)$ and the corresponding norm group $T_{\mathfrak{f}}(L/K) = A_{\mathfrak{f}}(L/K)$.

1. [Compute $\mathfrak{d}(L/K)$] Using Algorithm 2.4.9 and its subalgorithms, compute the relative discriminant ideal $\mathfrak{d}(L/K)$.

2. [Compute ramified real places] Set $\mathfrak{m}_\infty \leftarrow \varnothing$, and for each real embedding σ_i of K, using Sturm's algorithm ([Coh0, Algorithm 4.1.11]) test whether T^{σ_i} has only real roots. For each i for which this is not the case, do as follows. Test whether all the roots are nonreal. If at least one root is real, L/K is not a normal extension, so terminate the algorithm. Otherwise, set $\mathfrak{m}_\infty \leftarrow \mathfrak{m}_\infty \cup \{\sigma_i\}$.
3. [Compute norm group] Set $\mathfrak{m} \leftarrow \mathfrak{d}(L/K)\mathfrak{m}_\infty$. Using Algorithm 4.3.1 compute the ray class group $Cl_{\mathfrak{m}}(K)$, then execute Algorithm 4.4.5 above. If the algorithm fails, L/K is not an Abelian extension, so terminate the algorithm. Otherwise, set $C \leftarrow T_{\mathfrak{m}}(L/K)$ as computed by the algorithm.
4. [Compute conductor] Using Algorithm 4.4.2, compute the conductor $(\mathfrak{f}, CP_{\mathfrak{f}})$ of $(\mathfrak{m}, C)$, output a message saying that L/K is Abelian with Galois group isomorphic to $Cl_{\mathfrak{m}}(K)/\overline{C}$, output $\mathfrak{f}(L/K) \leftarrow \mathfrak{f}$ and $T_{\mathfrak{f}}(L/K) \leftarrow CP_{\mathfrak{f}}$, and terminate the algorithm.

This algorithm thus gives an answer to the question of whether or not L/K is Abelian, sometimes unconditionally (steps 4 and 5 of Algorithm 4.4.5 or step 2 of the above algorithm), and sometimes conditionally under the GRH (step 2 of Algorithm 4.4.5). Since we can have good confidence in the validity of the GRH, if the result is conditional, we can assume that the conclusion of the algorithm is probably true and then proceed to *prove* it using other methods.

Finally, recall that Proposition 3.5.8 and Theorem 3.5.11 give us efficient formulas allowing us to compute the signature and the relative or absolute discriminant of L/K. Thus, by far the most important point that we have not solved is the computation of an explicit relative (or absolute) defining polynomial for the Abelian extension L/K corresponding to the (equivalence class of the) congruence subgroup $(\mathfrak{m}, C)$. This will be considered in Chapters 5 and 6.

4.4.3 Conductors of Characters

The formulas given in Theorem 3.5.11 have the great advantage that we do not need to compute the conductors of individual characters. In this subsection, we explain how to do this if these conductors are really needed.

Let

$$Cl_{\mathfrak{m}}(K) = (G, D_G) = \bigoplus_{1 \le i \le k} (\mathbb{Z}/d_i\mathbb{Z})g_i$$

be the SNF of $Cl_{\mathfrak{m}}(K)$. Denote by ζ_n the specific primitive nth root of unity $\exp(2i\pi/n)$ and let $\zeta = \zeta_{d_1}$ (recall that d_i divides d_1 for all i). A character χ is uniquely defined by a vector $(a_1, \ldots, a_k)$ with $a_i \in \mathbb{Z}/d_i\mathbb{Z}$, so that

$$\chi\Big(\prod_i g_i^{x_i}\Big) = \prod_i \zeta_{d_i}^{a_i x_i} = \zeta^{\sum_i (d_1/d_i) a_i x_i} .$$

By definition, the conductor of χ is equal to the conductor of the congruence subgroup $C = \mathrm{Ker}(\chi)$. Since this is a congruence subgroup, we can use the above methods to compute its conductor. The only problem is to put this group into an algorithmic form, in other words to compute the matrix H associated to C by Proposition 4.1.6.

We have $\chi(\prod_i g_i^{x_i}) = 1$ if and only if there exists an integer y such that

$$\sum_i \frac{d_1}{d_i} a_i x_i + d_1 y = 0 \ .$$

This is an instance of the integer kernel problem. We have seen in [Coh0, Section 2.4.3] and in Section 4.1.6 how to solve it. In the present case, this gives the following.

Set $b_i = (d_1/d_i)a_i$, and let $B = [b_1, \ldots, b_k, d_1]$, considered as a one-row matrix. Using the Hermite normal form algorithm, we can compute a unimodular matrix U such that $BU = [0, \ldots, 0, d]$ for some d (equal to the GCD of the entries of B). Write in block matrix form $U = \left(\begin{smallmatrix} U_1 & V \\ R & a \end{smallmatrix}\right)$, where U_1 is a $k \times k$ matrix, V is a one-column matrix, and R is a one-row matrix. The column vectors $X = (x_i)$ such that there exists a y satisfying our equality above are then exactly the $\mathbb{Z}$-linear combinations of the columns of the matrix U_1. This means that the kernel of χ is defined by the matrix U_1, or if we want it in normalized form, by the HNF of $(U_1|D_G)$. We can then compute the conductor as usual.

Formally, this can be written as follows.

Algorithm 4.4.7 (Conductor of a Character). Let

$$Cl_{\mathfrak{m}}(K) = (G, D_G) = \bigoplus_{1 \le i \le k} (\mathbb{Z}/d_i\mathbb{Z})g_i$$

be the SNF of the ray class group $Cl_{\mathfrak{m}}(K)$, and let χ be a character defined by $\chi(\prod_i g_i^{x_i}) = \prod_i \zeta_{d_i}^{a_i x_i}$. This algorithm computes the conductor of χ (which is a modulus of K).

1. [Apply HNF] For all $i \le k$, set $b_i \leftarrow (d_1/d_i)a_i$, and set $B \leftarrow (b_1, \ldots, b_k, d_1)$, considered as a one-row matrix. Using the HNF algorithm, find a unimodular matrix U such that $BU = (0, \ldots, 0, d)$.
2. [Compute H] Let U_1 be the upper-left $k \times k$ submatrix of U, let H be the HNF of $(U_1|D_G)$, and call C the corresponding congruence subgroup.
3. [Terminate] Using Algorithm 4.4.2, compute the conductor $\mathfrak{f}$ of the congruence subgroup $(\mathfrak{m}, C)$, output $\mathfrak{f}$, and terminate the algorithm.

4.5 Exercises for Chapter 4

1. Using Algorithm 4.1.11, give an algorithm for computing the group $U_{\mathfrak{m}}(K)$ of units congruent to 1 (mod $^*\mathfrak{m}$) as a subgroup of $U(K)$.

2. Prove the validity of the remark made after Algorithm 4.1.8.
3. Let $\mathcal{B} = (B, D_B)$ and $\mathcal{C} = (C, D_C)$ be two known Abelian groups in SNF, let $\mathcal{A}$ be a subgroup of $\mathcal{C}$ given by a left divisor H_C of D_C, and let ϕ be an effective group homomorphism from $\mathcal{B}$ to $\mathcal{C}/\mathcal{A}$. Show that we can use Algorithm 4.1.11 to compute $\mathrm{Ker}(\phi)$ if in step 1 we simply replace $\phi(B) = CP$ by $\phi(B) = \pi(C)P$, where π denotes the canonical surjection from $\mathcal{C}$ to $\mathcal{C}/\mathcal{A}$.
4. Let $\mathcal{B} = (B, D_B)$ and $\mathcal{C} = (C, D_C)$ be two known Abelian groups and let ϕ be an effective group homomorphism from $\mathcal{B}$ to $\mathcal{C}$. Give an algorithm for computing the *cokernel* of ϕ, in other words the quotient $\mathcal{C}/\phi(\mathcal{B})$.
5. With the notation of Lemma 4.1.12, give an explicit formula for U_1^{-1} in terms of the block matrix decomposition of U^{-1}.
6. With the notation of the proof of Algorithm 4.1.13, show that if $\psi(\alpha) = BY$, the vector $H_B^{-1}Y$ has integral entries.
7. Prove the validity of Algorithm 4.1.14.
8. Prove the validity of Algorithm 4.1.15.
9. Let $D_C = \mathrm{diag}(c_1, \ldots, c_n)$ be a diagonal matrix in SNF, and let $H = (e_{i,j})$ be an $n \times n$ matrix in HNF.
   a) If $n = 2$, show that H is a left divisor of D_C if and only if $e_{i,i} \mid c_i$ for $i = 1$ and $i = 2$, and if $e_{1,2} = ke_{1,1}/\gcd(e_{1,1}, c_2/e_{2,2})$ with $0 \leq k < \gcd(e_{1,1}, c_2/e_{2,2})$.
   b) If $n \geq 3$, show that for all $i \leq n$ we must have $e_{i,i} \mid c_i$, and for all $i < n$

$$e_{i,i+1} \equiv 0 \pmod{e_{i,i}/\gcd(e_{i,i}, c_{i+1}/e_{i+1,i+1})} \ ,$$

   but that these conditions are not sufficient.
10. Write and implement a formal algorithm for computing all subgroups of a given algorithm using Birkhoff's Theorem 4.1.18. In particular, determine whether it is more efficient to choose first the y_i and then the permutation, as written in the text, or to do the reverse.
11. Prove Proposition 4.1.19.
12. Give a complete description of the subgroups of index n of a given Abelian group, in the style of Proposition 4.1.19, for $n = 4$ and $n = 6$, and more generally for $n = p^2$ and $n = pq$ when p and q are primes, and write the corresponding algorithms analogous to Algorithm 4.1.20.
13. Prove the validity of Algorithm 4.1.21.
14. Write an algorithm for solving a mixed system of linear equations and linear congruences by first solving the linear equations, and plugging the result into the linear congruences, instead of using Algorithm 4.1.23 given in the text. Compare the efficiency of both algorithms.
15. Prove the validity of Algorithm 4.1.23.
16. Extend Proposition 1.2.11 and Lemma 4.2.1 to the case where $\mathfrak{a}$ and $\mathfrak{c}$ are coprime moduli, in other words to the case where $\mathfrak{a}_0 + \mathfrak{c}_0 = \mathbb{Z}_K$ and $\mathfrak{a}_\infty \cap \mathfrak{c}_\infty = \varnothing$.
17. Show that $v_p(i!) \leq (i-1)/(p-1)$, and determine exactly the cases where there is equality.
18. Let $\mathfrak{a}$ be an integral ideal of a number field K.
   a) Show that $\mathbb{Z}_K/\mathfrak{a}$ is cyclic if and only if every prime ideal $\mathfrak{p}$ dividing $\mathfrak{a}$ has residual degree equal to 1, every prime ideal $\mathfrak{p}$ such that $\mathfrak{p}^2 \mid \mathfrak{a}$ is unramified and if $\mathfrak{p}$ and $\mathfrak{q}$ are distinct prime ideals dividing $\mathfrak{a}$, then $\mathfrak{p}$ and $\mathfrak{q}$ are not above the same prime number of $\mathbb{Z}$.

b) If K is a quadratic field, show that $\mathbb{Z}_K/\mathfrak{a}$ is cyclic if and only if $\mathfrak{a}$ is a primitive ideal (in other words, an integral ideal not divisible by an element of $\mathbb{Z}$ other than ± 1).

c) Let $\mathfrak{p}$ be a prime ideal of K, let p be the prime ideal below $\mathfrak{p}$, and let $e = e(\mathfrak{p}/p)$ and $f = f(\mathfrak{p}/p)$. Show that $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ is cyclic if and only if either $k = 1$; or $k = 2$ and $f = 1$; or $k \geq 3$, $e = f = 1$, and $p \geq 3$; or $k = 3$, $f = 1$, $e \geq 2$, and $p = 2$.

d) Deduce from this a necessary and sufficient condition for $(\mathbb{Z}_K/\mathfrak{a})^*$ to be cyclic, and specialize to the quadratic case.

19. Extend the table of Proposition 4.2.12 up to $k = 9$.

20. Let $\mathfrak{p}$ be a prime ideal above a prime p, of ramification index $e = e(\mathfrak{p}/p)$ and degree $f = f(\mathfrak{p}/p)$, and let $k \geq 1$ be an integer. Denote by $\log_p(x) = \log(x)/\log(p)$ the ordinary logarithm of x to base p. Prove the following strengthening of the first statement of Proposition 4.2.12: the group $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$ is killed by p^s, where the integer s is given as follows.

a) If $e < p-1$, then

$$s = \left\lceil \frac{k-1}{e} \right\rceil .$$

b) If

$$p-1 \leq e < (p-1)p^{\lfloor \log_p k \rfloor} ,$$

then

$$s = \left\lceil \frac{k - p^{\lceil \log_p(e/(p-1)) \rceil}}{e} \right\rceil + \left\lceil \log_p\left(\frac{e}{p-1}\right) \right\rceil .$$

c) If $e \geq (p-1)p^{\lfloor \log_p k \rfloor}$, then

$$s = \lceil \log_p k \rceil .$$

21. Let p be a prime. Define the Artin–Hasse logarithm $\log_a$ by the formula

$$\log_a(1+x) = \sum_{k=1}^{p-1} (-1)^{k-1} \frac{x^k}{k} .$$

a) Using combinatorial identities, show that formally

$$\log_a((1+x)(1+y)) - \log_a(1+x) - \log_a(1+y)$$

is a polynomial whose nonzero monomials are of the form $x^m y^n$ with $m+n \geq p$.

b) Define in a similar manner the Artin–Hasse exponential $\exp_a$ and prove its basic properties and relations with $\log_a$.

c) Deduce from this that if $a < b \leq pa$ and $\mathfrak{p}$ is an ideal above p, the map $(1+x) \mapsto \log_a(x)$ induces a group isomorphism from the multiplicative group $(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$ to the additive group $\mathfrak{p}^a/\mathfrak{p}^b$, and in particular from $(1+\mathfrak{p})/(1+\mathfrak{p}^p)$ to $\mathfrak{p}/\mathfrak{p}^p$ (note that this gives an alternate proof of Corollary 4.2.11 when $p \geq k$).

22. Write and implement an algorithm for computing the group $(\mathbb{Z}_K/\mathfrak{m})^*$ using a combination of $\mathfrak{p}$-adic logarithm techniques and the induction method, as suggested after Algorithm 4.2.17.

23. Assume that $e(\mathfrak{p}/p) = p - 1$ as in Proposition 4.2.19. Prove directly that the exact sequence

$$1 \longrightarrow (1+\mathfrak{p}^2)/(1+\mathfrak{p}^p) \longrightarrow (1+\mathfrak{p})/(1+\mathfrak{p}^p) \longrightarrow (1+\mathfrak{p})/(1+\mathfrak{p}^2) \longrightarrow 1$$

is split.

24. Prove the validity of Algorithm 4.2.22.

25. Prove the validity of Algorithms 4.3.3 and 4.3.4.

26. Let $(\mathfrak{m}_1, C_1)$ and $(\mathfrak{m}_2, C_2)$ be two equivalent congruence subgroups represented by triplets (G_1, D_1, H_1) and (G_2, D_2, H_2) as explained in the text. Give an algorithm that computes the GCD $(\mathfrak{n}, C)$ of these two congruence subgroups in the sense of Proposition 3.3.9.

5. Computing Defining Polynomials Using Kummer Theory

Class field theory deals with Abelian extensions of base fields. It gives complete answers to the existence of Abelian extensions with given relative or absolute discriminants. However, the algorithmic construction of these extensions is not completely straightforward. There are several ways to do this, but at present the most efficient general method is the use of Kummer extensions. In the next chapter, we will describe two other methods using analytic techniques, one using Stark units and Stark's conjecture, the other using complex multiplication. Both of these methods impose restrictions on the base field, but when they are applicable they are much more efficient.

5.1 General Strategy for Using Kummer Theory

If we look at the main theorem of Kummer theory (Theorem 10.2.5), we see that we have at our disposal a powerful tool to construct all Abelian extensions of a base field with given Galois group G, assuming that this base field contains sufficiently many roots of unity (more precisely contains ζ_n, where n is the exponent of G). To be able to use this, we must in general adjoin ζ_n to the base field K, hence take as new base field $K_z = K(\zeta_n)$, and use Kummer theory over K_z. Once the desired Abelian extension L_z/K_z is obtained, we must then come back to the desired Abelian extension L/K, which can be done using several methods. The aim of this chapter is to explain all this in great detail.

5.1.1 Reduction to Cyclic Extensions of Prime Power Degree

Let K be a number field, and let $(\mathfrak{m}, C)$ be a congruence subgroup modulo $\mathfrak{m}$, where we need not assume for the moment that $\mathfrak{m}$ is the conductor. The aim of this chapter is to find an explicit defining polynomial for the extension L/K corresponding to $(\mathfrak{m}, C)$ by Takagi's existence theorem in class field theory. We can easily compute $Cl_{\mathfrak{m}}(K)/\overline{C}$ in SNF as

$$Cl_{\mathfrak{m}}(K)/\overline{C} = \bigoplus_i (\mathbb{Z}/c_i\mathbb{Z})\overline{\mathfrak{c}_i} \ .$$

It is, however, often useful (and, in fact, essential if we use Kummer theory) to split this group even more into its cyclic components of prime power order. This is easily done by using the following equality. If $d = \prod_{1 \le i \le k} d_i$ with the d_i pairwise coprime, then

$$(\mathbb{Z}/d\mathbb{Z})g = \bigoplus_{1 \le i \le k} (\mathbb{Z}/d_i\mathbb{Z})g^{d/d_i}$$

(see Exercise 1), from which it follows that we can write

$$Cl_{\mathfrak{m}}(K)/\overline{C} = \bigoplus_i \bigoplus_{p^{v_p} \| c_i} (\mathbb{Z}/p^{v_p}\mathbb{Z})\overline{\mathfrak{c}_i}^{c_i/p^{v_p}} .$$

With a suitable change of notation, we will write this as

$$Cl_{\mathfrak{m}}(K)/\overline{C} = \bigoplus_{1 \le j \le s} (\mathbb{Z}/b_j\mathbb{Z})\overline{\mathfrak{b}_j} ,$$

where the b_j are (not necessarily distinct) prime powers.

Proposition 5.1.1. *Keep all the above notation. Let C_j be the congruence subgroup modulo $\mathfrak{m}$ generated by C and by the $\mathfrak{b}_i$ for $i \neq j$, and let L_j be the subfield of $K(\mathfrak{m})$ corresponding to the congruence subgroup $(\mathfrak{m}, C_j)$ under the Takagi correspondence.*

(1) *The group $\mathrm{Gal}(L_j/K)$ is isomorphic via the Artin map to $Cl_{\mathfrak{m}}(K)/\overline{C_j} = (\mathbb{Z}/b_j\mathbb{Z})\overline{\mathfrak{b}_j}$, and in particular L_j/K is a cyclic extension of prime power degree b_j.*
(2) *The compositum in $K(\mathfrak{m})$ of the L_j is equal to the class field L corresponding to the congruence subgroup $(\mathfrak{m}, C)$.*

Proof. By Galois theory we have $L_j = K(\mathfrak{m})^{\mathrm{Art}(C_j)}$, hence via the Artin map, $\mathrm{Gal}(L_j/K)$ is isomorphic to $Cl_{\mathfrak{m}}/\overline{C_j} = (\mathbb{Z}/b_j\mathbb{Z})\overline{\mathfrak{b}_j}$, so L_j/K is a cyclic extension of prime power degree b_j, proving the first statement. Furthermore, by Galois theory the compositum of the L_j in $K(\mathfrak{m})$ corresponds to the congruence subgroup $(\mathfrak{m}, \bigcap_j C_j)$. But since the only relations satisfied by the $\overline{\mathfrak{b}_i}$ in $Cl_{\mathfrak{m}}(K)/\overline{C}$ are $\overline{\mathfrak{b}_i}^{b_i} = \overline{1}$, it follows that $\mathfrak{b}_j \notin C_j$, hence that $\bigcap_j C_j = C$, proving the proposition. □

This proposition can be translated into the following algorithm.

Algorithm 5.1.2 (Splitting Class Field Extensions). Let K be a number field and let $(\mathfrak{m}, C)$ be a congruence subgroup modulo $\mathfrak{m}$. This algorithm computes a list of congruence subgroups $(\mathfrak{m}_j, C_j)$ of conductor $\mathfrak{m}_j$ dividing $\mathfrak{m}$ such that the compositum in $K(\mathfrak{m})$ of the class fields L_j corresponding to $(\mathfrak{m}_j, C_j)$ is equal to the class field L corresponding to $(\mathfrak{m}, C)$ and such that the L_j/K are cyclic extensions of prime power degree.

1. [Initializations] Using Algorithm 4.3.1, compute the SNF (A, D_A) of the ray class group $Cl_{\mathfrak{m}}(K)$, and let H_C be the HNF matrix defining the congruence subgroup C on the generators A. Using an SNF algorithm, compute unimodular matrices U_1 and V_1 such that $U_1 H_C V_1 = \mathrm{diag}(c_1, \ldots, c_s)$ is a diagonal matrix in SNF, possibly with ones on the diagonal, and let r be the largest index i such that $c_i > 1$. Finally, let U be the matrix obtained by keeping the first r columns of U_1^{-1} (thus, the ith column U_i of U expresses the ith generator $\overline{\mathfrak{c}_i}$ of $Cl_{\mathfrak{m}}(K)/\overline{C}$ on the generators A of $Cl_{\mathfrak{m}}(K)$).
2. [Split the Galois group] For each $i \leq r$ and each $p^{v_p} \| c_i$, do as follows. Let $U_{i,p}$ be the matrix obtained from U by replacing the ith column U_i by $p^{v_p} U_i \bmod c_i$. Compute the HNF $H_{i,p}$ of $(H_C | U_{i,p})$, which corresponds to a congruence subgroup $(\mathfrak{m}, C'_{i,p})$ modulo $\mathfrak{m}$. Then, using Algorithm 4.4.2, compute the conductor $\mathfrak{m}_{i,p}$ of $(\mathfrak{m}, C'_{i,p})$ and the congruence subgroup $(\mathfrak{m}_{i,p}, C_{i,p})$ equivalent to $(\mathfrak{m}, C'_{i,p})$, output the $(\mathfrak{m}_{i,p}, C_{i,p})$, and terminate the algorithm.

Proof. Write $A = (\overline{\mathfrak{a}_1}, \ldots, \overline{\mathfrak{a}_m})$. Using the remark following Algorithm 4.1.7, we see that the computation done in step 1 gives

$$Cl_{\mathfrak{m}}(K)/\overline{C} = \bigoplus_{1 \leq i \leq r} (\mathbb{Z}/c_i\mathbb{Z})\overline{\mathfrak{c}_i} \ ,$$

where $\overline{\mathfrak{c}_i}$ is given on A by the ith column U_i of the matrix U. Since

$$(\mathbb{Z}/c_i\mathbb{Z}) = \bigoplus_{p^{v_p} \| c_i} (\mathbb{Z}/p^{v_p}\mathbb{Z})\overline{\mathfrak{c}_i^{c_i/p^{v_p}}} \ ,$$

we see that if we set $\overline{\mathfrak{c}_{i,p}} = \overline{\mathfrak{c}_i^{c_i/p^{v_p}}}$, then $\mathfrak{c}_{i,p}$ is given on A by $(c_i/p^{v_p})U_i$. Thus, as a subgroup of $Cl_{\mathfrak{m}}$, the subgroup generated by $\overline{C}$ and by all the $\mathfrak{c}_{j,q}$ except $\mathfrak{c}_{i,p}$ is defined by the HNF of the concatenation of the matrix H_C with the column vectors U_j for $j \neq i$, as well as the column vectors $(c_i/q^{v_q})U_i$ for $q \neq p$. But the GCD of the c_i/q^{v_q} for $q \neq p$ is clearly equal to p^{v_p}, hence the subgroup generated by the column vectors $(c_i/q^{v_q})U_i$ for $q \neq p$ is equal to the subgroup generated by the single vector $p^{v_p}U_i$. Thus, the HNF matrices $H_{i,p}$ computed in step 2 correspond to the desired subgroups given in Proposition 5.1.1, proving the algorithm's validity. We perform additional conductor computations at the end, since in most class field computations it is simpler to start with a congruence subgroup of known conductor, and since these conductor computations are in any case much faster than the ray class field computations themselves. □

Remarks

(1) Thanks to this algorithm, we see that we can always reduce to the case where the desired ray class field extension L/K is cyclic of prime power degree p^r, so that the corresponding congruence subgroup $(\mathfrak{m}, C)$ has

conductor $\mathfrak{m}$ with $Cl_{\mathfrak{m}}(K)/\overline{C}$ cyclic of prime power order. We will make this assumption when using Kummer theory, but not necessarily when using other methods.

(2) It is easy to modify the above algorithm if we want a coarser splitting: for example, if we want to split only according to the c_i, instead of the $U_{i,p}$, for each i we use the single matrix $U_{i,0}$ obtained from U by removing the ith column (see Exercise 2).

5.1.2 The Four Methods

From now on, we have a base field K, a congruence subgroup $(\mathfrak{m}, C)$ of conductor $\mathfrak{m}$ such that $Cl_{\mathfrak{m}}(K)/\overline{C}$ is a cyclic group of order $n = \ell^r$ for some prime number ℓ. Our goal is to use Kummer theory to compute a defining polynomial for the Abelian extension L/K corresponding to $(\mathfrak{m}, C)$ by Takagi's theorem. We refer to Section 10.2 for detailed proofs of the results that we will use.

To be able to use Kummer theory, the base field K must contain ζ_n, a primitive nth root of unity. Thus, we will proceed in two steps. We begin (if necessary) by adjoining ζ_n to K; in other words, we set $K_z = K(\zeta_n)$, we "lift" the problem to K_z, and as a first step we must construct a suitable extension L_z/K_z. As a second step, we must come back down from L_z/K_z to the desired extension L/K.

For both steps, there are essentially two methods. Let L/K be a cyclic extension of degree $n = \ell^r$ for some prime ℓ corresponding to a congruence subgroup $(\mathfrak{m}, C)$ of conductor $\mathfrak{m}$, and assume that $\zeta_n \in K$. Then by definition L/K is a Kummer extension. The main theorem of Kummer theory (Theorem 10.2.5) tells us that $L = K(\theta)$ with $\theta^n = \alpha$ for some $\alpha \in \mathbb{Z}_K$. To apply class field theory to this situation, we have two possibilities.

A first possibility is to use information on the ramification of prime ideals in L/K and the relative discriminant $\mathfrak{d}(L/K)$. Indeed, using Theorems 3.5.3 and 3.5.11, we can easily compute such information from the congruence subgroup $(\mathfrak{m}, C)$. To be able to find a suitable α, we need to compute similar information if the field L is given as $L = K(\theta)$ as above, using only the base field K, the degree $n = \ell^r$, and the element $\alpha \in \mathbb{Z}_K$. This is quite a bit harder and in fact can be done in practice only for $r = 1$, that is, for cyclic extensions of prime degree. This is exactly the content of Hecke's theorem (Theorem 10.2.9). Although any cyclic extension of prime power degree can be considered as a tower of cyclic extensions of prime degree, the need to compute in number fields of much larger degree makes this method unfeasible if $\ell^r > 10$, say.

A second method, introduced by C. Fieker (see [Fie]), is to use directly the properties of the Artin map to construct the needed extension L/K. Of course, the Artin map contains the ramification and discriminant information, but it is in fact richer both in theory and in algorithmic practice. Indeed, Fieker's method has several advantages compared to the method using

Hecke's theorem. The first and most important one is that it is not limited to extensions of prime degree, and the second is that it is not difficult to describe and to implement.

For performing the second step (coming down from the extension L_z/K_z to the extension L/K), there are also two methods. One is the use of so-called Lagrange resolvents, and the other one, also due to C. Fieker, is once again the explicit use of the Artin map. Since both steps are mostly independent, the methods may be mixed if desired.

The main disadvantage of Fieker's methods is the necessity to introduce large moduli and the corresponding ray class groups (but fortunately not explicitly the corresponding ray class fields). Thus, although his method for the first step is usually superior, in some cases and also for the second step the other methods can be better; hence it is interesting to study all the methods. In addition, this study introduces some interesting new concepts such as ℓ-virtual units and the ℓ-Selmer group of a number field.

5.2 Kummer Theory Using Hecke's Theorem When $\zeta_\ell \in K$

Let ℓ be a prime number, and let K be a number field such that $\zeta_\ell \in K$. Hecke's theorem (see Section 10.2.3) gives us complete information on the ramification and relative discriminant for cyclic extensions of K of degree ℓ. In this section, we will show how Hecke's theorem allows us to find explicitly the Abelian extension L/K corresponding to a given congruence subgroup $(\mathfrak{m}, C)$ by Takagi's existence theorem and gives us a complete algorithm for this as long as the degree of L/K is equal to ℓ (see, in particular, Algorithm 5.2.14).

5.2.1 Characterization of Cyclic Extensions of Conductor $\mathfrak{m}$ and Degree ℓ

Let $\mathfrak{m}$ be a modulus, and let C be a congruence subgroup modulo $\mathfrak{m}$ such that

$$h_{\mathfrak{m},C} = \left|Cl_{\mathfrak{m}}(K)/\overline{C}\right| = |I_{\mathfrak{m}}/C| = \ell \ ,$$

so that the Abelian extension L/K corresponding to $(\mathfrak{m}, C)$ by class field theory is cyclic of degree ℓ.

Definition 5.2.1. *For a prime ideal $\mathfrak{p}$ dividing ℓ, denote by $z(\mathfrak{p}, \ell)$ the quantity*

$$z(\mathfrak{p}, \ell) = \ell \frac{e(\mathfrak{p}/\ell)}{\ell - 1} + 1$$

(see Theorem 10.2.9). We divide the prime ideals $\mathfrak{p}$ of K into six sets, as follows.

(1) *The set* $S_{\mathfrak{m},\ell,1}$ *(resp.,* $S_{\mathfrak{m},\ell,2}$*; resp.,* $S_{\mathfrak{m},\ell,3}$*) is the set of all prime ideals* $\mathfrak{p}$ *of* K *dividing both* $\mathfrak{m}$ *and* ℓ *and such that* $v_{\mathfrak{p}}(\mathfrak{m}) = z(\mathfrak{p},\ell)$ *(resp.,* $v_{\mathfrak{p}}(\mathfrak{m}) < z(\mathfrak{p},\ell)$*; resp.,* $v_{\mathfrak{p}}(\mathfrak{m}) > z(\mathfrak{p},\ell)$*).*
(2) *The set* S_ℓ *(resp.,* $S_{\mathfrak{m}}$*) is the set of all prime ideals* $\mathfrak{p}$ *of* K *dividing* ℓ *and not* $\mathfrak{m}$ *(resp.,* $\mathfrak{m}$ *and not* ℓ*).*
(3) *The set* $S_\varnothing$ *is the set of all prime ideals* $\mathfrak{p}$ *of* K *not dividing* $\mathfrak{m}$ *or* ℓ*.*

The main result, which is an easy consequence of Hecke's theorem, is as follows.

Theorem 5.2.2. *With the above notation, the field* $L = K(\sqrt[\ell]{\alpha})$ *with* $\alpha \in K^* \setminus K^{*\ell}$ *is a cyclic extension of* K *of conductor equal to* $\mathfrak{m}$ *and degree* ℓ *if and only if the following ten conditions hold.*

(1) $S_{\mathfrak{m},\ell,3} = \varnothing$.
(2) *If* $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$*, then* $v_{\mathfrak{p}}(\mathfrak{m}) \not\equiv 1 \pmod{\ell}$ *and, in particular,* $v_{\mathfrak{p}}(\mathfrak{m}) \geq 2$.
(3) *If* $\mathfrak{p} \in S_{\mathfrak{m}}$*, then* $v_{\mathfrak{p}}(\mathfrak{m}) = 1$.
(4) *If* $\mathfrak{p} \in S_{\mathfrak{m},\ell,1}$*, then* $\ell \nmid v_{\mathfrak{p}}(\alpha)$.
(5) *If* $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$*, then* $\ell \mid v_{\mathfrak{p}}(\alpha)$ *and the largest* k *such that the congruence*

$$\alpha \equiv x^\ell \pmod{\mathfrak{p}^{v_{\mathfrak{p}}(\alpha)+k}}$$

has a solution must be equal to $z(\mathfrak{p},\ell) - v_{\mathfrak{p}}(\mathfrak{m})$.
(6) *If* $\mathfrak{p} \in S_\ell$*, then* $\ell \mid v_{\mathfrak{p}}(\alpha)$ *and the congruence*

$$\alpha \equiv x^\ell \pmod{\mathfrak{p}^{v_{\mathfrak{p}}(\alpha)+z(\mathfrak{p},\ell)-1}}$$

has a solution.
(7) *If* $\mathfrak{p} \in S_{\mathfrak{m}}$*, then* $\ell \nmid v_{\mathfrak{p}}(\alpha)$.
(8) *If* $\mathfrak{p} \in S_\varnothing$*, then* $\ell \mid v_{\mathfrak{p}}(\alpha)$.
(9) *If* $\sigma \in \mathfrak{m}_\infty$*, then* $\sigma(\alpha) < 0$.
(10) *If* σ *is a real embedding that is not in* $\mathfrak{m}_\infty$*, then* $\sigma(\alpha) > 0$.

Remarks

(1) The first three conditions are only on the modulus $\mathfrak{m}$, while the others are on α.
(2) The last two conditions are used only if $\ell = 2$, since otherwise the condition $\zeta_\ell \in K$ implies that K is totally complex.

Proof. Assume first that L/K is of conductor equal to $\mathfrak{m}$. Then by Corollary 3.5.12 (1), we know that $\mathfrak{d}(L/K) = \mathfrak{m}_0^{\ell-1}$, where as usual $\mathfrak{m}_0$ is the finite part of $\mathfrak{m}$. By Theorem 10.2.9, we thus have the following.

(1) If $\ell \nmid v_{\mathfrak{p}}(\alpha)$, then $v_{\mathfrak{p}}(\mathfrak{m}) = z(\mathfrak{p},\ell)$.
(2) If $\mathfrak{p} \nmid \ell$ and $\ell \mid v_{\mathfrak{p}}(\alpha)$, then $v_{\mathfrak{p}}(\mathfrak{m}) = 0$.

(3) If $\mathfrak{p} \mid \ell$, $\mathfrak{p} \mid \mathfrak{m}$, and $\ell \mid v_\mathfrak{p}(\alpha)$, then $v_\mathfrak{p}(\mathfrak{m}) = 0$ if $a \geq z(\mathfrak{p},\ell) - 1$, $v_\mathfrak{p}(\mathfrak{m}) = z(\mathfrak{p},\ell) - a$ if $a < z(\mathfrak{p},\ell)$, where a is the largest value of k for which the congruence

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$$

has a solution.

By Theorem 10.2.9, we also know that $a \geq 1$ and $\ell \nmid a$. Since we want all the places of K dividing $\mathfrak{m}$, and only those, to ramify, this implies immediately all the necessary conditions on α. It also implies that $S_{\mathfrak{m},\ell,3} = \varnothing$. Finally, the two other conditions (2) and (3) on the modulus $\mathfrak{m}$ are immediate consequences of Corollary 3.5.12 (2).

Conversely, let $\mathfrak{m}$ and α be such that the conditions of the theorem are satisfied. Let $\mathfrak{d}'$ be the relative discriminant ideal of L/K, where $L = K(\sqrt[\ell]{\alpha})$. Theorem 10.2.9 allows us to compute $\mathfrak{d}'$ as $\mathfrak{d}' = P_1P_2P_3$, where

$$P_1 = \prod_{\mathfrak{p}\in S_{\mathfrak{m},\ell,1}} \mathfrak{p}^{(\ell-1)z(\mathfrak{p},\ell)} = \prod_{\mathfrak{p}\in S_{\mathfrak{m},\ell,1}} \mathfrak{p}^{(\ell-1)v_\mathfrak{p}(\mathfrak{m})} ,$$

$$P_2 = \prod_{\mathfrak{p}\in S_{\mathfrak{m},\ell,2},\ v_\mathfrak{p}(\mathfrak{m})\geq 2} \mathfrak{p}^{(\ell-1)(z(\mathfrak{p},\ell)-(z(\mathfrak{p},\ell)-v_\mathfrak{p}(\mathfrak{m})))} = \prod_{\mathfrak{p}\in S_{\mathfrak{m},\ell,2},\ v_\mathfrak{p}(\mathfrak{m})\geq 2} \mathfrak{p}^{(\ell-1)v_\mathfrak{p}(\mathfrak{m})} ,$$

$$P_3 = \prod_{\mathfrak{p}\in S_\mathfrak{m}} \mathfrak{p}^{\ell-1} .$$

The restriction $v_\mathfrak{p}(\mathfrak{m}) \geq 2$ in the product P_2 comes from the fact that, if $v_\mathfrak{p}(\mathfrak{m}) = 1$, then by Theorem 10.2.9, $\mathfrak{p}$ is unramified.

After simplifications, we obtain $\mathfrak{d}' = \mathfrak{m}_0^{\ell-1}/P_4P_5$ with

$$P_4 = \prod_{\mathfrak{p}\in S_\mathfrak{m}} \mathfrak{p}^{(\ell-1)(v_\mathfrak{p}(\mathfrak{m})-1)} \quad \text{and} \quad P_5 = \prod_{\mathfrak{p}\in S_{\mathfrak{m},\ell,2},\ v_\mathfrak{p}(\mathfrak{m})=1} \mathfrak{p}^{\ell-1} .$$

Conditions (2) and (3) on the modulus imply that $P_5 = \mathbb{Z}_K$, $P_4 = \mathbb{Z}_K$, respectively; hence $\mathfrak{d}' = \mathfrak{m}_0^{\ell-1}$. Since, by Corollary 3.5.12 (1), we also have $\mathfrak{d}' = \mathfrak{f}_0^{\ell-1}$, where $\mathfrak{f}$ is the conductor, we deduce that $\mathfrak{m}_0 = \mathfrak{f}_0$. Finally, the last conditions on the signatures imply that the ramified real places are exactly those in $\mathfrak{m}_\infty$, so we have $\mathfrak{m}_\infty = \mathfrak{f}_\infty$, hence $\mathfrak{m} = \mathfrak{f}$ as desired. □

5.2.2 Virtual Units and the ℓ-Selmer Group

To use this theorem in practice, we must introduce some notation and definitions. Let

$$Cl(K) = \bigoplus_{1\leq i\leq g_c} (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$$

be the SNF of the class group of K, where the $\mathfrak{a}_i$ are ideals of K. If r_c is the largest index such that $\ell \mid d_i$, then we clearly have

$$Cl(K)/Cl(K)^\ell = \bigoplus_{1\le i\le r_c} (\mathbb{Z}/\ell\mathbb{Z})\overline{\mathfrak{a}_i} \ ,$$

hence r_c is the ℓ-rank of the group $Cl(K)$. It follows that if I is an ideal of K, we can write $\overline{I} = \prod_{1\le i\le r_c} \overline{\mathfrak{a}_i}^{x_i}$ for $0 \le x_i < \ell$, where $\overline{\ }$ denotes the class in $Cl(K)/Cl(K)^\ell$. Lifting to $Cl(K)$, then to the ideals of K themselves, it follows that any ideal can be written in the form

$$I = \beta\mathfrak{q}^\ell \prod_{1\le i\le r_c} \mathfrak{a}_i^{x_i} \quad \text{with} \quad 0 \le x_i < \ell \ ,$$

and the x_i are unique.

Note that, thanks to Corollary 1.3.9, we could assume that the representatives $\mathfrak{a}_i$ of the ideal classes $\overline{\mathfrak{a}_i}$ are chosen in such a way as to be coprime with anything in sight, here coprime to ℓ and $\mathfrak{m}$. However, this would be inefficient for algorithmic purposes, especially since we will see that we do not need this hypothesis, so we do *not* assume that the $\mathfrak{a}_i$ are necessarily coprime with ℓ and $\mathfrak{m}$.

We define elements α_i, $\beta_\mathfrak{p}$ of K^* and integers $p_{i,\mathfrak{p}}$ for $1 \le i \le r_c$ and $\mathfrak{p} \in S = S_\mathfrak{m} \cup S_{\mathfrak{m},\ell,1}$ by the following formulas:

$$\mathfrak{a}_i^{d_i} = \alpha_i\mathbb{Z}_K \quad \text{for} \quad 1 \le i \le r_c \ ,$$
$$\mathfrak{p} = \beta_\mathfrak{p}\mathfrak{q}_\mathfrak{p}^\ell \prod_{1\le i\le r_c} \mathfrak{a}_i^{p_{i,\mathfrak{p}}} \quad \text{for } \mathfrak{p} \in S \ .$$

We will see in the next section how to compute such elements, but for now simply note their existence. We may, of course, assume if desired that $0 \le p_{i,\mathfrak{p}} < \ell$ for all i.

Proposition 5.2.3. *Let $\gamma \in K^*$. The following two properties are equivalent.*

(1) *There exists an ideal $\mathfrak{q}$ such that $\gamma\mathbb{Z}_K = \mathfrak{q}^\ell$.*
(2) *The element γ belongs to the group generated by the units, the α_i defined above for $1 \le i \le r_c$, and the ℓth powers of elements of K^*.*

Proof. Since for $i \le h$, we have $\alpha_i\mathbb{Z}_K = (\mathfrak{a}_i^{d_i/\ell})^\ell$, it is clear that if γ belongs to the group mentioned in the proposition, then $\gamma\mathbb{Z}_K$ is the ℓth power of an ideal. Conversely, assume that $\gamma\mathbb{Z}_K = \mathfrak{q}^\ell$. Then, if $\mathfrak{q} = \beta\prod_{1\le i\le g_c} \mathfrak{a}_i^{x_i}$, we have $\gamma\mathbb{Z}_K = \beta^\ell \prod_{1\le i\le g_c} \mathfrak{a}_i^{\ell x_i}$, hence $d_i \mid \ell x_i$ for all i, so $d_i \mid x_i$ for $i > r_c$, while $(d_i/\ell) \mid x_i$ for $i \le r_c$. It follows that

$$\gamma\mathbb{Z}_K = \beta^\ell \prod_{1\le i\le r_c} \alpha_i^{n_i} \prod_{r_c<i\le g_c} \alpha_i^{n_i\ell}$$

with $n_i = x_i/d_i$ for $i > r_c$ and $n_i = x_i/(d_i/\ell)$ for $i \le r_c$, thus proving the proposition. □

Note that we have set $\alpha_i \mathbb{Z}_K = \mathfrak{a}_i^{d_i}$ also for $i > r_c$, but these α_i do not occur in any of the definitions, only in proofs or in algorithms. In particular, they are *not* virtual units in the sense of the following definition. □

Definition 5.2.4. (1) *An element $\gamma \in K^*$ satisfying one of the two equivalent conditions of the above proposition will be called an ℓ-*virtual unit*, or more simply a* virtual unit *if there is no risk of confusion.*

(2) *The set of virtual units forms a multiplicative group, which we will denote by $V_\ell(K)$.*

(3) *The quotient group $V_\ell(K)/K^{*\ell}$ will be called the ℓ-*Selmer group *of the number field K and denoted $S_\ell(K)$.*

Proposition 5.2.5. *Let $r_u = r_1 + r_2 - 1$ be the rank of the torsion-free part of $U(K)$. Recall that we denote by r_c the ℓ-rank of $Cl(K)$. We denote by $(\varepsilon_j)_{1 \le j \le r_u}$ a system of fundamental units, and by ε_0 a generator of the group of roots of unity in K of order $w(K)$.*

(1) *The quotient group $U(K)/U(K)^\ell$ is a $\mathbb{Z}/\ell\mathbb{Z}$-vector space of dimension $r_u + 1$, a basis consisting of the classes of the ε_j for $0 \le j \le r_u$.*

(2) *The quotient group $V_\ell(K)/K^{*\ell}$ is a $\mathbb{Z}/\ell\mathbb{Z}$-vector space of dimension $r_v = r_c + r_u + 1$, a basis consisting of the classes of the ε_j for $0 \le j \le r_u$ and of the α_i for $1 \le i \le r_c$.*

Proof. Since $V_\ell(K)$ is generated by the α_i, the ε_j, and $K^{*\ell}$, we must simply find the dependencies between the α_i and ε_j in $V_\ell(K)/K^{*\ell}$ as a $\mathbb{Z}/\ell\mathbb{Z}$-vector space. Hence assume that

$$\prod_{0 \le j \le r_u} \varepsilon_j^{x_j} \prod_{1 \le i \le r_c} \alpha_i^{n_i} = \gamma^\ell$$

for some $\gamma \in K^*$. By definition of α_i, this implies

$$\prod_{1 \le i \le r_c} \mathfrak{a}_i^{d_i n_i} = \gamma^\ell \mathbb{Z}_K \ ;$$

hence $\mathfrak{b}^\ell = \gamma^\ell \mathbb{Z}_K = (\gamma \mathbb{Z}_K)^\ell$ with

$$\mathfrak{b} = \prod_{1 \le i \le r_c} \mathfrak{a}_i^{(d_i/\ell) n_i} \ .$$

Thus, $\mathfrak{b} = \gamma \mathbb{Z}_K$ is a principal ideal; hence $d_i \mid (d_i/\ell) n_i$ for $1 \le i \le r_c$ or, equivalently, $n_i \equiv 0 \pmod{\ell}$, so the virtual units α_i do not enter into our dependency. Thus, it is enough to prove (1).

So assume that

$$\prod_{0\le j\le r_u} \varepsilon_j^{x_j} = \varepsilon^\ell$$

for some $\varepsilon \in K^*$. Since ε is a root of the monic polynomial $X^\ell - \varepsilon^\ell = 0$ with coefficients in $\mathbb{Z}_K$, it follows that it is an algebraic integer (see, for example, [Coh0, Corollary 4.1.5]), and since $\varepsilon \in K^*$, we have $\varepsilon \in \mathbb{Z}_K$. Furthermore, the absolute norm of ε is in $\mathbb{Z}$, and its ℓth power is equal to ± 1, from which it follows that $\mathcal{N}(\varepsilon) = \pm 1$, hence ε is a unit of K. For future reference, we isolate this as a lemma.

Lemma 5.2.6. *We have* $U(K) \cap K^{*\ell} = U(K)^\ell$.

Thus, if $\varepsilon = \prod_{0\le j\le r_u} \varepsilon_j^{y_j}$, we have $x_j = \ell y_j$ for $j \ge 1$, while $x_0 \equiv \ell y_0 \pmod{w(K)}$. Thus, for $j \ge 1$, $x_j \equiv 0 \pmod{\ell}$ so the fundamental units ε_j do not enter into our dependency. Furthermore, since $\zeta_\ell \in K$, we have $\ell \mid w(K)$, so we also have $x_0 \equiv 0 \pmod{\ell}$, proving the proposition. □

Remark. In the case where $\zeta_\ell \notin K$ which is not considered here, we have $\ell \nmid w(K)$. Hence a primitive $w(K)$th root of unity is an ℓ-power, so the rank of $U(K)/U(K)^\ell$ is only equal to r_u, and that of $V_\ell(K)/K^{*\ell}$ is equal to $r_c + r_u$.

Definition 5.2.7. *We will write* $v_i = \alpha_i$ *for* $1 \le i \le r_c$, *and* $v_{i+r_c} = \varepsilon_{i-1}$ *for* $1 \le i \le r_u + 1$. *Thus, the* $\overline{v_i}$ *form a* $\mathbb{Z}/\ell\mathbb{Z}$*-basis for the* ℓ*-Selmer group* $V_\ell(K)/K^{*\ell}$.

Proposition 5.2.8. *We have the following exact sequence:*

$$1 \longrightarrow \mu_\ell(K) \longrightarrow U(K) \xrightarrow{[\ell]} U(K) \longrightarrow \frac{V_\ell(K)}{K^{*\ell}} \xrightarrow{\phi} Cl(K) \xrightarrow{[\ell]} Cl(K) \longrightarrow \frac{Cl(K)}{Cl(K)^\ell} \longrightarrow 1 \ .$$

Here, $\mu_\ell(K)$ *is the group of* ℓ*th roots of unity in* K, $[\ell]$ *denotes the map that raises to the* ℓ*th power (in* $Cl(K)$ *or* $U(K)$*), and* ϕ *is the map that sends the class of a virtual unit* v *to the ideal class of the ideal* $\mathfrak{q}$ *such that* $v\mathbb{Z}_K = \mathfrak{q}^\ell$.

In particular, we have the following short exact sequence:

$$1 \longrightarrow \frac{U(K)}{U(K)^\ell} \longrightarrow \frac{V_\ell(K)}{K^{*\ell}} \xrightarrow{\phi} Cl(K)[\ell] \longrightarrow 1 \ ,$$

where $Cl(K)[\ell]$ *denotes the subgroup of* $Cl(K)$ *of ideal classes killed by* ℓ.

Proof. The proof is straightforward and is left to the reader (Exercise 3). Note also that this proposition shows once again that the ℓ-rank of $V_\ell(K)/K^{*\ell}$ is equal to the sum of the ℓ-ranks of $U(K)/U(K)^\ell$ and of $Cl(K)/Cl(K)^\ell$. □

Remark. We have a very similar situation in the case of elliptic curves over some fixed number field, say $\mathbb{Q}$. The group of units $U(K)$ is analogous to the Mordell–Weil group $E(\mathbb{Q})$, the rank $r_u + 1$ of $U(K)/U(K)^\ell$ is analogous to the rank of $E(\mathbb{Q})/\ell E(\mathbb{Q})$, and the rank r_c of the ℓ-part of the class group is analogous to the rank of the ℓ-part of the Tate–Shafarevitch group of the curve. I refer to [Sil2] for these notions. Thus, it is perfectly reasonable to call $V_\ell(K)/K^{*\ell}$ the ℓ-Selmer group of the number field K as we have done above. The above exact sequences are analogs of the corresponding exact sequences for ℓ-Selmer groups of elliptic curves.

5.2.3 Construction of Cyclic Extensions of Prime Degree and Conductor $\mathfrak{m}$

The following theorem is the basis of our explicit Kummer algorithms when $\zeta_\ell \in K$.

Theorem 5.2.9. *Keep the above notation, and in particular recall that we write* $z(\mathfrak{p},\ell) = \ell e(\mathfrak{p}/\ell)/(\ell-1)+1$ *and* $S = S_\mathfrak{m} \cup S_{\mathfrak{m},\ell,1}$.

Let L/K be a cyclic extension of K of degree ℓ and of conductor equal to $\mathfrak{m}$. Then $\mathfrak{m}$ satisfies conditions (1), (2), *and* (3) *of Theorem 5.2.2, and up to Kummer-equivalence, we can choose $L = K(\sqrt[\ell]{\alpha})$ with α of the following form:*

$$\alpha = \prod_{\mathfrak{p}\in S} \beta_\mathfrak{p}^{x_\mathfrak{p}} \prod_{i=1}^{r_c+r_u+1} v_i^{n_i} \ ,$$

with the following additional conditions.

(1) *For all $\mathfrak{p} \in S$, we have $1 \le x_\mathfrak{p} \le \ell-1$; for all i, we have $0 \le n_i \le \ell-1$.*
(2) *For $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$, the largest k such that the congruence*

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$$

has a solution must be equal to $z(\mathfrak{p},\ell) - v_\mathfrak{p}(\mathfrak{m})$.
(3) *If S is not empty, we may fix any one (but only one) of the $x_\mathfrak{p}$ equal to* 1.
(4) *For each $\mathfrak{p} \in S_\ell$, the congruence*

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{z(\mathfrak{p},\ell)-1+v_\mathfrak{p}(\alpha)}}$$

has a solution.
(5) *For each $i \le r_c$, we must have*

$$\sum_{\mathfrak{p}\in S} x_\mathfrak{p} p_{i,\mathfrak{p}} \equiv 0 \pmod{\ell} \ .$$

(6) *For each $\sigma \in \mathfrak{m}_\infty$, $\sigma(\alpha) < 0$, while for each real embedding $\sigma \notin \mathfrak{m}_\infty$, we have $\sigma(\alpha) > 0$.*

Conversely, if $\mathfrak{m}$ *satisfies conditions* (1), (2), *and* (3) *of Theorem 5.2.2, if the above conditions are satisfied, and if* $\alpha \neq 1$, *then* $L = K(\sqrt[\ell]{\alpha})$ *is a cyclic extension of degree* ℓ *and of conductor* $\mathfrak{m}$.

Proof. Since by Theorem 5.2.2, we have $S_{\mathfrak{m},\ell,3} = \varnothing$, we can write

$$\alpha\mathbb{Z}_K = \prod_{\mathfrak{p}\in S} \mathfrak{p}^{x_\mathfrak{p}} \prod_{\mathfrak{p}\in S_{\mathfrak{m},\ell,2}} \mathfrak{p}^{x_\mathfrak{p}} \prod_{\mathfrak{p}\in S_\ell} \mathfrak{p}^{x_\mathfrak{p}} \prod_{\mathfrak{p}\in S_\varnothing} \mathfrak{p}^{x_\mathfrak{p}} .$$

By Theorem 5.2.2, when $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$, $\mathfrak{p} \in S_\ell$, or $\mathfrak{p} \in S_\varnothing$, we must have $\ell \mid v_\mathfrak{p}(\alpha) = x_\mathfrak{p}$. By the approximation theorem, we can find $\gamma \in K^*$ such that $v_\mathfrak{p}(\gamma) = -x_\mathfrak{p}/\ell$ for $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$ and $\mathfrak{p} \in S_\ell$, $v_\mathfrak{p}(\gamma) = -\lfloor x_\mathfrak{p}/\ell \rfloor$ for $\mathfrak{p} \in S = S_\mathfrak{m} \cup S_{\mathfrak{m},\ell,1}$, and no special conditions for $\mathfrak{p} \in S_\varnothing$. Since α is Kummer-equivalent to $\alpha\gamma^\ell$, we may thus replace α by $\alpha\gamma^\ell$, hence for this new α we will have $x_\mathfrak{p} = 0$ for $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$ and $\mathfrak{p} \in S_\ell$, and also $1 \leq x_\mathfrak{p} \leq \ell - 1$ for $\mathfrak{p} \in S$ (since by Theorem 5.2.2 we must have $\ell \nmid v_\mathfrak{p}(\alpha)$ for $\mathfrak{p} \in S$). To summarize, up to Kummer-equivalence, we have

$$\alpha\mathbb{Z}_K = \mathfrak{q}^\ell \prod_{\mathfrak{p}\in S} \mathfrak{p}^{x_\mathfrak{p}} ,$$

where $\mathfrak{q}$ is an ideal coprime to $\mathfrak{m}$ and ℓ.

Replacing the prime ideals $\mathfrak{p} \in S$ by their expressions in $Cl(K)/Cl(K)^\ell$, we obtain

$$\alpha\mathbb{Z}_K = \mathfrak{q}_1^\ell \prod_{\mathfrak{p}\in S} \beta_\mathfrak{p}^{x_\mathfrak{p}} \prod_{1\leq i\leq r_c} \mathfrak{a}_i^{y_i} ,$$

with

$$y_i = \sum_{\mathfrak{p}\in S} x_\mathfrak{p} p_{i,\mathfrak{p}}$$

and some other ideal $\mathfrak{q}_1$ (this time not necessarily coprime to $\mathfrak{m}$ and ℓ).

In the quotient group $Cl(K)/Cl(K)^\ell$, we thus have

$$\prod_{1\leq i\leq r_c} \overline{\mathfrak{a}_i}^{y_i} = 1 ,$$

and since the $\overline{\mathfrak{a}_i}$ form a $\mathbb{Z}/\ell\mathbb{Z}$-basis, we have $\ell \mid y_i$ for each i such that $1 \leq i \leq r_c$. Thus, we have shown that

$$\alpha\mathbb{Z}_K = \mathfrak{q}_2^\ell \prod_{\mathfrak{p}\in S} \beta_\mathfrak{p}^{x_\mathfrak{p}}$$

for some ideal $\mathfrak{q}_2$.

Since $\mathfrak{q}_2^\ell$ is both a principal ideal and the ℓth power of an ideal, by Proposition 5.2.3 it is of the form $v\mathbb{Z}_K$ for $v \in V_\ell(K)$, showing that α is of the form given in the theorem.

We have seen that $1 \le x_\mathfrak{p} \le \ell - 1$ and that for $1 \le i \le r_c$, we must have

$$y_i = \sum_{\mathfrak{p} \in S} x_\mathfrak{p} p_{i,\mathfrak{p}} \equiv 0 \pmod{\ell} \ .$$

Up to Kummer-equivalence, we may of course also choose $0 \le n_i \le \ell - 1$.

Finally, in the condition of Kummer-equivalence, we are allowed another degree of freedom in addition to multiplying by an ℓth power: we may also raise α to some power coprime to ℓ. If S is nonempty, this can be used to fix one (and only one) of the $x_\mathfrak{p}$ equal to 1, since we know that they are not divisible by ℓ.

The other conditions of the corollary follow immediately from Theorem 5.2.2, and hence we have proved that all the conditions of the corollary are necessary.

Let us prove the converse. Assume that all the conditions are satisfied. Since by definition,

$$\mathfrak{p} = \beta_\mathfrak{p} \mathfrak{q}_\mathfrak{p}^\ell \prod_{1 \le i \le r_c} \mathfrak{a}_i^{p_{i,\mathfrak{p}}} \ ,$$

it follows that for any prime ideal $\mathfrak{q}$, we have

$$v_\mathfrak{q}(\beta_\mathfrak{p}) \equiv \delta_{\mathfrak{p},\mathfrak{q}} - \sum_{1 \le i \le r_c} p_{i,\mathfrak{p}} v_\mathfrak{p}(\mathfrak{a}_i) \pmod{\ell} \ ,$$

where $\delta_{\mathfrak{p},\mathfrak{q}}$ is the Kronecker symbol. In addition, for $1 \le i \le r_c$, we have $v_\mathfrak{q}(\alpha_i) = d_i v_\mathfrak{q}(\mathfrak{a}_i) \equiv 0 \pmod{\ell}$. Thus, for any prime ideal $\mathfrak{q}$ we have

$$v_\mathfrak{q}(\alpha) \equiv \sum_{\mathfrak{p} \in S} x_\mathfrak{p} v_\mathfrak{q}(\beta_\mathfrak{p}) \equiv \sum_{\mathfrak{p} \in S} x_\mathfrak{p} \delta_{\mathfrak{p},\mathfrak{q}} - \sum_{1 \le i \le r_c} v_\mathfrak{p}(\mathfrak{a}_i) \sum_{\mathfrak{p} \in S} x_\mathfrak{p} p_{i,\mathfrak{p}} \pmod{\ell} \ ;$$

hence by condition (5), we have $v_\mathfrak{q}(\alpha) \equiv 0 \pmod{\ell}$ if $\mathfrak{q} \notin S$, while $v_\mathfrak{q}(\alpha) \equiv x_\mathfrak{q} \not\equiv 0 \pmod{\ell}$ if $\mathfrak{q} \in S$ by condition (1). This and the other conditions imply that all the conditions of Theorem 5.2.2 are satisfied. To finish the proof, we must show that $\alpha \notin K^{*\ell}$. Indeed, if α is an ℓth power, then $S = \varnothing$ (otherwise $x_\mathfrak{p} \equiv 0 \pmod{\ell}$); hence α is a virtual unit that is equal to an ℓth power of an element. It follows from Proposition 5.2.5 that $n_i = 0$ for all i, hence that $\alpha = 1$, contrary to the assumption of the theorem. □

The conditions of the theorem already restrict α to a finite set of cardinality at most equal to $(\ell - 1)^{|S|} \ell^{r_v}$. If α belongs to this finite set, we will know that $L_\alpha = K(\sqrt[\ell]{\alpha})$ is a cyclic extension of degree ℓ, conductor $\mathfrak{m}$, hence relative discriminant $\mathfrak{d}(L/K)$, and correct signature. Of course, there may be several fields L_α satisfying all these conditions. To terminate, we compute the norm group for each of the possible α in our finite set, and exactly one will be equal to $(\mathfrak{m}, C)$. The α we find are in one-to-one correspondence with congruence subgroups $(\mathfrak{m}, C)$ such that $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$.

5.2.4 Algorithmic Kummer Theory When $\zeta_\ell \in K$ Using Hecke

In this section, we simply put in formal algorithmic form the results of the preceding sections. We assume as above that K is a number field and that ℓ is a prime number such that $\zeta_\ell \in K$.

We keep the notation of the preceding section. In particular, if

$$Cl(K) = \bigoplus_{1 \le i \le g_c} (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$$

is an SNF for the class group $Cl(K)$, we set $\mathfrak{a}_i^{d_i} = \alpha_i \mathbb{Z}_K$, and r_c denotes the largest index such that $\ell \mid d_i$, in other words the ℓ-rank of $Cl(K)$.

The elements α_i that are needed to use Theorem 5.2.9 are found by directly using the principal ideal algorithm ([Coh0, Algorithm 6.5.10]). For the integers $p_{i,\mathfrak{p}}$ and the elements $\beta_\mathfrak{p}$, we use the following general algorithm.

Algorithm 5.2.10 (Decomposition of an Ideal in $Cl(K)/Cl(K)^\ell$). Keep the above notation and let $\mathfrak{b}$ be an ideal of K. This algorithm computes $\beta \in K^*$ and integers b_i such that there exists an ideal $\mathfrak{q}$ (which is not computed) such that

$$\mathfrak{b} = \beta \mathfrak{q}^\ell \prod_{1 \le i \le r_c} \mathfrak{a}_i^{b_i} \ .$$

1. [Use principal ideal algorithm] Using [Coh0, Algorithm 6.5.10], compute α and integers b_i such that $\mathfrak{b} = \alpha \prod_{1 \le i \le g_c} \mathfrak{a}_i^{b_i}$.
2. [Compute α_i for $i > r_c$] Using the same algorithm, for each i such that $r_c < i \le g_c$, compute $\alpha_i \in K$ such that $\mathfrak{a}_i^{d_i} = \alpha_i \mathbb{Z}_K$.
3. [Compute β] For each i such that $r_c < i \le g_c$ (equivalently, such that $\ell \nmid d_i$), using Euclid's extended algorithm, compute an integer u_i such that $u_i d_i \equiv 1 \pmod{\ell}$, set
$$\beta \leftarrow \alpha \prod_{r_c < i \le g_c} \alpha_i^{b_i u_i \bmod \ell} \ ,$$
output β and the b_i for $i \le r_c$, and terminate the algorithm.

Proof. Note that for $i > r_c$ we have $(d_i, \ell) = 1$; hence by Euclid's extended algorithm, we can find u_i and v_i such that $u_i d_i + v_i \ell = 1$. It follows that for $i > r_c$, we have

$$\mathfrak{a}_i^{b_i} = \alpha_i^{u_i b_i} \mathfrak{q}_i^\ell \quad \text{with} \quad \mathfrak{q}_i = \mathfrak{a}_i^{v_i b_i} \ ,$$

which shows the algorithm's validity. □

Remark. Since we will in practice apply this algorithm for many ideals $\mathfrak{b}$, we compute the α_i and u_i once and for all. Please recall once again that the α_i for $i > r_c$ are *not* ℓ-virtual units.

When we apply this algorithm to $\mathfrak{b} = \mathfrak{p}$, we obtain an element $\beta_\mathfrak{p} = \beta$, exponents $p_{i,\mathfrak{p}} = b_i$, and an ideal $\mathfrak{q}_\mathfrak{p} = \mathfrak{q}$. Since the ideals $\mathfrak{q}_\mathfrak{p}$ play no practical

role (although we must keep them for the theoretical analysis), while the $\beta_\mathfrak{p}$ are essential, it is useful to choose $\beta_\mathfrak{p}$ as simple as possible. Indeed, we have some freedom in choosing $\beta_\mathfrak{p}$. By definition, we may replace $\beta_\mathfrak{p}$ by any $\beta'_\mathfrak{p}$ such that $\beta'_\mathfrak{p}/\beta_\mathfrak{p}$ is the ℓth power of *an ideal* and not only of an element (since this simply changes the ideal $\mathfrak{q}_\mathfrak{p}$); in other words, we may multiply $\beta_\mathfrak{p}$ by any virtual unit if desired. This can be done in several ways, which are rather technical, so we will not give any details here, but simply note that in practice this reduction should be attempted.

To be able to give an algorithm corresponding to Theorem 5.2.9, we must explain how to check whether or not a congruence of the form

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$$

has a solution, where it is known that $\ell \mid v_\mathfrak{p}(\alpha)$. For this, it is useful to generalize the notion of discrete logarithm in $(\mathbb{Z}_K/\mathfrak{m})^*$ to elements not coprime to $\mathfrak{m}$.

Definition 5.2.11. *Let $\mathfrak{m} = \mathfrak{m}_0\mathfrak{m}_\infty$ be a modulus and let $\mathfrak{m}_0 = \prod_\mathfrak{p} \mathfrak{p}^{k_\mathfrak{p}}$ be the factorization of $\mathfrak{m}_0$ into prime ideals. For each $\mathfrak{p} \mid \mathfrak{m}_0$, let $\pi_\mathfrak{p}$ be an element of $\mathfrak{p} \smallsetminus \mathfrak{p}^2$ not belonging to $\mathfrak{q}$ for any prime ideal $\mathfrak{q} \mid \mathfrak{m}_0$ different from $\mathfrak{p}$, and let*

$$(\mathbb{Z}_K/\mathfrak{m})^* = \bigoplus_{1\le i\le s} (\mathbb{Z}/c_i\mathbb{Z})\overline{g_i}$$

be the SNF of $(\mathbb{Z}_K/\mathfrak{m})^*$. If $\alpha \in K^*$, we say that $((v_\mathfrak{p}),(a_1,\dots,a_s))$ is a discrete logarithm *for α with respect to the generators g_i and the uniformizers $\pi_\mathfrak{p}$ if*

$$\alpha = \beta \prod_\mathfrak{p} \pi_\mathfrak{p}^{v_\mathfrak{p}} \prod_{1\le i\le s} g_i^{a_i}$$

with $\beta \equiv 1 \pmod{{}^*\mathfrak{m}}$.

Remarks

(1) A discrete logarithm always exists. Indeed, we must take $v_\mathfrak{p} = v_\mathfrak{p}(\alpha)$ for all $\mathfrak{p} \mid \mathfrak{m}$. Then $\alpha/\prod_\mathfrak{p} \pi_\mathfrak{p}^{v_\mathfrak{p}}$ is coprime to $\mathfrak{m}$; hence its usual discrete logarithm in $(\mathbb{Z}_K/\mathfrak{m})^*$ is well-defined.

(2) It is also clear that if $((v'_\mathfrak{p}),(a'_1,\dots,a'_s))$ is another discrete logarithm, then $v'_\mathfrak{p} = v_\mathfrak{p} = v_\mathfrak{p}(\alpha)$; hence for all i we have $a'_i \equiv a_i \pmod{c_i}$.

(3) If $\mathfrak{m} = \mathfrak{p}^k$, we evidently have

$$\alpha \equiv \pi_\mathfrak{p}^{v_\mathfrak{p}} \prod_{1\le i\le s} g_i^{a_i} \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}} \ .$$

The following proposition easily answers our congruence problem.

Proposition 5.2.12. *Let $\mathfrak{p}$ be a prime ideal, let $k \geq 1$ be an integer, and let $(\mathbb{Z}_K/\mathfrak{p}^k)^* = \bigoplus_{1\leq i\leq s}(\mathbb{Z}/c_i\mathbb{Z})\overline{g_i}$ as above. Let t be the largest index i such that $\ell \mid c_i$. If $\alpha \in K^*$, the congruence*

$$x^\ell \equiv \alpha \pmod{{}^*\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$$

has a solution if and only if

$$(v_\mathfrak{p}, (a_1, \ldots, a_t)) \equiv 0 \pmod{\ell} \ ,$$

where $(v_\mathfrak{p}, (a_1, \ldots, a_s))$ is the discrete logarithm of α as defined above.

Proof. Write $\alpha = \beta\pi_\mathfrak{p}^{v_\mathfrak{p}} \prod_{1\leq i\leq s} g_i^{a_i}$ and $x = \gamma\pi_\mathfrak{p}^{w} \prod_{1\leq i\leq s} g_i^{x_i}$ as above. Since $k \geq 1$, $x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}}}$ implies that $v_\mathfrak{p} = w\ell$, hence that $v_\mathfrak{p} \equiv 0 \pmod{\ell}$. In addition, since $\beta \equiv \gamma \equiv 1 \pmod{{}^*\mathfrak{p}^k}$, we must have $a_i \equiv \ell x_i \pmod{c_i}$ for all i. For $i \leq t$, the existence of x_i is equivalent to $\ell \mid a_i$, while for $i > t$, x_i always exists since c_i is coprime to ℓ, proving the proposition. □

This leads us to introduce the following notation.

Definition 5.2.13. *Let $\mathfrak{m}$ be a modulus, let $(\mathbb{Z}_K/\mathfrak{m})^* = \bigoplus_{1\leq i\leq s}(\mathbb{Z}/c_i\mathbb{Z})\overline{g_i}$ be in SNF, let $\alpha \in K^*$, and let $((v_\mathfrak{p}), (a_1, \ldots, a_s))$ be the discrete logarithm of α with respect to the generators g_i and the uniformizers $\pi_\mathfrak{p}$, as defined above. If t is the largest index (possibly 0) such that $\ell \mid c_i$ (t is the ℓ-rank of $(\mathbb{Z}_K/\mathfrak{m})^*$), then we set $L_\mathfrak{m}(\alpha) = (a_1, \ldots, a_t)$ and call it the* short *discrete logarithm of α.*

Thus, if we know that $\ell \mid v_\mathfrak{p}(\alpha)$, the above proposition tells us that the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$ has a solution if and only if $L_{\mathfrak{p}^k}(\alpha) \equiv 0 \pmod{\ell}$.

Remarks

(1) To compute the short discrete logarithm, we use Algorithm 4.2.24, or Algorithm 4.2.18 if $\mathfrak{m}$ is the power of a prime ideal. An important special case, however, is case (4) of Theorem 5.2.9. We frequently have $e(\mathfrak{p}, \ell) = \ell - 1$, hence $z(\mathfrak{p}, \ell) - 1 = \ell$, and so we may apply Proposition 4.2.19 instead of the general algorithm.

(2) It is important to use the short discrete logarithm in order to find suitable elements satisfying our congruence conditions (see steps 6 to 8 of Algorithm 5.2.14 below). If, however, we only need to test whether a congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$ has a solution for a given α, it is probably faster to use the methods explained in Section 10.2.4.

We can now give the complete algorithm for computing explicitly the Abelian extension L/K using Hecke's theorem when $\zeta_\ell \in K$.

Algorithm 5.2.14 (Kummer Extensions of Prime Degree When $\zeta_\ell \in K$ Using Hecke). Let K be a number field and ℓ be a prime number such that $\zeta_\ell \in K$. We assume that the groups $Cl(K)$ and $U(K)$ have been explicitly computed, as well as the α_i for $1 \le i \le r_c$ (using the above notation). As above, let $(v_i)_{1\le i\le r_c+r_u+1}$ be a generating set for the group of virtual units $V_\ell(K)$ (modulo ℓth powers) generated by $U(K)$ and the α_i. Let $\mathfrak{m}$ be an arbitrary modulus of K. This algorithm outputs defining polynomials for all the Abelian extensions L/K of degree ℓ and of conductor equal to $\mathfrak{m}$.

1. [Factor $\mathfrak{m}$ and ℓ] Using Algorithm 2.3.22 (in the absolute case), find the prime ideal factorization of the finite part of the modulus $\mathfrak{m}_0 = \prod_{\mathfrak{p}|\mathfrak{m}_0} \mathfrak{p}^{v_\mathfrak{p}(\mathfrak{m}_0)}$, and using [Coh0, Algorithm 6.2.9], compute the prime ideal factorization of $\ell\mathbb{Z}_K$.
2. [Compute sets of prime ideals] Compute the finite sets $S_\mathfrak{m}$, S_ℓ, and $S_{\mathfrak{m},\ell,i}$ for $i = 1, 2, 3$ according to Definition 5.2.1.
3. [Test conditions on $\mathfrak{m}$] If $S_{\mathfrak{m},\ell,3} \neq \varnothing$, or if there exists $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$ such that $v_\mathfrak{p}(\mathfrak{m}) \equiv 1 \pmod{\ell}$, or if there exists $\mathfrak{p} \in S_\mathfrak{m}$ such that $v_\mathfrak{p}(\mathfrak{m}) \ge 2$, there are no suitable Abelian extensions L/K, so terminate the algorithm.
4. [Compute $\beta_\mathfrak{p}$ and $p_{i,\mathfrak{p}}$] Using Algorithm 5.2.10, for each $\mathfrak{p} \in S = S_\mathfrak{m} \cup S_{\mathfrak{m},\ell,1}$, compute $\beta_\mathfrak{p} \in K^*$ and integers $p_{i,\mathfrak{p}}$ such that for some ideal $\mathfrak{q}_\mathfrak{p}$ we have $\mathfrak{p} = \beta_\mathfrak{p}\mathfrak{q}_\mathfrak{p}^\ell \prod_{1\le i\le r_c} \mathfrak{a}_i^{p_{i,\mathfrak{p}}}$.
5. [Introduce notation] (This is a notational step, not really anything to be done.) To ease notation, set $r_v \leftarrow r_c + r_u + 1$, and for $1 \le j \le r_v$ let v_j be virtual units such that the $(\overline{v_j})_{1\le j\le r_v}$ form a $\mathbb{Z}/\ell\mathbb{Z}$-basis of $V_\ell(K)/K^{*\ell}$ as in Definition 5.2.7. For $1 \le j \le s$, let $\mathfrak{p}_j$ be the prime ideals in S, set $v_{j+r_v} \leftarrow \beta_{\mathfrak{p}_j}$ for $1 \le j \le s$, and set $r_w \leftarrow s + r_v$ (this will be the number of columns of a matrix that we will construct).
   On the other hand, let $(\mathfrak{m}_i)_{1\le i\le m}$ be the following moduli (in any order): $\mathfrak{p}^{z(\mathfrak{p},\ell)-v_\mathfrak{p}(\mathfrak{m})}$ for $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$, $\mathfrak{p}^{z(\mathfrak{p},\ell)-1}$ for $\mathfrak{p} \in S_\ell$; and in the case $\ell = 2$, $\mathfrak{m}'_\infty$, complement of $\mathfrak{m}_\infty$ in the set of real places. Finally, set $R \leftarrow m + r_c$, where m is the number of moduli just computed (this will be the number of *blocks* of rows).
6. [Compute discrete logarithms] Using Algorithms 4.2.17 and 4.2.18, compute the SNF of $(\mathbb{Z}_K/\mathfrak{m}_i)^*$ as well as $L_{\mathfrak{m}_i}(v_j)$ and $L_{\mathfrak{m}_i}(\beta_{\mathfrak{p}_{j'}})$ for all i such that $1 \le i \le m$, for all j such that $1 \le j \le r_v$, and for all j' such that $1 \le j' \le s$.
7. [Create big matrix] Construct a matrix M as follows. Let M_j be the jth column of M. Then M_j is obtained by concatenating the $L_{\mathfrak{m}_i}(v_j)$ for $1 \le i \le m$ (considered as column vectors), together with the zero vector with r_c components if $j \le r_v$, or with the r_c-component column vector $(p_{i,\mathfrak{p}_{j-r_v}})_{1\le i\le r_c}$ if $r_v < j \le r_w$. Finally, denote by $\overline{M}$ the matrix M reduced modulo ℓ, considered as a matrix with entries in $\mathbb{Z}/\ell\mathbb{Z}$.
8. [Compute kernel] Using [Coh0, Algorithm 2.3.1], compute the kernel $\mathcal{K}$ of the matrix $\overline{M}$ as a $\mathbb{Z}/\ell\mathbb{Z}$-vector space. If this kernel is reduced to $\{0\}$, there are no suitable Abelian extensions L/K, so terminate the algorithm. Otherwise,

let $d_{\mathcal{K}} \leftarrow \dim(\mathcal{K})$ be the dimension of this kernel, and denote by $(K_j)_{1\le j\le d_{\mathcal{K}}}$ a $\mathbb{Z}/\ell\mathbb{Z}$-basis of $\mathcal{K}$, where the K_j are considered as r_w-component column vectors. Finally, set $c \leftarrow d_{\mathcal{K}}$.

9. [Compute more discrete logarithms] Let $(\mathfrak{m}'_i)_{1\le i\le m'}$ be the following moduli (in any order): $\mathfrak{p}^{z(\mathfrak{p},\ell)-v_{\mathfrak{p}}(\mathfrak{m})+1}$ for $\mathfrak{p}\in S_{\mathfrak{m},\ell,2}$; and in the case $\ell=2$, $\mathfrak{m}_\infty$. As in step 6, compute the SNF of $(\mathbb{Z}_K/\mathfrak{m}'_i)^*$ as well as $L_{\mathfrak{m}'_i}(v_j)$ and $L_{\mathfrak{m}'_i}(\beta_{\mathfrak{p}_{j'}})$ for $1\le i\le m'$, $1\le j\le r_v$, and $1\le j'\le s$. For $1\le i\le m'$, let M'_i be the matrix with r_w columns, each column containing $L_{\mathfrak{m}'_i}(v_j)$ and $L_{\mathfrak{m}'_i}(\beta_{\mathfrak{p}_{j'}})$ as above. Do *not* put the matrices M'_i together by rows as above.
10. [Initialize backtracking] (In what follows, $c\ge 1$ and y will be a row vector with $c-1$ components.) Set $y\leftarrow(0,\dots,0)$ (vector with $c-1$ components).
11. [Compute trial vector] Let $X\leftarrow K_c+\sum_{1\le j<c} y_jK_j$. Apply Subalgorithm 5.2.15 below to see if X corresponds to a suitable Abelian extension L/K. If it does, set $\alpha=\prod_{1\le j\le r_w} v_j^{x_j}$ (where $X=(x_1,\dots,x_{r_w})^t$), and output the defining polynomial $X^\ell-\alpha=0$ (do not terminate the algorithm).
12. [Backtracking I] Set $i\leftarrow c$.
13. [Backtracking II] Set $i\leftarrow i-1$. If $i>0$, go to step 14. Otherwise, set $c\leftarrow c-1$. If $c>0$, go to step 10; otherwise, terminate the algorithm.
14. [Backtracking III] Set $y_i\leftarrow y_i+1$, and if $i<c-1$, set $y_{i+1}\leftarrow 0$. If $y_i\ge\ell$, go to step 13; otherwise, go to step 11.

Subalgorithm 5.2.15 (Is X Suitable?). Given a vector $X=(x_1,\dots,x_{r_w})^t$ found in step 11 of Algorithm 5.2.14, this subalgorithm determines whether X corresponds to a suitable Abelian extension L/K. We use all the quantities computed in the main algorithm.

1. [Test conditions on $x_{\mathfrak{p}}$] If any of the x_i for $r_v<i\le r_w$ is equal to zero modulo ℓ, X is not suitable, so terminate the subalgorithm.
2. [Test $\mathfrak{m}'_i$] For $1\le i\le m'$, compute $Y_i\leftarrow M'_iX$. If for any i we have $Y_i\equiv 0\pmod{\ell}$, then X is not suitable. Otherwise (in other words, if for all $i\le m'$ we have $Y_i\not\equiv 0\pmod{\ell}$), X is suitable. Terminate the subalgorithm.

Proof. Although the algorithm looks complicated, it is very little else than the exact algorithmic translation of Theorem 5.2.9. Thus, we simply make a few comments. We first want the modulus to satisfy conditions (1), (2), and (3) of Theorem 5.2.2. This is ensured by step 3. If

$$\alpha=\prod_{1\le j\le r_v} v_j^{n_j}\prod_{\mathfrak{p}\in S}\beta_{\mathfrak{p}}^{x_{\mathfrak{p}}}\ ,$$

we want a number of congruences and noncongruences to be satisfied, as well as conditions on the $x_{\mathfrak{p}}$. If we set $X=(n_1,\dots,n_{r_v},x_{\mathfrak{p}_1},\dots,x_{\mathfrak{p}_s})^t$, then it is easily seen that $X\in\mathrm{Ker}(\overline{M})$ is equivalent to the congruences that α must satisfy, together with the conditions $v_{\mathfrak{p}}(\alpha)\equiv 0\pmod{\ell}$ for $\mathfrak{p}\in S$, and so to

condition (5) of Theorem 5.2.9. Thus, at the end of step 8, all the elements X of the kernel $\mathcal{K}$ correspond to elements α satisfying the congruences.

We must now add some negative conditions: the noncongruences satisfied for prime ideals $\mathfrak{p} \in S_{\mathfrak{m},\ell,2}$, for $\sigma \in \mathfrak{m}_\infty$, and the conditions $x_\mathfrak{p} \not\equiv 0 \pmod{\ell}$ for $\mathfrak{p} \in S$. Instead of leading to the intersection of subspaces as before, this corresponds to the intersection of *complements* of subspaces. This is no longer linear algebra, and there does not seem to be any better method than complete enumeration, which at this stage should be rather short.

This is achieved by a standard backtracking procedure (steps 10 to 14), and the negative conditions for each trial vector are tested in Subalgorithm 5.2.15. Note that condition (3) of Theorem 5.2.9 (the possibility of setting some $x_\mathfrak{p} = 1$ if $S \neq \varnothing$) was included only to make the search faster by dehomogenizing the solution to the congruences, which is allowed up to Kummer-equivalence. In the above algorithm we proceeded differently (and more efficiently if $S = \varnothing$): if $X = \sum_{1 \le i \le c} y_i K_i$, we ask that $y_i = 1$ for the largest index i such that $y_i \neq 0$ (see also remark (3) below).

In Subalgorithm 5.2.15, we must also ensure that $\alpha \notin K^{*\ell}$. By Theorem 5.2.9, it is sufficient to ensure that $\alpha \neq 1$, and this is indeed excluded by our backtracking procedure. □

Remarks

(1) By Proposition 3.3.12, we know that $\mathfrak{m}$ must be a modulus for the congruence subgroup $P_\mathfrak{m}$, hence we can begin by checking that $h_{\mathfrak{m}/\mathfrak{p}} < h_\mathfrak{m}$ for all places $\mathfrak{p}$ dividing $\mathfrak{m}$, since if this not the case, there are no suitable extensions L. Computing $h_\mathfrak{m}$ and all the $h_{\mathfrak{m}/\mathfrak{p}}$ may, however, be costly, so it is not certain that this is worthwhile.

(2) Instead of putting all the discrete logarithm data into a big matrix, we could also consider computing the product of all the prime ideal powers modulo which congruences have to be taken, and compute a single discrete logarithm. This would almost certainly be slower than the method given above.

(3) The algorithm given in [Coh0, Algorithm 2.3.1] gives a basis of the kernel in column echelon form. If $S \neq \varnothing$, the last coordinate of the vector X is one of the $x_\mathfrak{p}$ and hence must be nonzero, so in step 13, when i gets down to zero it is not necessary to continue the backtracking with $c \leftarrow c - 1$, since all subsequent vectors X will be excluded by the subalgorithm. We could have included this remark explicitly in the algorithm, but its validity would then have been dependent on the algorithm chosen for computing the kernel in step 8.

We thus have finished our description of algorithmic Kummer theory when $\zeta_\ell \in K$ using Hecke's theorem.

5.3 Kummer Theory Using Hecke When $\zeta_\ell \notin K$

In this section, we no longer assume that $\zeta_\ell \in K$ (in fact, we explicitly assume that $\zeta_\ell \notin K$), and we want to give an algorithmic description of cyclic extensions L/K of degree ℓ and given conductor $\mathfrak{m}$.

Our first problem is simply to be able to *describe* them, in other words, to give explicit defining polynomials for cyclic extensions of degree ℓ of K. If $\zeta_\ell \in K$, Kummer theory tells us that any such extension is of the form $K(\sqrt[\ell]{\alpha})$ for some α, but here the situation is less simple. Even after this problem is solved, we need to control the ramification of prime ideals, and this is difficult to do directly.

Thus, the only method used in practice is to *adjoin* to K a primitive ℓth root of unity ζ_ℓ, thus obtaining a larger field $K_z = K(\zeta_\ell)$. We then apply Kummer theory to the field K_z, obtaining a cyclic extension L_z of K_z, of degree ℓ and having suitable properties, and we finally must go back down from L_z to the desired extension L. We will see in detail how this is done. Before doing this, however, we recall some basic facts about idempotents.

5.3.1 Eigenspace Decomposition for the Action of τ

We first need a well-known result in Galois theory, which we state as follows (recall that since we are in characteristic 0, the notions of normal and Galois coincide).

Proposition 5.3.1. *Let L be a number field, let L_1 and L_2 be two extensions of L included in a fixed algebraic closure $\overline{L}$ of L, and let L_1L_2 be the compositum of L_1 and L_2 in $\overline{L}$.*

(1) *If L_1/L and L_2/L are normal extensions, then L_1L_2/L is also a normal extension.*
(2) *If L_1/L and L_2/L are Abelian extensions, then L_1L_2/L is also an Abelian extension.*
(3) *Assume only that L_1/L is a normal extension with Galois group G_1. Then L_1L_2/L_2 is a normal extension whose Galois group H_1 can be canonically identified with a subgroup of G_1. Furthermore, H_1 is isomorphic to G_1 if and only if $L_1 \cap L_2 = L$.*

Proof. Let N be the normal closure of L_1L_2 (or any field containing N and normal over L), and let $\mathcal{G} = \mathrm{Gal}(N/L)$ be the Galois group of N/L. For $i = 1, 2$, let $\mathcal{G}_i = \mathrm{Gal}(N/L_i)$ so that $\mathcal{G}_i$ is a normal subgroup of $\mathcal{G}$ with $\mathrm{Gal}(L_i/L) \simeq \mathcal{G}/\mathcal{G}_i$. By Galois theory, subfields of N containing both L_1 and L_2 are in one-to-one correspondence with subgroups of $\mathcal{G}$ contained in $\mathcal{G}_1 \cap \mathcal{G}_2$, hence $\mathrm{Gal}(N/L_1L_2) = \mathcal{G}_1 \cap \mathcal{G}_2$. Since $\mathcal{G}_1 \cap \mathcal{G}_2$ is the intersection of two normal subgroups, it is also a normal subgroup, hence L_1L_2/L is normal with Galois group isomorphic to $\mathcal{G}/\mathcal{G}_1 \cap \mathcal{G}_2$, proving (1).

For (2), note that L_1/L is Abelian if and only if $\overline{xyx^{-1}y^{-1}} = \overline{1}$ in $\mathcal{G}/\mathcal{G}_1$ for all x, y in $\mathcal{G}$, hence if and only if $[\mathcal{G},\mathcal{G}] \subset \mathcal{G}_1$, where $[\mathcal{G},\mathcal{G}]$ is the commutator subgroup of $\mathcal{G}$ (in other words, L_1 is a subfield of the maximal Abelian extension of L included in N). Thus, if L_1/L and L_2/L are both Abelian, we have $[\mathcal{G},\mathcal{G}] \subset \mathcal{G}_1 \cap \mathcal{G}_2$, so $\mathrm{Gal}(L_1L_2/L)$ is Abelian, proving (2).

For (3), since $\mathcal{G}_1$ is a normal subgroup of $\mathcal{G}$, $\mathcal{G}_1 \cap \mathcal{G}_2$ is a normal subgroup of $\mathcal{G}_2$, so

$$H_1 \simeq \frac{\mathcal{G}_2}{\mathcal{G}_1 \cap \mathcal{G}_2} \simeq \frac{\mathcal{G}_1 + \mathcal{G}_2}{\mathcal{G}_1} \subset \frac{\mathcal{G}}{\mathcal{G}_1} \simeq G_1 \ ,$$

showing the result. In addition, we have the isomorphism $H_1 \simeq G_1$ if and only if $\mathcal{G}_1 + \mathcal{G}_2 = \mathcal{G}$, hence if and only if $L_1 \cap L_2 = L$, as claimed.

More directly, let $\sigma \in H_1$ and let $\tau = \sigma^{-1}$. If s (resp., t) is the restriction of σ (resp., τ) to L_1, then $s \circ t = t \circ s = Id$, hence s and t are automorphisms of L_1. Since σ and τ fix L_2 pointwise, s and t fix L pointwise, so $s \in G_1$, and thus this defines a canonical map from H_1 to G_1. If $s = Id$, then σ is the identity on L_1 and on L_2 by assumption, so σ is the identity on L_1L_2, hence $\sigma = Id$, showing that the map is injective, so H_1 can indeed be considered as a subgroup of G_1. If, in addition, L_1 and L_2 are linearly disjoint over L, then σ can be defined (uniquely) from the knowledge of s and t, so our map is also surjective, and the converse clearly also holds. □

We now come back to our specific situation. Let ℓ be a prime number, let K be a number field such that $\zeta_\ell \notin K$ (hence, in particular, $\ell > 2$), and set $K_z = K(\zeta_\ell)$. In the rest of this chapter we choose once and for all a primitive root g_0 modulo ℓ.

Proposition 5.3.2. *The extension $K_z/K = K(\zeta_\ell)/K$ is a cyclic extension of degree $d = (\ell-1)/m$ for some divisor m of $\ell - 1$ such that $m < \ell - 1$. The Galois group $\mathrm{Gal}(K_z/K)$ is generated by the automorphism τ of order d defined by $\tau(\zeta_\ell) = \zeta_\ell^g$ and $\tau(x) = x$ for $x \in K$, where $g = g_0^m$.*

Proof. We apply Proposition 5.3.1 (3) to the case $L = \mathbb{Q}$, $L_1 = \mathbb{Q}(\zeta_\ell)$, $L_2 = K$, hence $L_1L_2 = K(\zeta_\ell) = K_z$. Thus K_z/K is normal, and its Galois group can be identified with a subgroup of $\mathrm{Gal}(\mathbb{Q}(\zeta_\ell)/\mathbb{Q}) \simeq (\mathbb{Z}/\ell\mathbb{Z})^*$. Since this is a cyclic group of order $\ell-1$, $\mathrm{Gal}(K_z/K)$ is a cyclic group of order dividing $\ell-1$, hence of order $(\ell-1)/m$ for some $m < \ell - 1$, since we have assumed $\zeta_\ell \notin K$. Since $(\mathbb{Z}/\ell\mathbb{Z})^*$ has a unique subgroup of given order $(\ell-1)/m$, generated by $g = g_0^m$, the proposition follows. □

We will denote by G the Galois group of K_z/K, so that $G = \langle\tau\rangle$ is a cyclic group generated by τ of order $d = (\ell-1)/m$. The diagram of fields is as follows:

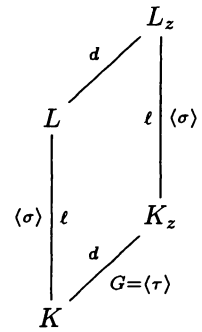

Let W be an $\mathbb{F}_\ell$-vector space (not necessarily finite dimensional) on which the group G operates. With our applications in mind, the Abelian group law of W will be written multiplicatively.

Then $\tau \in G$ acts as an endomorphism t of W, of order dividing d. Since $d \mid (\ell - 1)$ hence is coprime to ℓ, $X^d - 1$ is a squarefree polynomial in $\mathbb{F}_\ell[X]$, hence t is diagonalizable. Furthermore, the eigenvalues of t (in $\overline{\mathbb{F}_\ell}$) are among the roots of $X^d - 1 = 0$, hence are among the elements of $\mathbb{F}_\ell^*$ which are roots of this polynomial, and these are the powers of $g = g_0^m$. Therefore, we can write $W = \bigoplus_{0 \le k < d} W_k$, where W_k is the eigenspace corresponding to the eigenvalue g^k of τ acting on W.

For $0 \le k < d$, set

$$e_k = \frac{1}{d} \sum_{0 \le a < d} g^{-ka} \tau^a = -m \sum_{0 \le a < d} g^{-ka} \tau^a \in \mathbb{F}_\ell[G] \ .$$

Lemma 5.3.3. *The e_k for $0 \le k < d$ form a* complete set of orthogonal idempotents *for the action of G. In other words:*

(1) *if $k_1 \ne k_2$, then $e_{k_1} e_{k_2} = 0$;*
(2) *$e_k^2 = e_k$;*
(3) *we have $\sum_{0 \le k < d} e_k = 1$;*
(4) *we have $\tau e_k = g^k e_k$.*

Proof. The proof is a trivial direct verification: we have

$$\begin{aligned} e_{k_1} e_{k_2} &= d^2 \sum_{a,b} g^{-(k_1 a + k_2 b)} \tau^{a+b} \\ &= d^2 \sum_A \tau^A \sum_a g^{-(k_1 a + k_2 (A-a))} \\ &= d^2 \sum_A \tau^A g^{-k_2 A} \sum_a g^{a(k_2 - k_1)} \ , \end{aligned}$$

and the inner sum is a geometric series that vanishes if $k_2 \ne k_1$ and is equal to d if $k_2 = k_1$, showing (1) and (2). Statement (3) also follows immediately by summing a geometric series. Finally,

$$\tau e_k = -m \sum_a g^{-ka} \tau^{a+1} = -m \sum_a g^{-k(a-1)} \tau^a = g^k e_k \ ,$$

proving (4). □

Recall that it is common usage, and very useful, to use exponential notation for the action of group rings, so that if $x \in W$ and $e = \sum_{\sigma \in G} a_\sigma \sigma$, then

$$x^e = \prod_{\sigma \in G} \sigma(x)^{a_\sigma} \ .$$

Corollary 5.3.4. *With the above notation, the eigenspace W_k is equal to $e_k W = \{x^{e_k} / \ x \in W\}$.*

Proof. We have

$$\tau(x^{e_k}) = x^{\tau e_k} = x^{g^k e_k} = x^{e_k g^k} \ .$$

Thus, $x^{e_k} \in W_k$; hence $e_k W \subset W_k$. It follows that

$$\bigoplus_{0 \le k < d} e_k W \subset \bigoplus_{0 \le k < d} W_k = W \ .$$

Since the e_k form a complete set of orthogonal idempotents, we have $W = \bigoplus_{0 \le k < d} e_k W$ (since $x = \prod_{0 \le k < d} x^{e_k}$), and so we must have the equality $e_k W = W_k$ for all k, proving the corollary. □

We will use these results mainly for $W = K_z^*/K_z^{*\ell}$, $W = U(K_z)/U(K_z)^\ell$, $W = V_\ell(K_z)/K_z^{*\ell}$, and $W = Cl(K_z)/Cl(K_z)^\ell$. All these groups are $\mathbb{F}_\ell$-vector spaces that are stable by τ, and hence by G. Indeed, this is clear for K_z^*, the units, and the class group, while for the virtual units it follows from the characterization of virtual units given in Proposition 5.2.3 as elements generating the ℓth power of an *ideal*.

The basic theorem we will use is the following.

Theorem 5.3.5. *Let K be a number field, and let L be a cyclic extension of K of degree ℓ. Assume that K does not contain ζ_ℓ, and let $K_z = K(\zeta_\ell)$ and $L_z = L(\zeta_\ell)$. Let g_0 be a primitive root modulo ℓ, let $d = [K_z : K] = (\ell - 1)/m$, and $g = g_0^m$ as above. Finally, let $W = K_z^*/K_z^{*\ell}$.*

(1) *Any α such that $L_z = K_z(\sqrt[\ell]{\alpha})$ belongs to the eigenspace $e_1 W = W_1$ of W (and such α exist by Kummer theory).*
(2) *If $L_z = K_z(\theta)$ with $\theta = \sqrt[\ell]{\alpha}$ as in (1), then $L = K(\eta)$ with*

$$\eta = \mathrm{Tr}_{L_z/L}(\theta) = \sum_{0 \le a < d} \tau^a(\theta) \ ,$$

where τ is any extension to L_z of the K-automorphism τ of K_z.

(3) *A defining polynomial for L/K is given by the polynomial*

$$P(X) = \prod_{0 \le j < \ell} \left(X - \sum_{0 \le a < d} \zeta_\ell^{jg^a} \tau^a(\theta) \right) \in K[X] \ .$$

(4) *We have*

$$\theta = \frac{1}{\ell} \sum_{0 \le j < \ell} \zeta_\ell^{-j} \sigma^j(\eta) \ .$$

(5) *Conversely, if we are given a cyclic extension L of K of degree ℓ by $L = K(\eta)$ and if we define θ by the above formula, then $\alpha = \theta^\ell \in K(\zeta_\ell)$ and $\alpha \in W_1$.*

Proof. The extensions K_z/K and L/K are cyclic and have coprime degrees, hence by Proposition 5.3.1 L_z/K is an Abelian extension and $\mathrm{Gal}(L_z/L) \simeq \mathrm{Gal}(K_z/K) = \langle \tau \rangle$, and $\mathrm{Gal}(L_z/K_z) \simeq \mathrm{Gal}(L/K) = \langle \sigma \rangle$ for some σ of order ℓ. For any K-automorphism s of L and K-automorphism t of K_z, there exists a unique K-automorphism of L_z that extends both s and t, and it is defined in a natural way.

By a natural abuse of notation, we will denote by τ the unique K-automorphism of L_z that extends the K-automorphism τ of K_z and is the identity on L, and similarly we will denote by σ the unique K-automorphism of L_z that extends the K-automorphism σ of L and is the identity on K_z.

This being noted, let us prove the theorem. Since L_z is a cyclic extension of K_z of degree ℓ, by Kummer theory (Corollary 10.2.7 in this case) we know that there exists $\alpha \in K_z^*$ not in $K_z^{*\ell}$ such that $L_z = K_z(\theta)$ with $\theta = \sqrt[\ell]{\alpha}$. Since τ is an automorphism of L_z, we have $\tau(\theta) \in L_z$ and $\tau^{-1}(\theta) \in L_z$, from which it follows that $L_z = K_z(\tau(\theta))$. Since $\tau(\theta)^\ell = \tau(\alpha)$, Corollary 10.2.7 (2) tells us that there exists j coprime to ℓ (hence of the form g_0^a) and $\gamma \in K_z$ such that

$$\tau(\alpha) = \alpha^j \gamma^\ell = \alpha^{g_0^a} \gamma^\ell \ .$$

I first claim that a is a multiple of m (or, equivalently, g_0^a is a power of g). Indeed, if we compute $\tau^d(\alpha) = \alpha$, we obtain

$$\alpha = \tau^d(\alpha) = \alpha^{g_0^{da}} \delta^\ell$$

for some other $\delta \in K_z^*$. Thus $\alpha^{g_0^{da}-1}$ is the ℓth power of an element. If $g_0^{da} - 1$ is coprime to ℓ, then we can write $1 = u\ell + v(g_0^{da} - 1)$ for some u and v, and raising α to the power both sides, we see that α is an ℓth power, which is absurd. Thus $g_0^{da} \equiv 1 \pmod{\ell}$ or, equivalently, since g_0 is a primitive ℓth root of unity, $m \mid a$, as claimed.

We thus have proved that $\tau(\alpha) = \alpha^{g^k} \gamma^\ell$ for some integer k, so that $\alpha \in W_k$, the subspace of $W = K_z^*/K_z^{*\ell}$ corresponding to the eigenvalue g^k. Note that to prove this we have only used the fact that L_z/K is a *normal* extension.

We now use the fact that it is an *Abelian* extension to show that, in fact, $k = 1$.

Indeed, since $\tau(\alpha) = \alpha^{g^k}\gamma^\ell$, we have $\tau(\theta)^\ell = \theta^{\ell g^k}\gamma^\ell$; hence $\tau(\theta) = \theta^{g^k}\gamma'$ for some other $\gamma' \in K_z$ since $\zeta_\ell \in K_z$. On the other hand, θ is a root of the irreducible polynomial $X^\ell - \alpha \in K_z[X]$, hence $\sigma(\theta) = \zeta_\ell^r\theta$ for some r, and r is necessarily coprime to ℓ, otherwise σ would be equal to the identity. In fact, if desired, by changing the generator σ, we may assume that $r = 1$.

Since $\tau(\zeta_\ell) = \zeta_\ell^g$, we obtain

$$\tau(\sigma(\theta)) = \zeta_\ell^{gr}\theta^{g^k}\gamma'$$

while

$$\sigma(\tau(\theta)) = \zeta_\ell^{g^k r}\theta^{g^k}\gamma' \ .$$

Thus, since σ and τ commute, we obtain $(g^k - g)r \equiv 0 \pmod{\ell}$, hence $g^k \equiv g \pmod{\ell}$ since r is coprime to ℓ, and hence $k \equiv 1 \pmod{m}$ since g_0 is a primitive root modulo ℓ, proving (1).

For (2), we clearly have

$$\eta = \mathrm{Tr}_{L_z/L}(\theta) = \sum_{0 \le a < d} \tau^a(\theta) \in L \ .$$

Since L/K is of prime degree ℓ, to show that $L = K(\eta)$, we must simply show that $\eta \notin K$. Assume the contrary. We then have $\sigma^k(\eta) = \eta$ for all k. Since τ and σ commute, we obtain the system of equations

$$\sum_{0 \le a < d} \zeta_\ell^{kg^a}\tau^a(\theta) = \eta \quad \text{for} \quad 0 \le k \le \ell - 1 \ .$$

If we restrict to the first d equations, we have a system of d equations in the d unknowns $\tau^a(\theta)$ whose determinant is the Vandermonde determinant corresponding to the variables $\zeta_\ell^{g^a}$ for $0 \le a < d$. Since $g = g_0^m$ and g_0 is a primitive root modulo ℓ, these variables are distinct, hence the determinant is nonzero. Since $\zeta_\ell \in K_z$ and $\eta \in K$ by assumption, it follows that for all a, $\tau^a(\theta) \in K_z$ and in particular $\theta \in K_z$, so $\alpha \in K_z^\ell$, which is absurd, proving (2).

For (3), note that the minimal polynomial of η in $K[X]$ is given by $P(X) = \prod_{0 \le j < \ell}(X - \sigma^j(\eta))$. As already mentioned, we may choose σ so that $\sigma(\theta) = \zeta_\ell\theta$. It follows that

$$\sigma^j(\eta) = \sigma^j\left(\sum_a \tau^a(\theta)\right) = \sum_a \tau^a(\zeta_\ell^j\theta) = \sum_a \zeta_\ell^{jg^a}\tau^a(\theta) \ ,$$

as claimed.

For (4), a direct computation using the commutativity of σ and τ gives

$$\begin{aligned}\sum_j \zeta_\ell^{-j}\sigma^j(\eta) &= \sum_j \zeta_\ell^{-j}\sum_a \sigma^j(\tau^a(\theta)) &&= \sum_j \zeta_\ell^{-j}\sum_a \tau^a(\sigma^j(\theta))\\ &= \sum_j \zeta_\ell^{-j}\sum_a \tau^a(\zeta_\ell^j\theta) &&= \sum_j \zeta_\ell^{-j}\sum_a \zeta_\ell^{jg^a}\tau^a(\theta)\\ &= \sum_a \tau^a(\theta)\sum_j \zeta_\ell^{j(g^a-1)} \ .\end{aligned}$$

The inner sum vanishes unless $g^a - 1 \equiv 0 \pmod{\ell}$, hence unless $a = 0$, so $\sum_j \zeta_\ell^{-j}\sigma^j(\eta) = \ell\theta$, as claimed.

For (5), we have

$$\sigma(\theta) = \frac{1}{\ell}\sum_j \zeta_\ell^{-(j-1)}\sigma^j(\eta) = \zeta_\ell\theta \ .$$

It follows that $\alpha = \theta^\ell$ is invariant by σ, hence by Galois theory it belongs to K_z. Similarly, we see that $\beta_1 = \tau(\theta)/\theta^g$ is invariant by σ, hence belongs to K_z. Thus, $\tau(\alpha)/\alpha^g = \beta_1^\ell \in K_z^{*\ell}$, so $\alpha \in W_1$, finishing the proof of the theorem. □

Remark. The generating element θ of L_z, constructed either via K_z and Kummer theory or directly using (4), is called the *Lagrange resolvent* of the extension L/K.

5.3.2 Lift in Characteristic 0

To perform actual algorithmic computations, we must look explicitly at the situation in characteristic 0 and not in characteristic ℓ; in other words, we must consider $\mathbb{Z}[G]$-modules and not $\mathbb{F}_\ell[G]$-modules.

Consider first the generator $g = g_0^m$ of the subgroup of order m of $(\mathbb{Z}/\ell\mathbb{Z})^*$. In the preceding section, we could consider at will g to be an element of $\mathbb{Z}$ or an element of $\mathbb{Z}/\ell\mathbb{Z}$. From now on, we specifically ask that g be considered as an element of $\mathbb{Z}$. In fact, to simplify many of the computations, we will ask in addition that $g > 0$. As we will see in the next proposition, we also need the technical condition that $g^d \not\equiv 1 \pmod{\ell^2}$. This is easily achieved, since if $g^d \equiv 1 \pmod{\ell^2}$, then

$$(g+\ell)^d \equiv 1 + d\ell g^{d-1} \not\equiv 1 \pmod{\ell^2} \ ,$$

so we simply replace g by $g + \ell$. When $d = 1$ or, equivalently, when $g = g_0$ we are simply asking that g_0 is a primitive root modulo ℓ^2. Note that we cannot start from a primitive root modulo ℓ^2 and define g as $g_0^m \bmod \ell$, since we reduce modulo ℓ and not modulo ℓ^2.

Next, recall that the idempotent e_1 is given by $e_1 = -m\sum_{0\le a<d} g^{-a}\tau^a \in \mathbb{F}_\ell[G]$. Set

$$\lambda_0 = \sum_{0\le a<d} g^{d-1-a}\tau^a \ .$$

If λ is any element in $\mathbb{Z}[G]$ such that $\lambda \equiv \lambda_0 \pmod{\ell}$, then since g and m are coprime to ℓ, we have $\overline{\lambda} \equiv ce_1 \pmod{\ell}$ for some $c \in \mathbb{F}_\ell^*$. Hence if W is an $\mathbb{F}_\ell$-vector space, we have $e_1 W = \overline{\lambda} W$. We will use λ instead of e_1 in our statements, and in the next section we will discuss how to choose it efficiently. For now, the choice is not important.

Let $\alpha \in K_z^*$. Assume that, considered as an element of $K_z^*/K_z^{*\ell}$, we know that $\alpha \in W_1$, the g^1-eigenspace of W for the action of τ. These are exactly the α that we need to construct our extensions L_z/K_z. On the level of K_z^* itself, this means that $\tau(\alpha) = \alpha^g\gamma^\ell$ for some $\gamma \in K_z^*$.

Since $W_1 = e_1 W = \overline{\lambda} W$, we know that $\alpha = \beta^\lambda \delta^\ell$ for some $\beta \in K_z^*$. We want to compute β explicitly, and for this we prove the following proposition.

Proposition 5.3.6. *Let $\alpha \in K_z^*$ be such that $\tau(\alpha) = \alpha^g\gamma^\ell$ for some $\gamma \in K_z^*$. Then $\alpha = \beta^\lambda\delta^\ell$, where*

$$\beta = \gamma^{-a}\zeta_\ell^k \quad \textit{with} \quad a = \left(\frac{g^d-1}{\ell}\right)^{-1} \pmod{\ell}$$

for some integer k and some $\delta \in K_z^*$.

Conversely, let β be given such that $\alpha = \beta^\lambda$, and let $\lambda = \lambda_0 + \nu\ell$ for $\nu \in \mathbb{Z}[G]$ with λ_0 as above. Then $\tau(\alpha) = \alpha^g\gamma^\ell$ with

$$\gamma = \beta^{-(g^d-1)/\ell+(\tau-g)\nu} \ .$$

Proof. It is easy to prove by induction on a that for any $a \ge 0$,

$$\tau^a(\alpha) = \alpha^{g^a}\gamma^{\ell(g^a-\tau^a)/(g-\tau)} \ .$$

Applying this formula to $a = d$, we obtain

$$\alpha = \alpha^{g^d}\gamma^{\lambda_0\ell}$$

with λ_0 as above. Thus,

$$\alpha^{(g^d-1)/\ell} = \gamma^{-\lambda_0}\zeta_\ell^u = \gamma^{-\lambda}\varepsilon^\ell\zeta_\ell^u$$

for some integer u and some $\varepsilon \in K_z^*$.

We now use for the first and only time the technical condition $g^d \not\equiv 1 \pmod{\ell^2}$ imposed above on the generator g. This implies that $(g^d-1)/\ell \not\equiv 0 \pmod{\ell}$, so we can find integers a and b such that $a(g^d-1)/\ell + b\ell = 1$. It follows that

$$\alpha = \alpha^{a(g^d-1)/\ell}\alpha^{b\ell} = \gamma^{-a\lambda}\delta^{\ell}\zeta_{\ell}^{au}$$

for some $\delta \in K_z^*$.

Finally, note that $\zeta_\ell^\lambda = \zeta_\ell^{dg^{d-1}}$, so that $\zeta_\ell^{-gm\lambda} = \zeta_\ell$. We thus obtain

$$\alpha = (\gamma^{-a}\zeta_\ell^{-gmau})^\lambda \delta^\ell \ ,$$

proving the first part of the proposition.

For the converse, we have

$$\frac{\tau(\alpha)}{\alpha^g} = \alpha^{\tau-g} = \beta^{(\tau-g)\lambda} = \beta^{(\tau-g)(\lambda_0+\ell\nu)} \ .$$

Since $(\tau - g)\lambda_0 = 1 - g^d$, we have $\tau(\alpha)/\alpha^g = \gamma^\ell$ with $\gamma = \beta^{-(g^d-1)/\ell+(\tau-g)\nu}$, as claimed. □

Corollary 5.3.7. *Keep the notation of Theorem 5.3.5 and Proposition 5.3.6. Set* $\mu = -(g^d - 1)/\ell + (\tau - g)\nu$. *Then if* $\alpha = \beta^\lambda$, *we can take* $\tau(\theta) = \theta^g\beta^\mu$.

Proof. We have $\theta^\ell = \alpha$. With any initial choice of τ, we have $\tau(\theta)^\ell = \tau(\alpha) = \alpha^g\gamma^\ell$ with $\gamma = \beta^\mu$. Thus, $\tau(\theta) = \theta^g\beta^\mu\zeta_\ell^k$ for some integer k. If we set $\tau' = \sigma^{-g^{-1}k}\tau$, then τ' also extends τ, and we have

$$\tau'(\theta) = \theta^g\zeta_\ell^{-gg^{-1}k}\zeta_\ell^k\beta^\mu = \theta^g\beta^\mu \ ,$$

so the result follows. Of course, the fixed field of L_z by τ' is still equal to L since L is normal over K. □

Let us now discuss the choice of λ. Theoretically, it has no importance, but algorithmically, the situation is different. Since we will take $\alpha = \beta^\lambda$, we must choose λ as simple as possible. A reasonable choice is to ask that the coefficients of τ^a in λ are all between 0 and $\ell - 1$. In other words, if we set

$$r_a = g^a \bmod \ell = g^a - \ell\left\lfloor\frac{g^a}{\ell}\right\rfloor \ ,$$

it is reasonable to set

$$\lambda = \sum_{0\le a<d} r_{d-1-a}\tau^a \ .$$

It is an easy exercise (Exercise 5) to show that this definition is independent of the choice of the primitive root g (recall that τ depends also on g). With this choice, we have the following.

Lemma 5.3.8. *Keep all the above notation, and choose*

$$\lambda = \sum_{0 \le a < d} r_{d-1-a} \tau^a \ .$$

Then $\tau^b(\theta) = \theta^{r_b} \beta^{\mu_b}$ *with*

$$\mu_b = -\sum_{0 \le a < d} \left\lfloor \frac{r_b r_{d-1-a}}{\ell} \right\rfloor \tau^a \ .$$

Proof. This follows from a direct computation. Note that r_a is periodic of period d for $a \ge 0$, hence it is reasonable to extend it by periodicity to all integral a. By the above corollary, we have $\tau(\theta) = \theta^g \beta^\mu$, with

$$\mu = -\frac{g^d - 1}{\ell} + (\tau - g)\frac{\lambda - \lambda_0}{\ell} = -\sum_{0 \le a < d} \frac{g r_{d-1-a} - r_{d-a}}{\ell} \tau^a \ .$$

By induction, we immediately obtain

$$\tau^b(\theta) = \theta^{g^b} \beta^{(g^b - \tau^b)/(g-\tau)\mu} = \theta^{r_b} \beta^{\mu_b}$$

with

$$\mu_b = \frac{g^b - r_b}{\ell}\lambda + \frac{g^b - \tau^b}{g - \tau}\mu \ .$$

Since

$$\tau^k \mu = -\sum_a \frac{g r_{d-1-a+k} - r_{d-a+k}}{\ell} \tau^a \ ,$$

the series $\sum_{0 \le k \le b-1} g^{b-1-k} \tau^k \mu$ telescopes and we obtain

$$\frac{g^b - \tau^b}{g - \tau}\mu = \sum_{0 \le k \le b-1} g^{b-1-k} \tau^k \mu = -\sum_a \frac{g^b r_{d-1-a} - r_{d-1-a+b}}{\ell} \tau^a \ .$$

It follows that

$$\mu_b = -\sum_a \frac{r_b r_{d-1-a} - r_{d-1-a+b}}{\ell} \tau^a = -\sum_a \left\lfloor \frac{r_b r_{d-1-a}}{\ell} \right\rfloor \tau^a \ ,$$

as claimed. □

We can now give a relatively explicit form of the polynomial $P(X)$ given in Theorem 5.3.5.

Proposition 5.3.9. *Keep the notation of Theorem 5.3.5. For* $2 \le k \le \ell$, *set*

$$t(b_1, \ldots, b_{k-1}) = \frac{1}{\ell} \sum_{0 \le a < d} \tau^a \left(r_{d-1-a} - \ell + \sum_{1 \le i \le k-1} r_{d-1-a+b_i} \right)$$

and define γ_k *by the formula*

$$\gamma_k = \sum_{\substack{b_1 \le \cdots \le b_{k-1} \\ r_{b_1}+\cdots+r_{b_{k-1}}+1 \equiv 0 \pmod{\ell}}} \frac{(k-1)!}{\prod m_j!} \beta^{t(b_1,\ldots,b_{k-1})} ,$$

where the m_j *denote the multiplicities of the* b_i. *Set* $e = \mathcal{N}_{K_z/K}(\beta)$ *and* $S_k = e\ell \operatorname{Tr}_{K_z/K}(\gamma_k)$. *If we define* $\mathfrak{S}_k$ *by the usual Newton recursion*

$$k\mathfrak{S}_k = \sum_{1 \le i \le k} (-1)^{i+1} S_i \mathfrak{S}_{k-i} ,$$

a defining polynomial $P(X)$ *for the* L/K *is given by*

$$P(X) = \sum_{0 \le k \le \ell} (-1)^k \mathfrak{S}_k X^{\ell-k} .$$

Proof. By Theorem 5.3.5 and Lemma 5.3.8, the kth power sum of the roots of the polynomial $P(X)$ is given by

$$\begin{aligned} S_k &= \sum_{0 \le j < \ell} \left(\sum_{0 \le b < d} \zeta_\ell^{jg^b} \theta^{r_b} \beta^{\mu_b} \right)^k \\ &= \sum_{0 \le j < \ell} \left(\sum_{b_1,\ldots,b_k} \zeta_\ell^{j(g^{b_1}+\cdots g^{b_k})} \theta^{r_{b_1}+\cdots+r_{b_k}} \beta^{\mu_{b_1}+\cdots+\mu_{b_k}} \right) \\ &= \ell \sum_{r_{b_1}+\cdots+r_{b_k} \equiv 0 \pmod{\ell}} \beta^{e(b_1,\ldots,b_k)} , \end{aligned}$$

where

$$\begin{aligned} e(b_1,\ldots,b_k) &= \sum_{0 \le a < d} \tau^a \left(\sum_{1 \le i \le k} \frac{r_{b_i} r_{d-1-a}}{\ell} - \sum_{1 \le i \le k} \left\lfloor \frac{r_{b_i} r_{d-1-a}}{\ell} \right\rfloor \right) \\ &= \frac{1}{\ell} \sum_{0 \le a < d} \tau^a \left(\sum_{1 \le i \le k} r_{d-1-a+b_i} \right) . \end{aligned}$$

It follows in particular from this formula that

$$\tau e(b_1,\ldots,b_k) = e(b_1+1,\ldots,b_k+1) ,$$

hence that

$$S_k = \ell \operatorname{Tr}_{K_z/K} \left(\sum_{r_{b_1}+\cdots+r_{b_{k-1}}+r_0 \equiv 0 \pmod{\ell}} \beta^{e(b_1,\ldots,b_{k-1},0)} \right) ,$$

so that

$$S_k = \ell \operatorname{Tr}_{K_z/K}\left(\sum_{r_{b_1}+\cdots+r_{b_{k-1}}+1\equiv 0 \pmod{\ell}} \beta^{t'(b_1,\ldots,b_{k-1})}\right)$$

with

$$t'(b_1,\ldots,b_{k-1}) = \frac{1}{\ell}\sum_{0\le a<d} \tau^a\left(r_{d-1-a} + \sum_{1\le i\le k-1} r_{d-1-a+b_i}\right) .$$

To finish the proof, we simply notice that each coefficient of τ is strictly positive, hence we can factor out the norm $N_{K_z/K}(\beta) = \beta^{\sum_{0\le a<d}\tau^a}$, and furthermore the summands are symmetrical in the b_i, so it is enough to sum for $b_1 \le \cdots \le b_{k-1}$, except that we must compensate by the multinomial coefficient that counts the number of $(k-1)$-tuples $(b_1,\ldots,b_{k-1})$ corresponding to a given nondecreasing sequence of b_i. □

Remark. We have chosen r_b in the interval $0 \le r_b < \ell$ and not, for instance, in the interval $-\ell/2 < r_b < \ell/2$, so that the coefficients of τ^a in $t(b_1,\ldots,b_{k-1})$ are nonnegative. If we had done otherwise, we would have obtained relative and absolute defining polynomials with nonintegral coefficients.

Examples.
Let us look at the simplest cases of Theorem 5.3.5 and Proposition 5.3.9.

(1) If $\ell = 3$ and $d = 2$, we take $g = 2$ and we have $\alpha = \beta^{2+\tau}$, $\tau(\theta) = \theta^2/\beta$, $\eta = \theta + \theta^2/\beta$, and a computation gives

$$P(X) = X^3 - 3eX - e\operatorname{Tr}_{K_z/K}(\beta)$$

with $e = \beta^{1+\tau} = \mathcal{N}_{K_z/K}(\beta)$. This is exactly the formula of [Coh0, Lemma 6.4.5] for cyclic cubic fields, but this time for cyclic cubic extensions of any base field.

(2) If $\ell = 5$ and $d = 4$, we take $g = 2$, and we have $\alpha = \beta^{3+4\tau+2\tau^2+\tau^3}$, $\tau(\theta) = \theta^2/\beta^{1+\tau}$, $\tau^2(\theta) = \theta^4/\beta^{2+3\tau+\tau^2}$, $\tau^3(\theta) = \theta^3/\beta^{1+2\tau+\tau^2}$,

$$\eta = \theta + \theta^2/\beta^{1+\tau} + \theta^3/\beta^{1+2\tau+\tau^2} + \theta^4/\beta^{2+3\tau+\tau^2} ,$$

and a computation gives

$$P(X) = X^5 - 10eX^3 - 5e\operatorname{Tr}_{K_z/K}(\beta^{1+\tau})X^2 + 5e(e - \operatorname{Tr}_{K_z/K}(\beta^{1+2\tau+\tau^2}))X - e\operatorname{Tr}_{K_z/K}(\beta^{2+3\tau+\tau^2})$$

with

$$e = \beta^{1+\tau+\tau^2+\tau^3} = \mathcal{N}_{K_z/K}(\beta) .$$

(3) If $\ell = 5$ and $d = 2$, we still take $g_0 = 2$, hence $g = g_0^2 = 4$, and we have $\alpha = \beta^{4+\tau}$, $\tau(\theta) = \theta^4/\beta^3$, $\eta = \theta + \theta^4/\beta^3$, and a computation gives

$$P(X) = X^5 - 5eX^3 + 5e^2X - e\,\mathrm{Tr}_{K_z/K}(\beta^3)$$

with $e = \beta^{1+\tau} = \mathcal{N}_{K_z/K}(\beta)$.

With the help of a computer algebra package, the reader can calculate the formulas for larger values of ℓ (see Exercises 6 and 7).

Before we give the theorems and algorithms that will enable us to compute cyclic extensions of a number field K when $\zeta_\ell \notin K$, we must explain in detail the action of τ on the different objects that we need, in other words the units, the virtual units, and the class group. To simplify notation, in the next subsections only, by abuse of notation, we simply write K instead of $K_z = K(\zeta_\ell)$.

5.3.3 Action of τ on Units

As usual, let ε_0 be a generator of the group of torsion units, and let $(\varepsilon_1, \ldots, \varepsilon_{r_u})$ be a system of fundamental units. If ε is a unit, then $\tau(\varepsilon)$ is also a unit, which we want to express on the ε_i. We have already implicitly mentioned this problem several times, including in [Coh0], but the algorithm is worth writing out explicitly.

Algorithm 5.3.10 (Discrete Logarithm in the Unit Group). Let $\boldsymbol{\mu}(K)$ be the group of roots of unity of K, let $w(K) = |\boldsymbol{\mu}(K)|$, let ε_0 be a generator of $\boldsymbol{\mu}(K)$ (a primitive $w(K)$th root of unity), and let $(\varepsilon_1, \ldots, \varepsilon_{r_u})$ be a system of fundamental units of K. If ε is a unit of K, this algorithm computes the discrete logarithm of ε with respect to the ε_i, in other words, exponents x_i such that $\varepsilon = \prod_{0 \le i \le r_u} \varepsilon_i^{x_i}$. We let σ_i for $1 \le i \le r_1 + 2r_2$ be the embeddings of K into $\mathbb{C}$ ordered in the usual way.

1. [Compute real logarithmic embeddings] Compute the $r_u \times r_u$ matrix $A = (n_i \log|\sigma_i(\varepsilon_j)|)_{1 \le i,j \le r_u}$ of the real logarithmic embeddings of the fundamental units, where as usual $n_i = 1$ if σ_i is a real embedding, and $n_i = 2$ otherwise. We omit one of the σ_i so as to have a square matrix of determinant equal (up to sign) to the regulator $R(K)$.
2. [Solve system] Similarly, compute the column vector $B = (n_i \log|\sigma_i(\varepsilon)|)_{1 \le i \le r_u}$ of the real logarithmic embeddings of ε, omitting the same embedding as in step 1. Using Gaussian elimination over $\mathbb{R}$, let Y be a solution of the linear system $AY = B$ (since $\det(A) = \pm R(K)$, this solution Y exists and is unique).
3. [Check correctness] If the entries of Y are not close to integers (say, further away than 10^{-5}), increase the accuracy of the computations and start again in step 1 (or output an error message saying that ε is not a unit). Otherwise, if $Y = (y_1, \ldots, y_{r_u})^t$ for $1 \le i \le r_u$, set $x_i \leftarrow \lfloor y_i \rceil$, and compute exactly (as an algebraic number) $\eta \leftarrow \varepsilon / \prod_{1 \le i \le r_u} \varepsilon_i^{x_i}$.

4. [Compute x_0] Let $a \leftarrow w(K)\log(\eta)/(2i\pi)$, $b \leftarrow w(K)\log(\varepsilon_0)/(2i\pi)$. If a is not close to an integer, increase the accuracy of the computations and start again in step 1 (or output an error message saying that ε is not a unit). Otherwise, set $a \leftarrow \lfloor a \rceil$, $b \leftarrow \lfloor b \rceil$, $x_0 \leftarrow ab^{-1} \bmod w(K)$.
5. [Terminate] If η is not equal to $\varepsilon_0^{x_0}$, increase the accuracy of the computations and start again in step 1. Otherwise, output $(x_0, x_1, \ldots, x_{r_u})$ and terminate the algorithm.

Remark. This algorithm's validity is clear. Note that in step 4 we do not check that b is close to an integer since ε_0 must be a generator of the roots of unity of K.

Using this algorithm, we can thus construct an $(r_u+1)\times(r_u+1)$ matrix $T_u = (t_{i,j})$ such that

$$\tau(\varepsilon_j) = \prod_{0\le i\le r_u} \varepsilon_i^{t_{i,j}}$$

for $0 \le j \le r_u$.

Let $E = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{r_u})$. We thus have $\tau(E) = ET_u$. The subspace $e_1U(K)/U(K)^\ell$, which is of interest to us, is the $\mathbb{F}_\ell$-vector space of classes of units ε such that $\varepsilon^{\tau-g} \in U(K)^\ell$ or, equivalently, by Lemma 5.2.6, such that $\varepsilon^{\tau-g} \in K^{*\ell}$. Thus, if $\varepsilon = EX$ for some integer column vector X, we want $E(T_u - gI_{r_u+1})X \equiv 0 \pmod{\ell}$, and since the classes of elements of E form a basis of $U(K)/U(K)^\ell$ by Proposition 5.2.5, this is equivalent to $\overline{X} \in \mathrm{Ker}\left(\overline{T_u - gI_{r_u+1}}\right)$, where $\overline{\ }$ denotes reduction modulo ℓ. Thus, using ordinary Gaussian pivoting in the field $\mathbb{F}_\ell$, we compute a basis of this kernel, and we obtain a basis of $e_1U(K)/U(K)^\ell$.

5.3.4 Action of τ on Virtual Units

We must now solve the same problems in the larger groups $V_\ell(K)$ and $V_\ell(K)/K^{*\ell}$. Let $(\alpha_1, \ldots, \alpha_{r_c}, \varepsilon_0, \ldots, \varepsilon_{r_u})$ be as above. If v is a virtual unit, then $\tau(v)$ generates the ℓth power of an ideal, hence it is also a virtual unit, and we want to express it on the α_i, ε_j, and ℓth powers of elements. For this, we use the following algorithm.

Algorithm 5.3.11 (Discrete Logarithm in the ℓ-Selmer Group). As above, let $Cl(K) = \bigoplus_{1\le i\le g_c}(\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$ and $\alpha_i\mathbb{Z}_K = \mathfrak{a}_i^{d_i}$, and let $v \in V_\ell(K)$ be a virtual unit. This algorithm computes the discrete logarithm of v with respect to the α_i for $1 \le i \le r_c$ and ε_j for $0 \le j \le r_u$, in other words, exponents y_i and x_j such that

$$v = \gamma^\ell \prod_{1\le i\le r_c} \alpha_i^{y_i} \prod_{0\le j\le r_u} \varepsilon_j^{x_j}$$

for some $\gamma \in K^{*\ell}$.

1. [Factor $v\mathbb{Z}_K$] Using Algorithm 2.3.22, factor the ideal $v\mathbb{Z}_K$ into a product of prime ideal powers as $v\mathbb{Z}_K = \prod_{\mathfrak{p}} \mathfrak{p}^{a_\mathfrak{p}}$. If any one of the $a_\mathfrak{p}$ is not divisible by

ℓ, output an error message saying that v is not a virtual unit, and terminate the algorithm.

2. [Compute ideal] Using ideal multiplication and powering algorithms (see Section 2.3.4), compute $\mathfrak{q} \leftarrow \prod_{\mathfrak{p}} \mathfrak{p}^{a_{\mathfrak{p}}/\ell}$.
3. [Use principal ideal algorithm] Using [Coh0, Algorithm 6.5.10], compute $\beta \in K^*$ and integers q_i such that $\mathfrak{q} = \beta \prod_{1 \le i \le g_c} \mathfrak{a}_i^{q_i}$.
4. [Compute γ and y_i] For $1 \le i \le r_c$, set $y_i \leftarrow q_i/(d_i/\ell)$, and set

$$\gamma \leftarrow \beta \prod_{r_c < i \le g_c} \alpha_i^{q_i/d_i} \quad \text{and} \quad \varepsilon \leftarrow v/(\gamma^{\ell} \prod_{1 \le i \le r_c} \alpha_i^{y_i}) \ .$$

5. [Terminate] (Here ε is a unit.) Using Algorithm 5.3.10, compute the discrete logarithm $(x_0, \ldots, x_{r_u})$ of ε. Output the x_j, y_i, γ, and terminate the algorithm.

Proof. Left to the reader (Exercise 8). □

Using this algorithm, we can thus construct an $r_v \times r_v$ matrix $T_v = (t_{i,j})$ such that

$$\tau(v_j) = \beta_j^{\ell} \prod_{1 \le i \le r_v} v_i^{t_{i,j}}$$

for some $\beta_j \in K^*$, where we use the notation v_i introduced in Definition 5.2.7 for the α_i and ε_j (recall that $r_v = r_c + r_u + 1$ is the ℓ-rank of the ℓ-Selmer group).

If $E_v = (v_1, \ldots, v_{r_v})$, we have $\tau(E_v) \equiv ET_v \pmod{\ell}$ (more precisely modulo ℓth powers of elements of K^*). Thus, if $v = E_v X$ is a virtual unit, $v \in e_1 V_\ell(K)/K^{*\ell}$ if and only if $E_v(T_v - gI_{r_v})X \equiv 0 \pmod{\ell}$, and since the classes of elements of E_v form a basis of $V_\ell(K)/K^{*\ell}$ by Proposition 5.2.5, as in the case of units we conclude that the kernel of $\overline{T_v - gI_{r_v}}$ gives us an $\mathbb{F}_\ell$-basis of $e_1 V_\ell(K)/K^{*\ell}$.

5.3.5 Action of τ on the Class Group

We must solve a similar problem as above, but now in the class group $Cl(K)$. The situation is more complicated in this case since the class group is already a set of equivalence classes. More precisely, if $Cl(K) = \bigoplus_{1 \le i \le g_c} (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$, the principal ideal algorithm ([Coh0, Algorithm 6.5.10]) allows us to write

$$\tau(\mathfrak{a}_j) = \beta_j \prod_{1 \le i \le g_c} \mathfrak{a}_i^{t_{i,j}}$$

for certain integers $t_{i,j}$ and certain elements $\beta_j \in K^*$. The presence of these elements β_j creates unnecessary complications, so we want to get rid of them. For this, we use the following proposition.

Proposition 5.3.12. *Let $Cl(K) = \bigoplus_{1\le i\le g_c}(\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$ be the SNF of the class group $Cl(K)$, and as usual let r_c be the largest index i such that $\ell \mid d_i$. There exist representatives $\mathfrak{b}_i$ of the ideal classes $\overline{\mathfrak{a}_i}$ such that for all j such that $1 \le j \le r_c$, we have*

$$\tau(\mathfrak{b}_j) = \mathfrak{q}_j^\ell \prod_{1\le i\le r_c} \mathfrak{b}_i^{t_{i,j}}$$

for some ideal $\mathfrak{q}_j$ and some integers $t_{i,j}$.

As mentioned above, the main point of this proposition is to make the elements β_j disappear. On the other hand, the occurrence of the ideal $\mathfrak{q}_j^\ell$ is perfectly natural since we work implicitly modulo ℓth powers of ideal classes.

Proof. For the (arbitrarily chosen) representatives $\mathfrak{a}_i$ of the ideal classes $\overline{\mathfrak{a}_i}$, we can write

$$\tau(\mathfrak{a}_j) = \beta_j \mathfrak{q}_j^\ell \prod_{1\le i\le r_c} \mathfrak{a}_i^{t_{i,j}}$$

for some elements $\beta_j \in K^*$, ideals $\mathfrak{q}_j$, and integers $t_{i,j}$. If we set $\mathfrak{b}_i = \gamma_i\mathfrak{a}_i$, we have

$$\tau(\mathfrak{b}_j) = \frac{\tau(\gamma_j)}{\prod_{1\le i\le r_c}\gamma_i^{t_{i,j}}}\beta_j\mathfrak{q}_j^\ell \prod_{1\le i\le r_c} \mathfrak{b}_i^{t_{i,j}} \ .$$

To satisfy the conditions of the proposition, we must show the existence (and compute explicitly) elements γ_i such that

$$\frac{\tau(\gamma_j)}{\prod_{1\le i\le r_c}\gamma_i^{t_{i,j}}} = \frac{w_j}{\beta_j}$$

for some virtual unit w_j, since by definition of virtual units, we have $w_j\mathbb{Z}_K = \mathfrak{q}^\ell$ for some ideal $\mathfrak{q}$.

Let $T = (t_{i,j})_{1\le i,j\le r_c}$ be the matrix of the $t_{i,j}$, let $B = (\beta_1,\dots,\beta_{r_c})$ (resp., $C = (\gamma_1,\dots,\gamma_{r_c})$) be the row vector of the β_i (resp., of the γ_i). Recall that we have set $[K_z : K] = d = (\ell-1)/m$. I claim that the elements γ_i defined by

$$C = d\cdot B\Big(\sum_{1\le a<d} a\tau^a T^{d-a-1}\Big)$$

satisfy our requirements (we use the notation $B\tau^a T^{d-a-1}$ to denote the row vector $\tau^a(B)T^{d-a-1}$).

Indeed,

$$\begin{aligned}(\tau - T)\sum_{1\le a<d} a\tau^a T^{d-a-1} &= \sum_{2\le a\le d}(a-1)\tau^a T^{d-a} - \sum_{1\le a<d} a\tau^a T^{d-a}\\ &= -\sum_{0\le a<d}\tau^a T^{d-a} + T^d - I_{r_c} + dI_{r_c}\ ;\end{aligned}$$

hence, if we set $U = \sum_{0 \le a < d} \tau^a T^{d-1-a}$, we have

$$\tau(C) - CT = d \cdot B(-UT + T^d - I_{r_c} + dI_{r_c}) \ .$$

On the other hand, if we let $\mathfrak{A} = (\mathfrak{a}_1, \ldots, \mathfrak{a}_{r_c})$ be the row vector of the ideals $\mathfrak{a}_i$, the formula for $\tau(\mathfrak{a}_j)$ given above can be written as $\tau(\mathfrak{A}) \equiv \mathfrak{A}T + B \pmod{\ell}$, from which it easily follows by induction that

$$\tau^a(\mathfrak{A}) \equiv \mathfrak{A}T^a + B\frac{\tau^a - T^a}{\tau - T} \pmod{\ell} \ .$$

Applying this to $a = d$, we obtain $\mathfrak{A} \equiv \mathfrak{A}T^d + BU \pmod{\ell}$. By the uniqueness of the representation on the $\mathfrak{a}_i$, it follows that $T^d \equiv I_{r_c} \pmod{\ell}$ and $BU \equiv 0 \pmod{\ell}$, hence the elements of BU are both elements of K^* and ℓth powers of ideals, so they are virtual units.

Since $dm \equiv -1 \pmod{\ell}$, we have $\tau(C) - CT \equiv W - B \pmod{\ell}$ for some vector $W = (w_j)_{1 \le j \le r_c}$ of virtual units, as was to be proved. □

Important Remark. The implicit operation used between elements to define the γ_i is always multiplication, no additions are involved here. In particular, the initial multiplication by m is in fact raising to the mth power, and multiplication by the powers of the matrix T are combinations of multiplications and powerings, *not* of additions and multiplications.

It is of course very easy to compute the discrete logarithm (more precisely, to solve the principal ideal problem) on the $\mathfrak{b}_i$. If $\mathfrak{a}$ is any ideal of K, we use Algorithm 5.2.10 to compute an element α and integers x_i such that $\mathfrak{a} = \alpha \mathfrak{q}^\ell \prod_{1 \le i \le r_c} \mathfrak{a}_i^{x_i}$ for some ideal $\mathfrak{q}$. It follows that

$$\mathfrak{a} = \frac{\alpha}{\prod_{1 \le i \le r_c} \gamma_i^{x_i}} \mathfrak{q}^\ell \prod_{1 \le i \le r_c} \mathfrak{b}_i^{x_i}$$

is the desired decomposition of $\mathfrak{a}$.

Thus, let $T_c = T = (t_{i,j})$ be the $r_c \times r_c$ matrix defined above, which in particular satisfies $T_c^d \equiv I_{r_c} \pmod{\ell}$. If $\mathfrak{B} = (\mathfrak{b}_1, \ldots, \mathfrak{b}_{r_c})$, we have $\tau(\mathfrak{B}) \equiv \mathfrak{B}T \pmod{\ell}$. As in the preceding cases, to obtain an $\mathbb{F}_\ell$-basis of $e_1 Cl(K)/Cl(K)^\ell$, we simply compute a matrix P whose columns give an $\mathbb{F}_\ell$-basis of the kernel of the matrix $\overline{T_c - gI_{r_c}}$. If δ is the dimension of this kernel, the desired basis is then

$$(\mathfrak{c}_1, \ldots, \mathfrak{c}_\delta) = \mathfrak{C} = \mathfrak{B}P \ .$$

We will also need the following technical but important proposition.

Proposition 5.3.13. *Let P (resp., Q) be a matrix whose columns give an $\mathbb{F}_\ell$-basis of the kernel of $\overline{T_c - gI_{r_c}}$ (resp., of $\overline{T_c^t - gI_{r_c}}$, where as usual T^t denotes the transposed matrix of T). Let $R = (Q^t P)^{-1}$. Then*

$$\mathfrak{B}^\lambda \equiv dg^{-1}\mathfrak{C}(RQ^t) \pmod{\ell} \ .$$

Remark. Since in general P is not a square matrix, we cannot write $RQ^t = P^{-1}(Q^t)^{-1}Q^t = P^{-1}$.

Proof. Since $\lambda(\tau - g) \equiv 0 \pmod{\ell}$, for any ideal $\mathfrak{a}$ we know that the class of $\mathfrak{a}^\lambda$ belongs to $e_1 Cl(K)/Cl(K)^\ell$; hence it is expressible on the ideals $\mathfrak{c}_i$.

Since $\tau(\mathfrak{B}) \equiv \mathfrak{B}T \pmod{\ell}$, we have

$$\mathfrak{B}^\lambda \equiv \mathfrak{B} \sum_{0 \le a < d} g^{d-1-a} T^a \pmod{\ell} \ .$$

Since $T^d \equiv I_{r_c} \pmod{\ell}$ and d is a divisor of $\ell - 1$, the matrix T is diagonalizable in $\mathbb{F}_\ell$. In other words, we can write $M^{-1}TM \equiv D \pmod{\ell}$ for some invertible matrix M and diagonal matrix D. Since the eigenvalues of T are powers of g, we can write D in block diagonal form:

$$D = \operatorname{diag}(g^0 I_{\delta_0}, g^1 I_{\delta_1}, \dots, g^{d-1} I_{\delta_{d-1}}) \ ,$$

where δ_k is the dimension of the eigenspace corresponding to the eigenvalue g^k, in other words, the dimension of $W_k = e_k Cl(K)/Cl(K)^\ell$. Correspondingly, we can write $M = (M_0, \dots, M_{d-1})$ and $(M^{-1})^t = (M'_0, \dots, M'_{d-1})$, where M_k and M'_k are $h \times \delta_k$ matrices.

We of course have $\delta_1 = \delta$, and we can choose $M_1 = P$. Thus,

$$\begin{aligned}
\sum_{0 \le a < d} g^{d-1-a} T^a &\equiv M \operatorname{diag}\left(\left(\sum_{0 \le a < d} g^{d-1-a} g^{ka} I_{\delta_k}\right)\right) M^{-1} \\
&\equiv M \operatorname{diag}(0, dg^{-1} I_\delta, 0, \dots, 0) M^{-1} \\
&\equiv dg^{-1}(0, M_1, 0, \dots, 0)(M'_0, \dots, M'_{d-1})^t \\
&\equiv dg^{-1} M_1 (M'_1)^t \pmod{\ell} \ .
\end{aligned}$$

Since we have chosen $M_1 = P$, we thus have

$$\mathfrak{B}^\lambda \equiv dg^{-1}\mathfrak{B}M_1(M'_1)^t \equiv dg^{-1}\mathfrak{C}(M'_1)^t \pmod{\ell} \ .$$

As the columns of both Q and of M'_1 give a basis of the kernel of $\overline{T_c^t - gI_{r_c}}$, there exists an invertible matrix R^t such that $M'_1 = QR^t$ or, equivalently, $(M'_1)^t = RQ^t$. On the other hand, the identity $M^{-1}M = I_{r_c}$ gives $(M'_i)^t M_j = \delta_{i,j} I_{\delta_j}$, where $\delta_{i,j}$ is the Kronecker symbol (no confusion should occur with δ_j). In particular, $(M'_1)^t M_1 = I_\delta$, and so we obtain $RQ^t M_1 = I_\delta$ so $R = (Q^t M_1)^{-1} = (Q^t P)^{-1}$. It follows that $(M'_1)^t = RQ^t$ with $R = (Q^t P)^{-1}$ so

$$\mathfrak{B}^\lambda \equiv dg^{-1}\mathfrak{C}RQ^t \pmod{\ell} \ ,$$

as claimed. Note that for $j \neq 1$, $0 = (M'_1)^t M_j = RQ^t M_j$, hence $Q^t M_j = 0$ since R is invertible. □

Corollary 5.3.14. *Let*

$$\mathfrak{p} = \beta \mathfrak{q}^{\ell} \prod_{1 \le i \le r_c} \mathfrak{b}_i^{x_i}$$

be the decomposition of some ideal $\mathfrak{p}$ *on the* $\mathfrak{b}_i$*, and let* X *be the column vector of the* x_i*. Then*

$$\mathfrak{p}^{\lambda} = \beta^{\lambda} \mathfrak{q}_1^{\ell} \prod_{1 \le i \le \delta} \mathfrak{c}_i^{y_i} \ ,$$

where, if Y *is the column vector of the* y_i *we have* $Y = dg^{-1} R Q^t X$*.*

Proof. In matrix terms, we have $\mathfrak{p} \equiv \beta + \mathfrak{B} X \pmod{\ell}$. Thus, by Proposition 5.3.13 we have

$$\mathfrak{p}^{\lambda} \equiv \beta^{\lambda} + \mathfrak{B}^{\lambda} X \equiv \beta^{\lambda} + dg^{-1} \mathfrak{C}(RQ^t X) \pmod{\ell} \ ,$$

proving the corollary. □

5.3.6 Algorithmic Kummer Theory When $\zeta_\ell \notin K$ Using Hecke

Thanks to the results of the preceding sections, we can compute relative defining polynomials for cyclic extensions of prime degree ℓ, even when $\zeta_\ell \notin K$. We now describe in detail the choice of α (or, equivalently, of β) in K_z, and then give the complete algorithm.

Let $(\mathfrak{m}, C)$ be a congruence subgroup of the base field K. We want to give an explicit defining polynomial for the (isomorphism class of) Abelian extension(s) L/K corresponding to $(\mathfrak{m}, C)$ by class field theory. By Proposition 3.5.5, the extension L_z/K_z considered above corresponds to the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K_z}, \mathcal{N}_{K_z/K}^{-1}(C))$ of K_z. Unfortunately, even if $\mathfrak{m}$ is the conductor of L/K, $\mathfrak{m}\mathbb{Z}_{K_z}$ will not necessarily be the conductor of L_z/K_z. Thus, before proceeding, we use Algorithm 4.4.2 to compute the conductor $\mathfrak{f}$ of the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K_z}, \mathcal{N}_{K_z/K}^{-1}(C))$ of K_z, which will be a divisor of $\mathfrak{m}\mathbb{Z}_{K_z}$. In addition, $\mathfrak{m}\mathbb{Z}_{K_z}$ is evidently invariant by τ (or, equivalently, by $\mathrm{Gal}(K_z/K)$), but it is easily shown that $\mathfrak{f}$ is also invariant by τ. Thus, the six sets $S_{\mathfrak{f},\ell,i}$, $S_{\mathfrak{f}}$, S_ℓ and $S_\varnothing$ are invariant by τ, and we can thus consider a system of representatives for the action of τ. Until further notice (more precisely, until we start going down to L), we work in the field K_z and not in the field K.

We keep all the notation of the preceding sections, in particular that concerning the action of τ. Thus, we write

$$Cl(K_z) = \bigoplus_{1 \le i \le g_c} (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i} \ ,$$

r_c is the largest index such that $\ell \mid d_i$, $\alpha_i \mathbb{Z}_{K_z} = \mathfrak{a}_i^{d_i}$, the class modulo $K_z^{*\ell}$ of $(\varepsilon_0, \ldots, \varepsilon_{r_u}, \alpha_1, \ldots, \alpha_{r_c})$ is an $\mathbb{F}_\ell$-basis of the ℓ-Selmer group of K_z, $\mathfrak{b}_i = \gamma_i \mathfrak{a}_i$

are the representatives of the ideal classes $\overline{\mathfrak{a}_i}$ defined above, and finally the class of $\mathfrak{C} = (\mathfrak{c}_1, \dots, \mathfrak{c}_\delta)$ is a basis of $e_1 Cl(K_z)/Cl(K_z)^\ell$ as obtained above.

Recall that $S = S_\mathfrak{f} \cup S_{\mathfrak{f},\ell,1}$. We will denote by $S/\langle\tau\rangle$ a system of representatives of S for the action of τ, and similarly for the other sets of prime ideals.

Theorem 5.2.9 generalizes as follows.

Theorem 5.3.15. *Keep the notation of the preceding section. In particular, K is a number field and ℓ is a prime number such that $\zeta_\ell \notin K$.*

Let L/K be a cyclic extension of degree ℓ corresponding to the congruence subgroup $(\mathfrak{m}, C)$, let $K_z = K(\zeta_\ell)$ and $L_z = L(\zeta_\ell)$, and let $\mathfrak{f} = \mathfrak{f}(L_z/K_z)$ be the conductor of the extension L_z/K_z (or, equivalently, of the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K_z}, \mathcal{N}_{K_z/K}^{-1}(C))$ of K_z).

Let $W_1 = e_1 V_\ell(K_z)/K_z^{*\ell}$ be the g^1-eigenspace of the ℓ-Selmer group of K_z for the action of τ, let d_v be its dimension, and let $(\overline{w_i})_{1\le i\le d_v}$ be an $\mathbb{F}_\ell$-basis of W_1 computed as explained in Section 5.3.4.

For each prime ideal $\mathfrak{p} \in S/\langle\tau\rangle$, write

$$\mathfrak{p} = \beta_\mathfrak{p} \mathfrak{q}_\mathfrak{p}^\ell \prod_{1\le i\le r_c} \mathfrak{b}_i^{p_{i,\mathfrak{p}}} .$$

Then $\mathfrak{f}$ satisfies the following conditions:

(1) $S_{\mathfrak{f},\ell,3} = \emptyset$;
(2) *if* $\mathfrak{p} \in S_{\mathfrak{f},\ell,2}/\langle\tau\rangle$, *then* $v_\mathfrak{p}(\mathfrak{f}) \not\equiv 1 \pmod{\ell}$;
(3) *if* $\mathfrak{p} \in S_\mathfrak{f}/\langle\tau\rangle$, *then* $v_\mathfrak{p}(\mathfrak{f}) = 1$.

In addition, up to Kummer-equivalence, we can take $\alpha = \beta^\lambda$ with β of the following form:

$$\beta = \prod_{1\le i\le d_v} w_i^{n_i} \prod_{\mathfrak{p}\in S/\langle\tau\rangle} \beta_\mathfrak{p}^{x_\mathfrak{p}} ,$$

with the following additional conditions.

(1) *For all $\mathfrak{p} \in S/\langle\tau\rangle$, we have $1 \le x_\mathfrak{p} \le \ell - 1$, and for all i, we have $0 \le n_i \le \ell - 1$.*
(2) *For $\mathfrak{p} \in S_{\mathfrak{f},\ell,2}/\langle\tau\rangle$, the largest k such that the congruence*

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+v_\mathfrak{p}(\alpha)}}$$

has a solution must be equal to $z(\mathfrak{p},\ell) - v_\mathfrak{p}(\mathfrak{f})$.
(3) *If S is not empty, we may fix any one (but only one) of the $x_\mathfrak{p}$ equal to 1.*
(4) *For all $\mathfrak{p} \in S_\ell/\langle\tau\rangle$, the congruence*

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{z(\mathfrak{p},\ell)-1+v_\mathfrak{p}(\alpha)}}$$

has a solution.

(5) *For all* $\mathfrak{p} \in S/\langle\tau\rangle$, *let* $X_{\mathfrak{p}}$ *be the column vector of the* $p_{i,\mathfrak{p}}$ *for* $1 \le i \le r_c$. *We have*

$$Q^t \sum_{\mathfrak{p}\in S/\langle\tau\rangle} x_{\mathfrak{p}} X_{\mathfrak{p}} \equiv 0 \pmod{\ell} \ ,$$

where Q *is as in Proposition 5.3.13.*

Conversely, if all the above conditions are satisfied, if $\alpha \neq 1$, *and if the norm group* $T_{\mathfrak{m}}(L/K)$ *is equal to* C, *then* $L = K(\mathrm{Tr}_{L_z/L}(\sqrt[\ell]{\alpha}))$ *is a cyclic extension of degree* ℓ *corresponding to the congruence subgroup* $(\mathfrak{m}, C)$.

Proof. It follows from Theorem 5.3.5 and Proposition 5.3.6 that, up to Kummer-equivalence, we can take α such that $\overline{\alpha}$ is in the g^1-eigenspace W_1 for the action of τ on $K_z^*/K_z^{*\ell}$.

On the other hand, following the same proof as for Theorem 5.2.9, we find that, up to Kummer-equivalence, we can take

$$\alpha\mathbb{Z}_{K_z} = \mathfrak{q}^{\ell} \prod_{\mathfrak{p}\in S} \mathfrak{p}^{x_{\mathfrak{p}}} \ ,$$

where $\mathfrak{q}$ is an ideal coprime to $\mathfrak{f}$ and ℓ and $1 \le x_{\mathfrak{p}} \le \ell - 1$. The condition $\overline{\alpha} \in W_1$ means that $\tau(\alpha)/\alpha^g$ is an ℓth power in K_z^*. Since $\mathfrak{f}$ is stable by τ, we have

$$\tau(\alpha)\mathbb{Z}_{K_z} = \tau(\mathfrak{q})^{\ell} \prod_{\mathfrak{p}\in S} \tau(\mathfrak{p})^{x_{\mathfrak{p}}} = \mathfrak{q}_1^{\ell} \prod_{\mathfrak{p}\in S} \mathfrak{p}^{x_{\tau^{-1}(\mathfrak{p})}}$$

for some ideal $\mathfrak{q}_1$. Since $\mathfrak{q}$ and $\mathfrak{q}_1$ are coprime to the ideals in S, it follows that

$$x_{\tau^{-1}(\mathfrak{p})} \equiv g x_{\mathfrak{p}} \pmod{\ell}$$

or, equivalently,

$$x_{\mathfrak{p}} \equiv g x_{\tau(\mathfrak{p})} \pmod{\ell} \ .$$

This immediately implies the following lemma.

Lemma 5.3.16. *Let* $\mathfrak{p} \in S = S_{\mathfrak{f}} \cup S_{\mathfrak{f},\ell,1}$, *and assume that there exists a cyclic extension* L/K *such that the corresponding extension* L_z/K_z *is of degree* ℓ *and conductor* $\mathfrak{f}$. *Then the ideals* $\tau^j(\mathfrak{p})$ *for* $0 \le j < d$ *are distinct; in other words, the prime ideal of* K *below* $\mathfrak{p}$ *is totally split in* K_z.

Proof. Let j be some index such that $\tau^j(\mathfrak{p}) = \mathfrak{p}$. Applying the recursion for $x_{\mathfrak{p}}$, we obtain

$$x_{\mathfrak{p}} \equiv g^j x_{\tau^j(\mathfrak{p})} = g^j x_{\mathfrak{p}} \pmod{\ell} \ ,$$

and since $\mathfrak{p} \in S$, we have $x_{\mathfrak{p}} \not\equiv 0 \pmod{\ell}$, so $g^j \equiv 1 \pmod{\ell}$ and j is a multiple of d, as claimed. □

Resuming the proof of the theorem, we thus have

$$\alpha\mathbb{Z}_{K_z} = \mathfrak{q}^\ell \prod_{\mathfrak{p}\in S} \mathfrak{p}^{x_\mathfrak{p}} = \mathfrak{q}_1^\ell \prod_{\mathfrak{p}\in S/\langle\tau\rangle} (\mathfrak{p}^\lambda)^{gx_\mathfrak{p}} \ .$$

Replacing $\mathfrak{p}^\lambda$ by the expression given by Corollary 5.3.14, we obtain

$$\alpha\mathbb{Z}_{K_z} = \mathfrak{q}_2^\ell \prod_{\mathfrak{p}\in S/\langle\tau\rangle} \beta_\mathfrak{p}^{\lambda g x_\mathfrak{p}} \prod_{1\le i\le\delta} \mathfrak{c}_i^{z_i} \ ,$$

where, if Z is the column vector of the z_i, we have

$$Z = dRQ^t \sum_{\mathfrak{p}\in S/\langle\tau\rangle} x_\mathfrak{p} X_\mathfrak{p} \ .$$

Since the $\overline{\mathfrak{c}_i}$ form an $\mathbb{F}_\ell$-basis of $e_1 Cl(K_z)/Cl(K_z)^\ell$, it follows that we must have $Z \equiv 0 \pmod{\ell}$, and since d and R are invertible modulo ℓ, this gives the condition $Q^t \sum_{\mathfrak{p}\in S/\langle\tau\rangle} x_\mathfrak{p} X_\mathfrak{p} \equiv 0 \pmod{\ell}$ of the theorem.

Hence, replacing $gx_\mathfrak{p}$ by $x_\mathfrak{p}$ (which is legitimate since it will still be nonzero modulo ℓ), we obtain

$$\alpha\mathbb{Z}_{K_z} = \mathfrak{q}_3^\ell \prod_{\mathfrak{p}\in S/\langle\tau\rangle} \beta_\mathfrak{p}^{\lambda x_\mathfrak{p}} \ .$$

Thus $\mathfrak{q}_3^\ell$ is the ℓth power of an ideal and is also a principal ideal; hence by Proposition 5.2.3, it is of the form $v_0\mathbb{Z}_K$ for some virtual unit $v_0 \in V_\ell(K)$. Hence, for some other virtual unit v we obtain

$$\alpha = v \prod_{\mathfrak{p}\in S/\langle\tau\rangle} \beta_\mathfrak{p}^{\lambda x_\mathfrak{p}} \ .$$

Since up to ℓth powers both α and the $\beta_\mathfrak{p}^\lambda$ belong to the g^1-eigenspace under the action of τ, it follows that v does also. Thus, up to ℓth powers, since e_1 is an idempotent and λ is proportional to e_1 modulo ℓ, we have $v = w^\lambda = (w^{e_1})^\lambda$, and we have obtained that, up to Kummer-equivalence, $\alpha = \beta^\lambda$ with

$$\beta = w^{e_1} \prod_{\mathfrak{p}\in S/\langle\tau\rangle} \beta_\mathfrak{p}^{x_\mathfrak{p}} \ ,$$

proving that α is of the given form.

The conditions satisfied by $\mathfrak{f}$ and α are of course the same as in Theorem 5.2.9.

Conversely, assume that the conditions of the theorem are satisfied. Since g is coprime to ℓ, the solubility of a congruence of the type $x^\ell \equiv \alpha^g \pmod{I}$ for some ideal I is equivalent to the solubility of the congruence $x^\ell \equiv \alpha \pmod{I}$. Since our six sets of primes are stable under τ, it follows from

$\tau(\alpha) = \alpha^g \gamma^\ell$ for some $\gamma \in K_z$ that the congruence conditions are stable by τ. Thus, it is enough to check the congruence conditions for a system of representatives of the prime ideals modulo the action of τ, which is precisely what the theorem states. We conclude by Theorem 5.2.9 that $K_z(\sqrt[\ell]{\alpha})/K_z$ is a cyclic extension of degree ℓ and conductor equal to $\mathfrak{f}$. By Theorem 5.3.5, $L = K(\mathrm{Tr}_{L_z/L}(\sqrt[\ell]{\alpha}))$ is a cyclic extension of degree ℓ corresponding to the congruence subgroup $\mathfrak{m}$. Finally, we must simply check that the norm group $T_{\mathfrak{m}}(L/K)$ is equal to C. □

Before giving the final algorithm, we must explain how to compute the invariants of the field K_z and, in particular, all the necessary tools for doing class field theory. It is essential to have a representation of K_z/K that is as simple as possible. There are at least three distinct methods for achieving this goal. The first one is blindly to use Algorithm 2.1.8, which tries elements of the form $\theta_1 + k\zeta_\ell$, combined with Theorem 2.1.14 for factoring the discriminant of the resulting polynomial. Once an initial polynomial is computed, it is absolutely necessary to use a reduction procedure such as the Polred algorithm or improvements based on a Fincke–Pohst enumeration.

A second method for computing K_z is to use Algorithm 2.1.9, which tries elements of the form $\zeta_\ell(\theta_1 + k)$, combined with the analog of Theorem 2.1.14 proved in Exercise 9 of Chapter 2. This almost always gives much simpler polynomials, and in addition we can usually take $k = 0$ (see Exercise 15). Even so, before starting the class group computations on K_z, we must still apply variants of the Polred algorithm.

A third method for computing K_z is to consider it as a *relative* extension K_z/K and do all the computations using relative algorithms. This is certainly the best of all methods. Indeed, the computation of an integral pseudo-basis becomes very simple since the discriminant of the polynomial defining K_z/K is a power of ℓ. It follows that essentially no factoring needs to be done. Once we have obtained an integral pseudo-basis, we may, if desired, compute an absolute pseudo-basis, as explained in Section 2.5.2. Once this computation is done, we do not need to factor any large numbers, but we still need to compute several invariants of K_z, by using either relative or absolute algorithms. For the class and unit group computations, we may use the relative methods briefly described in Chapter 7. However, for the ray class group computations, we need to write and implement specific algorithms using the relative representation. While not difficult, this requires quite a lot of work. Thus we suggest using the relative methods to compute the integral pseudo-basis and possibly also the class and unit group, then convert to the absolute representation and compute the necessary ray class groups using this representation. For simplicity of exposition, in the sequel we will describe the computation of the invariants of K_z using the second method, but a serious implementation should use the third method instead.

We are now ready to give the complete algorithm for computing a defining polynomial for a cyclic extension of prime degree and given conductor using Kummer theory. Contrary to Algorithm 5.2.14, which computed all the cyclic extensions of degree ℓ and given conductor, because of the bad behavior of the conductor at primes dividing ℓ when we adjoin ζ_ℓ, it is necessary to restrict to one congruence subgroup at a time.

Algorithm 5.3.17 (Kummer Extensions of Prime Degree When $\zeta_\ell \notin K$ Using Hecke). Let $K = \mathbb{Q}(\theta_1)$ be a number field and ℓ be a prime number such that $\zeta_\ell \notin K$. Let $(\mathfrak{m}, C)$ be a congruence subgroup of K such that $h_{\mathfrak{m},C} = |I_\mathfrak{m}/C| = \left|Cl_\mathfrak{m}(K)/\overline{C}\right| = \ell$. This algorithm outputs a defining polynomial for the Abelian extension L/K of degree ℓ (which is unique up to isomorphism) corresponding to the congruence subgroup $(\mathfrak{m}, C)$ by Takagi's theorem.

1. [Adjoin ζ_ℓ] Using Algorithm 2.1.9, compute the compositum K_z of K with $\mathbb{Q}(\zeta_\ell)$ given by the defining polynomial $x^{\ell-1} + \cdots + x + 1 = 0$ (there may be several factors in the compositum, but they all define isomorphic number fields so any one can be taken). The algorithm outputs an integer k, an irreducible monic polynomial $R(X) \in \mathbb{Z}[X]$ having $\theta = \zeta_\ell(\theta_1 + k)$ as a root such that $K_z = \mathbb{Q}(\theta)$, and polynomials $A_1(X)$ and $A_2(X)$ such that $\theta_1 = A_1(\theta)$ and $\zeta_\ell = A_2(\theta)$.
2. [Compute the action of τ on θ] Let $d \leftarrow [K_z : K]$, $m \leftarrow (\ell-1)/d$. Using a naive algorithm, compute a primitive root g_0 modulo ℓ and set $g \leftarrow g_0^m \bmod \ell$. If $g^d \equiv 1 \pmod{\ell^2}$, set $g \leftarrow g + \ell$. Finally, compute the polynomial $U(X) \leftarrow XA_2(X)^{g-1} \bmod R(X)$ (we will have $U(\theta) = \tau(\theta)$).
3. [Compute data for K_z] Using the standard algorithms for the absolute case given in [Coh0], compute an integral basis, the unit group $U(K_z)$ with generators $(\varepsilon_0, \ldots, \varepsilon_{r_u})$, the class group in SNF as $Cl(K_z) = \bigoplus_{1 \le i \le g_c} (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$, the ℓ-rank r_c of the class group (maximum of the i such that $\ell \mid d_i$), and the α_i such that $\alpha_i \mathbb{Z}_{K_z} = \mathfrak{a}_i^{d_i}$ for $1 \le i \le g_c$ (only the α_i for $1 \le i \le r_c$ will be virtual units, but we need the others in Algorithm 5.2.10). As usual, we let v_i be the virtual units $(\alpha_1, \ldots, \alpha_{r_c}, \varepsilon_0, \ldots, \varepsilon_{r_u})$. If we use absolute algorithms, it will be necessary to factor the discriminant of the polynomial $R(X)$, which can be done if necessary by using the analog of Theorem 2.1.14 proved in Exercise 9 of Chapter 2.
4. [Adjust representatives $\mathfrak{a}_i$] Using the polynomial U obtained in step 2, compute $\tau(\mathfrak{a}_j)$ for $1 \le j \le r_c$, and by using Algorithm 5.2.10, compute elements β_j and integers $t_{i,j}$ such that

$$\tau(\mathfrak{a}_j) = \beta_j \mathfrak{q}_j^\ell \prod_{1 \le i \le r_c} \mathfrak{a}_i^{t_{i,j}} \ .$$

Let $T_c = (t_{i,j})$, let $\mathcal{B} = (\beta_j)$, and define $\mathcal{C} = (\gamma_j)$ modulo ℓth powers by

$$\mathcal{C} \leftarrow d\left(\sum_{1 \le a < d} a\tau^a(\mathcal{B}) T_c^{d-a-1} \right) ,$$

where the action of τ is once again given by the polynomial U, and the implicit operations are multiplicative (see the remark following the proof of Proposition 5.3.12). We could now set $\mathfrak{b}_i = \gamma_i\mathfrak{a}_i$ for $1 \le i \le r_c$, but we do not need the ideals $\mathfrak{b}_i$ explicitly.

5. [Compute $e_1V_\ell(K_z)/K_z^{*\ell}$] Using the polynomial U obtained in step 2, compute the $\tau(v_j)$ for $1 \le j \le r_v = r_c + r_u + 1$, and express them in terms of the v_i using Algorithm 5.3.11. If $T_v = (t_{i,j})$ is the matrix giving the $\tau(v_j)$ in terms of the v_i, using [Coh0, Algorithm 2.3.1], compute an $\mathbb{F}_\ell$-basis $P = (p_{i,j})$ of the kernel of the $\mathbb{F}_\ell$-matrix $\overline{T_v - gI_{r_v}}$. Let d_v be the dimension of this kernel, and for $1 \le j \le d_v$ set $w_j \leftarrow \prod_{1\le i\le r_v} v_i^{p_{i,j}}$ so that $\left(\overline{w_j}\right)_{1\le j\le d_v}$ is an $\mathbb{F}_\ell$-basis of $e_1V_\ell(K_z)/K_z^{*\ell}$.

6. [Compute data for $e_1Cl(K_z)/Cl(K_z)^\ell$] In a similar manner, compute a matrix Q whose columns give an $\mathbb{F}_\ell$-basis of the kernel of the matrix $\overline{T_c^t - gI_{r_c}}$, where T_c is the matrix computed in step 4. Let d_c be the dimension of this kernel.

7. [Compute conductor of L/K] Using Algorithm 4.4.2, replace $(\mathfrak{m}, C)$ by the conductor of the class of $(\mathfrak{m}, C)$.

8. [Lift congruence subgroup] If $\mathfrak{m}$ is coprime to ℓ, set $\mathfrak{f} \leftarrow \mathfrak{m}\mathbb{Z}_{K_z}$. Otherwise, proceed as follows. Set $\mathfrak{m}' \leftarrow \mathfrak{m}\mathbb{Z}_{K_z}$ and compute the SNF of $Cl_{\mathfrak{m}'}(K_z)$. Using Algorithm 4.1.11, compute $C' \leftarrow \mathcal{N}_{K_z/K}^{-1}(C)$. Finally, using Algorithm 4.4.2, compute the conductor $\mathfrak{f}$ of the congruence subgroup $(\mathfrak{m}', C')$.

9. [Factor $\mathfrak{f}$ and ℓ] (From now on, the algorithm will be very similar to Algorithm 5.2.14.) Using Algorithm 2.3.22 (in the absolute case), find the prime ideal factorization of $\mathfrak{f}$, say $\mathfrak{f} = \prod_{\mathfrak{p}|\mathfrak{f}} \mathfrak{p}^{v_\mathfrak{p}(\mathfrak{f})}$, and using [Coh0, Algorithm 6.2.9], compute the prime ideal factorization of $\ell\mathbb{Z}_{K_z}$.

10. [Compute sets of prime ideals] Compute the finite sets $S_\mathfrak{f}/\langle\tau\rangle$, $S_\ell/\langle\tau\rangle$, and $S_{\mathfrak{f},\ell,i}/\langle\tau\rangle$ for $i = 1, 2, 3$ according to Definition 5.2.1 (in other words, compute the sets and keep only one representative modulo the action of τ).

11. [Test conditions on $\mathfrak{f}$] If $S_{\mathfrak{f},\ell,3}/\langle\tau\rangle \neq \emptyset$, or if there exists $\mathfrak{p} \in S_{\mathfrak{f},\ell,2}/\langle\tau\rangle$ such that $v_\mathfrak{p}(\mathfrak{f}) \equiv 1 \pmod{\ell}$, or if there exists $\mathfrak{p} \in S_\mathfrak{f}/\langle\tau\rangle$ such that $v_\mathfrak{p}(\mathfrak{f}) \ge 2$, there is an error (the extension L/K must exist), so terminate the algorithm (see Remark (1) below).

12. [Compute $\beta_\mathfrak{p}$ and $p_{i,\mathfrak{p}}$] Using Algorithm 5.2.10, for each $\mathfrak{p} \in S/\langle\tau\rangle = S_\mathfrak{f}/\langle\tau\rangle \cup S_{\mathfrak{f},\ell,1}/\langle\tau\rangle$, compute $\beta \in K_z^*$ and integers $p_{i,\mathfrak{p}}$ such that for some ideal $\mathfrak{q}$ we have $\mathfrak{p} = \beta\mathfrak{q}^\ell \prod_{1\le i\le r_c} \mathfrak{a}_i^{p_{i,\mathfrak{p}}}$. Then set $\beta_\mathfrak{p} \leftarrow \beta/\prod_{1\le i\le r_c} \gamma_i^{p_{i,\mathfrak{p}}}$ and

$$\alpha_\mathfrak{p} \leftarrow \prod_{0\le a<d} (\tau^a(\beta_\mathfrak{p}))^{g^{d-1-a} \bmod \ell} .$$

13. [Introduce notation] (This is mainly a notational step.) For $1 \le j \le s$, let $\mathfrak{p}_j$ be the prime ideals in $S/\langle\tau\rangle$, set $w_{j+d_v} \leftarrow \beta_{\mathfrak{p}_j}$ for $1 \le j \le s$, and set $d_w \leftarrow s + d_v$ (this will be the number of columns of a matrix that we will construct). On the other hand, let $(\mathfrak{m}_i)_{1\le i\le t}$ be the following moduli (in any order): $\mathfrak{p}^{z(\mathfrak{p},\ell)-v_\mathfrak{p}(\mathfrak{f})}$ for $\mathfrak{p} \in S_{\mathfrak{f},\ell,2}/\langle\tau\rangle$ and $\mathfrak{p}^{z(\mathfrak{p},\ell)-1}$ for $\mathfrak{p} \in S_\ell/\langle\tau\rangle$.

14. [Compute discrete logarithms] Using Algorithms 4.2.17 and 4.2.18, compute the SNF of $(\mathbb{Z}_{K_z}/\mathfrak{m}_i)^*$ as well as $L_{\mathfrak{m}_i}(w_j)$ and $L_{\mathfrak{m}_i}(\alpha_{\mathfrak{p}_{j'}})$ for all i such that $1 \le i \le t$, for all j such that $1 \le j \le d_v$, and for all j' such that $1 \le j' \le s$.

15. [Create big matrix] Construct a matrix M with d_w columns as follows. Let M_j be the jth column of M. Let g^{-1} be an inverse of g modulo ℓ. If $1 \le j \le d_v$, M_j is obtained by concatenating the $(dg^{-1})L_{\mathfrak{m}_i}(w_j)$ for $1 \le i \le t$ (considered as column vectors), together with the zero vector with d_c components. If $d_v < j \le d_w$, M_j is obtained by concatenating the $L_{\mathfrak{m}_i}(\alpha_{\mathfrak{p}_{j-d_v}})$ for $1 \le i \le t$ together with the d_c-component column vector $Q^t P_{j-d_v}$, where P_{j-d_v} denotes the r_c-component column vector of the $p_{i,\mathfrak{p}_{j-d_v}}$ for $1 \le i \le r_c$. Finally, denote by $\overline{M}$ the matrix M reduced modulo ℓ, considered as a matrix with entries in $\mathbb{F}_\ell$.

16. [Compute kernel] Using [Coh0, Algorithm 2.3.1], compute the kernel $\mathcal{K}$ of the matrix $\overline{M}$ as an $\mathbb{F}_\ell$-vector space. If this kernel is reduced to $\{0\}$, there is an error, so terminate the algorithm. Otherwise, let $d_{\mathcal{K}} \leftarrow \dim(\mathcal{K})$ be the dimension of this kernel, and denote by $(K_j)_{1 \le j \le d_{\mathcal{K}}}$ a $\mathbb{F}_\ell$-basis of $\mathcal{K}$, where the K_j are considered as d_w-component column vectors. Finally, set $c \leftarrow d_{\mathcal{K}}$.

17. [Compute more discrete logarithms] Let $(\mathfrak{m}'_i)_{1 \le i \le t'}$ be the moduli $\mathfrak{p}^{z_\mathfrak{p}}$ with $z_\mathfrak{p} = z(\mathfrak{p}, \ell) - v_\mathfrak{p}(\mathfrak{f}) + 1$, for all $\mathfrak{p} \in S_{\mathfrak{f},\ell,2}/\langle\tau\rangle$. As in step 14, compute the SNF of $(\mathbb{Z}_{K_z}/\mathfrak{m}'_i)^*$ as well as $L_{\mathfrak{m}'_i}(w_j)$ and $L_{\mathfrak{m}'_i}(\alpha_{\mathfrak{p}_{j'}})$ for $1 \le i \le t'$, $1 \le j \le d_v$, and $1 \le j' \le s$. For $1 \le i \le t'$, let M'_i be the matrix with d_w columns, each column containing $(dg^{-1})L_{\mathfrak{m}'_i}(w_j)$ for $1 \le j \le d_v$, or $L_{\mathfrak{m}'_i}(\alpha_{\mathfrak{p}_{j-d_v}})$ if $d_v < j \le d_w$. Do *not* put the matrices M'_i together by rows as above.

18. [Initialize backtracking] (In what follows, $c \ge 1$ and y will be a row vector with $c-1$ components.) Set $y \leftarrow (0, \ldots, 0)$ (vector with $c-1$ components).

19. [Compute trial vector] Let $X \leftarrow K_c + \sum_{1 \le j < c} y_j K_j$. Apply Subalgorithm 5.3.18 below to see if X corresponds to a suitable Abelian extension L/K. If it does, output the defining polynomial given by the subalgorithm and terminate the algorithm.

20. [Backtracking I] Set $i \leftarrow c$.

21. [Backtracking II] Set $i \leftarrow i-1$. If $i > 0$, go to step 22. Otherwise, set $c \leftarrow c-1$. If $c > 0$, go to step 18; otherwise, there is an error, so terminate the algorithm.

22. [Backtracking III] Set $y_i \leftarrow y_i + 1$, and if $i < c-1$, set $y_{i+1} \leftarrow 0$. If $y_i \ge \ell$, go to step 21; otherwise, go to step 19.

Subalgorithm 5.3.18 (Is X Suitable?). Given a vector $X = (x_1, \ldots, x_{d_w})^t$ found in step 19 of Algorithm 5.3.17, this subalgorithm determines whether X corresponds to a suitable Abelian extension L/K. If it does, it computes a defining polynomial for L/K. We use all the quantities computed in the main algorithm.

1. [Test conditions on $x_{\mathfrak{p}}$] If X is equal to zero modulo ℓ or if any of the x_i for $d_v < i \leq d_w$ is equal to zero modulo ℓ, X is not suitable, so terminate the subalgorithm.
2. [Test $\mathfrak{m}'_i$] For $1 \leq i \leq t'$, compute $Y_i \leftarrow M'_i X$. If for any i, $Y_i \equiv 0 \pmod{\ell}$, X is not suitable, so terminate the subalgorithm.
3. [Compute defining polynomial] (Here, for all $i \leq t'$ we have $Y_i \not\equiv 0 \pmod{\ell}$.) Compute $\beta \leftarrow \prod_{1 \leq i \leq d_w} w_i^{x_i}$, then try to reduce β as much as possible by multiplying by ℓth powers of elements. Finally, using Theorem 5.3.5 and Proposition 5.3.9, compute the defining polynomial P_β of the number field L_β corresponding to β.
4. [Compute norm group] Using Algorithm 4.4.3, compute the norm group $T_{\mathfrak{m}}(L_\beta/K)$ corresponding to the given initial modulus $\mathfrak{m}$. If this is not equal to the initial congruence subgroup C, X is not suitable; otherwise, output the defining polynomial found in step 3 and terminate the subalgorithm.

Remarks

(1) Step 11 is included for completeness (and takes no time), but it is not necessary since by Takagi's theorem we know that the extension L/K exists, hence the conductor $\mathfrak{f}$ of L_z/K_z will satisfy the conditions of step 11.

(2) In Algorithm 5.2.14 we looked for all extensions of degree ℓ corresponding to a given modulus, and we could find none, one, or several. In the present algorithm, we look at a specific congruence subgroup $(\mathfrak{m}, C)$ such that $h_{\mathfrak{m},C} = \ell$; hence by Takagi's existence theorem, we know that, up to isomorphism, there exists one and only one suitable Abelian extension L/K. Thus, if in steps 11, 16, or 21 the given conditions are not satisfied, this means that there is an error somewhere, either in the author's write-up of the above (admittedly extremely technical) algorithm or in the implementation.

(3) Steps 1 to 6 depend only on K and ℓ and not on the congruence subgroup $(\mathfrak{m}, C)$, and so should be done once and for all if several congruence subgroups are considered with the same base field K and degree ℓ.

(4) Since we explicitly assume that $\zeta_\ell \notin K$, the case $\ell = 2$ cannot occur; hence we do not have to deal with Archimedean conditions at the level of K_z which is totally complex.

(5) The case $\alpha = 1$ (or more generally $\alpha \in K_z^{*\ell}$) can only occur if no congruence conditions are tested, hence only if $S = S_{\mathfrak{f},\ell,2} = \varnothing$, which is equivalent to $\mathfrak{f} = \mathbb{Z}_{K_z}$.

(6) It is absolutely essential to try to reduce β in step 2 of Subalgorithm 5.3.18, otherwise the size of the coefficients of the defining polynomial will be too large. We are, of course, allowed to multiply β by any ℓth power of an element of K_z without changing the field L. Hence, one method is to multiply by ℓth power of units (which have been computed

anyway in step 3 of Algorithm 5.3.17) and do this recursively until it is no longer possible to reduce β. Of course, one needs a way to measure the "size" of β, but in the present problem any naive measure such as the T_2 norm or the L^2 norm of the coefficients of β in an integral basis of K_z is sufficient.

(7) In step 8 of the algorithm, we compute the conductor $\mathfrak{f}$ of the extension L_z/K_z. This is essential in order to be able to apply the necessary and sufficient conditions given by Theorem 5.3.15. Computing $\mathfrak{f}$ is trivial when $\mathfrak{m}$ is coprime to ℓ. On the other hand, if $\mathfrak{m}$ is not coprime to ℓ, the conductor may be a strict divisor of $\mathfrak{m}\mathbb{Z}_{K_z}$, although its prime to ℓ-part is the same as that of $\mathfrak{m}\mathbb{Z}_{K_z}$. Thus, there are two solutions to this problem. The first one, which we have chosen, is to compute the conductor of the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K_z}, \mathcal{N}_{K_z/K}^{-1}(C))$. However, this involves computing ray class groups in the large field K_z, and could be extremely costly. Another possible solution, which should seriously be considered, is to lose the necessity of the conditions of Theorem 5.3.15 and reformulate the theorem for a modulus $\mathfrak{m}$ that is a multiple of the conductor but not necessarily equal to it.

(8) In the case $\ell = 3$ (which is not difficult anyway), there is quite a different method for computing the defining polynomial of L/K from that of L_z/K_z instead of using Lagrange resolvents, Theorem 5.3.5, and Proposition 5.3.9. We can simply look for some α satisfying the ramification conditions but not necessarily the Galois conditions, and then use resultants (see step 10 of Algorithm 9.2.7).

Once we have the desired defining polynomial, we want to reduce it as much as possible and also to reduce the corresponding absolute polynomial. Some reduction has been done during the algorithm, but usually this is far from sufficient. As explained in Chapter 2, to go further, there are several possibilities. All of them are absolute or relative variants of polynomial reduction algorithms and hence involve computing an integral pseudo-basis of $\mathbb{Z}_L$. If done rashly, this may involve factoring large discriminants. The best method is certainly as follows. We first compute an integral pseudo-basis of $\mathbb{Z}_{L_z}/\mathbb{Z}_{K_z}$, which is easy since it is a Kummer extension. Then using the integral pseudo-basis of $\mathbb{Z}_{K_z}/\mathbb{Z}_K$ we can easily compute an integral pseudo-basis of $\mathbb{Z}_{L_z}/\mathbb{Z}_K$, and finally using Exercise 35 of Chapter 2, we can compute an integral pseudo-basis of $\mathbb{Z}_L/\mathbb{Z}_K$. After that, we can either use a relative version of the polynomial reduction algorithm (see Section 2.4.2), which improves the defining polynomial somewhat, or use an algorithm such as Algorithm 2.4.12 combined with the explicit knowledge of the integral pseudo-basis of $\mathbb{Z}_L/\mathbb{Z}_K$, and this usually gives good results.

5.4 Explicit Use of the Artin Map in Kummer Theory When $\zeta_n \in K$

The contents of this section as well as the next one are for the most part inspired by work of C. Fieker (see [Fie]), whom I gratefully thank for detailed comments.

As mentioned in Section 5.1.2, there are two other methods for constructing the desired Kummer extension L/K which use the explicit computation of the Artin map. As we shall see, these methods are preferable to the methods using Hecke's theorem since they do not require extensions of prime degree and since the corresponding algorithms are simpler.

We recall that, using Algorithm 5.1.2, we are reduced to the following situation. We have a base field K, a congruence subgroup $(\mathfrak{m}, C)$ of K of conductor $\mathfrak{m}$ such that $Cl_{\mathfrak{m}}(K)/\overline{C}$ is a cyclic group of order $n = \ell^r$ for some prime number ℓ. The fact that $\mathfrak{m}$ is the conductor is not essential for Fieker's method, but it speeds up the algorithm. As with Hecke's theorem, there are two distinct stages. In the first stage, we assume that $\zeta_n \in K$, so that we can use Kummer theory. In the second stage, we will as before adjoin ζ_n to K, determine the corresponding Kummer extension L_z/K_z, and then explain the method for coming down to the desired extension L/K corresponding to $(\mathfrak{m}, C)$ under the Takagi map.

5.4.1 Action of the Artin Map on Kummer Extensions

In this section, we consider a Kummer extension N/K which is assumed to be cyclic of exponent dividing n, and we assume known a suitable modulus $\mathfrak{m}_N$ in the sense of Definition 3.4.2.

Recall that the Artin map is the unique map from the ideals of K coprime to $\mathfrak{m}_N$ to the Galois group $\mathrm{Gal}(N/K)$, defined by multiplicativity on ideals, and such that for an unramified prime ideal $\mathfrak{p}$ we have $\mathrm{Art}(\mathfrak{p}) = \sigma_{\mathfrak{p}}$, where $\sigma_{\mathfrak{p}}$ is the Frobenius automorphism at $\mathfrak{p}$ characterized by $\sigma_{\mathfrak{p}}(x) \equiv x^{\mathcal{N}(\mathfrak{p})} \pmod{\mathfrak{p}\mathbb{Z}_N}$ for all $x \in \mathbb{Z}_N$. Since $\mathfrak{m}_N$ is suitable, the Artin map is surjective, and its kernel is equal to the congruence subgroup D modulo $\mathfrak{m}_N$ defining the extension N/K through the Takagi correspondence.

We want to determine the Artin map explicitly. More precisely, since $\zeta_n \in K$, we know by Kummer theory that $N = K(\theta)$ for some θ such that $\theta^n = \alpha \in \mathbb{Z}_K$. If $\mathfrak{a}$ is an ideal of K coprime to $\mathfrak{m}_N$, then $\mathrm{Art}(\mathfrak{a}) \in \mathrm{Gal}(N/K)$ is a K-automorphism of N, which we want to determine explicitly. Since the action of an automorphism is entirely determined by the image of θ, we must simply compute $\mathrm{Art}(\mathfrak{a})(\theta)$. By multiplicativity, it is sufficient to determine $\sigma_{\mathfrak{p}}(\theta)$ for a Frobenius automorphism $\sigma_{\mathfrak{p}}$ associated to an unramified prime ideal $\mathfrak{p}$. This is given by the following proposition, which we state for general cyclic Kummer extensions, not necessarily of prime power degree.

Proposition 5.4.1. *Let $n > 1$ be an arbitrary integer (not necessarily a prime power), let K be a number field such that $\zeta_n \in K$, let $N = K(\theta)$ with $\theta^n = \alpha \in \mathbb{Z}_K$ of degree n, and let $\mathfrak{p}$ be a prime ideal of K not dividing α and n (in particular, unramified in N/K). Finally, let g be a generator of the multiplicative group $(\mathbb{Z}_K/\mathfrak{p})^*$.*

(1) *There exists an integer $q \geq 1$ such that $\mathcal{N}(\mathfrak{p}) = qn + 1$.*
(2) *In the group $(\mathbb{Z}_K/\mathfrak{p})^*$, there exist integers y and z with $(z, n) = 1$ such that $\overline{\zeta_n} = g^{qz}$ and $\overline{\alpha} = g^y$.*
(3) *We have*

$$\sigma_{\mathfrak{p}}(\theta) = \zeta_n^{yz^{-1}}\theta \ ,$$

where z^{-1} denotes an inverse of z modulo n.

Proof. Since $\sigma_{\mathfrak{p}}(\theta)$ is a conjugate of θ, we must have $\sigma_{\mathfrak{p}}(\theta) = \zeta_n^{s_{\mathfrak{p}}}\theta$ for some integer $s_{\mathfrak{p}}$.

Let $\mathcal{N}(\mathfrak{p}) = qn + t$ with $0 \leq t < n$ be the Euclidean division of $\mathcal{N}(\mathfrak{p})$ by n. We obtain

$$\sigma_{\mathfrak{p}}(\theta) \equiv \theta^{\mathcal{N}(\mathfrak{p})} = (\theta^n)^q\theta^t = \alpha^q\theta^t \pmod{\mathfrak{p}\mathbb{Z}_N} \ ,$$

hence

$$\zeta_n^{s_{\mathfrak{p}}}\theta \equiv \alpha^q\theta^t \pmod{\mathfrak{p}\mathbb{Z}_N} \ .$$

Now

$$\operatorname{disc}(1, \theta, \ldots, \theta^{n-1}) = \operatorname{disc}(X^n - \alpha) = (-1)^{(n-1)(n-2)/2}n^n\alpha^{n-1} \ ,$$

and since we have assumed that $\mathfrak{p} \nmid \alpha$ and $\mathfrak{p} \nmid n$, we have $\operatorname{disc}(1, \theta, \ldots, \theta^{n-1}) \not\equiv 0 \pmod{\mathfrak{p}\mathbb{Z}_N}$, so the classes of $1, \theta, \ldots, \theta^{n-1}$ are $\mathbb{Z}_K/\mathfrak{p}$-linearly independent in $\mathbb{Z}_N/\mathfrak{p}\mathbb{Z}_N$. It follows that $t = 1$ and that $\zeta_n^{s_{\mathfrak{p}}} \equiv \alpha^q \pmod{\mathfrak{p}\mathbb{Z}_N}$. Since ζ_n and α belong to $\mathbb{Z}_K$ and $\mathfrak{p}\mathbb{Z}_N \cap \mathbb{Z}_K = \mathfrak{p}$, we therefore have the congruence

$$\zeta_n^{s_{\mathfrak{p}}} \equiv \alpha^q \pmod{\mathfrak{p}}$$

in $\mathbb{Z}_K$.

Since $t = 1$, the group $(\mathbb{Z}_K/\mathfrak{p})^*$ is of order qn. If g is a generator, we have $\overline{\alpha} = g^y$ for some y since $\mathfrak{p} \nmid \alpha$. On the other hand, $\overline{\zeta_n}$ is of order dividing n, hence we have $\overline{\zeta_n} = g^{qz}$ for some integer z. I claim that $\overline{\zeta_n}$ is of order exactly n. Indeed, the discriminant of the set of elements ζ_n^a for $(a, n) = 1$ is equal to the discriminant of the cyclotomic polynomial $\Phi_n(X)$, which is only divisible by primes dividing n (see Exercise 9) hence not divisible by $\mathfrak{p}$ by assumption, proving my claim. Thus we have $(z, n) = 1$, and the congruence for $\zeta_n^{s_{\mathfrak{p}}}$ that we have obtained gives

$$g^{qzs_{\mathfrak{p}}} = g^{qy} \ ,$$

and since g is of order exactly qn it follows that $qzs_{\mathfrak{p}} \equiv qy \pmod{qn}$, hence $s_{\mathfrak{p}} \equiv yz^{-1} \pmod{n}$, finishing the proof of the proposition. □

Note that if $\mathfrak{p}$ is an unramified prime ideal in the extension N/K, it is always possible to choose α such that $\mathfrak{p} \nmid \alpha$, so the (essential) restriction of the proposition is not restrictive in practice (see Exercise 10).

We can therefore give the following algorithm for the explicit action of the Artin map on Kummer extensions.

Algorithm 5.4.2 (Action of the Artin Map on Kummer Extensions). Let N/K be a cyclic Kummer extension of degree n, where $\zeta_n \in K$ and N given as $N = K(\theta)$ with $\theta^n = \alpha \in \mathbb{Z}_K$. Let $\mathfrak{a}$ be an ideal of K coprime to α and n. This algorithm computes the element $\mathrm{Art}(\mathfrak{a})(\theta)$ by means of an element e defined modulo n such that $\mathrm{Art}_{N/K}(\mathfrak{a})(\theta) = \zeta_n^e\theta$ (so that $\mathrm{Art}_{N/K}(\mathfrak{a})(P(\theta)) = P(\zeta_n^e\theta)$ if $P \in K[X]$).

1. [Factor $\mathfrak{a}$] Factor $\mathfrak{a}$ into a product of prime ideals as $\mathfrak{a} = \prod_{1\le i\le k} \mathfrak{p}_i^{v_i}$, and set $j \leftarrow 0$, $e \leftarrow 0$.
2. [Compute $(\mathbb{Z}_K/\mathfrak{p})^*$] Set $j \leftarrow j+1$. If $j > k$, output e mod n and terminate the algorithm. Otherwise, set $\mathfrak{p} \leftarrow \mathfrak{p}_j$, compute $\mathcal{N}(\mathfrak{p})$, a generator g of $(\mathbb{Z}_K/\mathfrak{p})^*$, and set $q \leftarrow (\mathcal{N}(\mathfrak{p}) - 1)/n$.
3. [Compute discrete logarithms] Use a discrete logarithm algorithm in $(\mathbb{Z}_K/\mathfrak{p})^*$ to compute y and z_1 such that $\overline{\alpha} = g^y$ and $\overline{\zeta_n} = g^{z_1}$, and set $z \leftarrow z_1/q$ (this must be an integer coprime to n).
4. [Compute $s_{\mathfrak{p}}$] Set $s \leftarrow yz^{-1} \pmod n$ and $e \leftarrow e + s$, and go to step 2.

Remark. This algorithm may be slow for two reasons. The first is that the computation of discrete logarithms in $(\mathbb{Z}_K/\mathfrak{p})^*$ may take time if the norm of $\mathfrak{p}$ is large, so we must avoid this if possible. The second, perhaps more subtle, reason is that the factorization of the ideal $\mathfrak{a}$ may be very costly, if not impossible. Thus we must also try to avoid this, and we shall see that in practice this is no problem.

5.4.2 Reduction to $\alpha \in U_S(K)/U_S(K)^n$ for a Suitable S

Keeping all our notation, our goal is to find $\alpha \in \mathbb{Z}_K$ such that $L = K(\theta)$ with $\theta^n = \alpha$ is the class field corresponding to the congruence subgroup $(\mathfrak{m}, C)$ of conductor $\mathfrak{m}$. We begin with the following lemma, which is a generalization of the easy part of Hecke's Theorem 10.2.9.

Lemma 5.4.3. *Let $L = K(\theta)$ with $\theta^n = \alpha \in K$ and $n = \ell^r$ be a Kummer extension as above. If $\mathfrak{p}$ is a prime ideal of K that satisfies $\ell^r \nmid v_{\mathfrak{p}}(\alpha)$, then $\mathfrak{p}$ is ramified in L/K. In other words, if $\mathfrak{p}$ is unramified, then $\ell^r \mid v_{\mathfrak{p}}(\alpha)$.*

Proof. Multiplying α by a suitable ℓ^rth power, we may assume that $\alpha \in \mathbb{Z}_K$. Indeed, this does not change the field L and does not change the condition $\ell^r \nmid v_{\mathfrak{p}}(\alpha)$.

Let $v_{\mathfrak{p}}(\alpha) = \ell^a w$ with $\ell \nmid w$. By assumption, we have $a < r$, and we can find nonnegative integers x and y such that $-x\ell^{r-a} + yw = 1$. If $\pi \in \mathfrak{p}^{-1} \setminus \mathfrak{p}^{-2}$,

then $\beta = \pi^{x\ell^r}\alpha^y$ satisfies $v_{\mathfrak{p}}(\beta) = -x\ell^r + yw\ell^a = \ell^a$, it is such that $\beta \in \mathbb{Z}_K$, and by Corollary 10.2.7, its nth root defines the same field as θ since $\ell \nmid y$. Thus, we may assume that $\alpha \in \mathbb{Z}_K$ and that $w = 1$, so that $v_{\mathfrak{p}}(\alpha) = \ell^a$ for $a < r$.

The ideal $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \theta\mathbb{Z}_L$ is not necessarily a prime ideal, but since $\alpha \in \mathbb{Z}_K$ it is an integral ideal of $\mathbb{Z}_L$. Using Proposition 2.3.15 we know that

$$\mathfrak{P}^{\ell^r} = \mathfrak{p}^{\ell^r}\mathbb{Z}_L + \alpha\mathbb{Z}_L = (\mathfrak{p}^{\ell^r} + \alpha\mathbb{Z}_K)\mathbb{Z}_L \ .$$

Since $\mathfrak{p}$ is a prime ideal of K, $v_{\mathfrak{q}}(\mathfrak{p}^{\ell^r} + \alpha\mathbb{Z}_K) = 0$ if $\mathfrak{q} \neq \mathfrak{p}$ while

$$v_{\mathfrak{p}}(\mathfrak{p}^{\ell^r} + \alpha\mathbb{Z}_K) = \min(\ell^r, \ell^a) = \ell^a \ ,$$

hence $\mathfrak{p}^{\ell^r} + \alpha\mathbb{Z}_K = \mathfrak{p}^{\ell^a}$ so that $\mathfrak{P}^{\ell^r} = (\mathfrak{p}\mathbb{Z}_L)^{\ell^a}$. It follows that $\mathfrak{P}^{\ell^{r-a}} = \mathfrak{p}\mathbb{Z}_L$, and since $a < r$ this shows that $\mathfrak{p}$ is ramified in L/K with ramification exponent equal to ℓ^b for some b such that $r - a \leq b \leq r$, proving the lemma. □

The following proposition is a generalization of part of Theorem 5.2.9.

Proposition 5.4.4. *Let $Cl(K) = \bigoplus_i (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$ be the SNF of the ordinary class group of K. Let S be the set of prime ideals of K dividing $\mathfrak{m}$ and the $\mathfrak{a}_i$. We may choose $\alpha \in \mathbb{Z}_K$ such that $L = K(\theta)$ with $\theta^n = \alpha$ with α an S-unit; in other words, α divisible only by prime ideals of S (see Definition 7.4.1).*

Proof. Since L/K is a cyclic Kummer extension of degree n, we know that there exists $\alpha \in K$ such that $L = K(\theta)$ with $\theta^n = \alpha$. We do not assume for the moment that $\alpha \in \mathbb{Z}_K$. We are going to modify α so that it satisfies the required properties. We prove the result by induction on the number k of prime ideals occurring in the prime decomposition of α and not belonging to S. If $k = 0$, we are done. Otherwise let $k \geq 1$, assume the proposition proved up to $k - 1$, let α have exactly k prime ideals not in S in its prime ideal decomposition, and let $\mathfrak{p}$ be such an ideal. Then $\mathfrak{p} \nmid \mathfrak{m}$ so $\mathfrak{p}$ is unramified, hence by the above lemma we know that $n \mid v_{\mathfrak{p}}(\alpha)$. Let $\mathfrak{a} = \mathfrak{p}^{v_{\mathfrak{p}}(\alpha)/n}$. By definition of the $\mathfrak{a}_i$ we can write $\mathfrak{a} = \beta\prod_i \mathfrak{a}_i^{x_i}$ for some $\beta \in K$. Thus,

$$\alpha\beta^{-n}\mathbb{Z}_K = \alpha\mathfrak{p}^{-v_{\mathfrak{p}}(\alpha)}\prod_i \mathfrak{a}_i^{nx_i} \ ,$$

so $\gamma = \alpha\beta^{-n}$ also defines the extension L/K and has exactly $k-1$ prime ideals not belonging to S in its prime ideal decomposition, proving our induction hypothesis.

We have thus shown that $\alpha \in K$ exists with the desired properties. Finally, let us prove that we can choose such an α belonging to $\mathbb{Z}_K$. Indeed, we can write in a unique way $\alpha\mathbb{Z}_K = \mathfrak{a}/\mathfrak{b}$ with $\mathfrak{a}$ and $\mathfrak{b}$ coprime integral ideals. Since α is an S-unit, $\mathfrak{a}$ and $\mathfrak{b}$ are divisible only by prime ideals of S. If $h(K) = |Cl(K)|$ is the class number of K, then $\mathfrak{b}^{h(K)} = \gamma\mathbb{Z}_K$ is a principal

ideal. It is clear that $\alpha\gamma^n$ also defines L/K, is still an S-unit, and is in $\mathbb{Z}_K$, proving the proposition. □

Thus, to find α we need to search only among S-units. I refer to Chapter 7 for definitions, results, and algorithms on S-units. In particular, we denote by $U_S(K)$ the multiplicative group of S-units, and we recall that $U_S(K)$ is a finitely generated Abelian group of rank equal to $r_1 + r_2 - 1 + |S|$ with cyclic torsion subgroup equal to $\mu(K)$. Because of the freedom given to us by Kummer theory (specifically Corollary 10.2.7), we may assume that α belongs to a fixed system of representatives of $U_S(K)/U_S(K)^n$. Since $\zeta_n \in K$, $w(K) = |\mu(K)|$ is divisible by n; hence it follows that $U_S(K)/U_S(K)^n$ is a free $\mathbb{Z}/n\mathbb{Z}$-module of rank $r_1 + r_2 + |S|$. As in the method using Hecke's theorem, we thus have only a finite number of possibilities for α (at most $n^{r_1+r_2+|S|}$). Obviously, we will not systematically explore all these possibilities. To the contrary, however, we will see that it is not too difficult to reduce the search to a much smaller number of cases, as we have done using Hecke's theorem. It is clear, however, that we should try to have $|S|$ as small as possible. Let us see how to achieve this.

First, $|S|$ must contain the primes dividing $\mathfrak{m}$. Thus, to have a minimal number of such primes, it is often a good idea to replace $\mathfrak{m}$ by the conductor of the congruence subgroup $(\mathfrak{m}, C)$. Once this is done, the prime ideals dividing $\mathfrak{m}$ are exactly the primes that ramify in the extension L/K, and from the proof of the proposition we see that these are necessary, so we cannot hope for a smaller set.

In addition, S must contain the prime ideals dividing the $\mathfrak{a}_i$, whose classes are the generators of the ordinary class group $Cl(K)$ of K. Here we do not have any freedom in the number of such ideals, but we do have great freedom in the choice of the generators. Indeed, the subgroup $(\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$ of $Cl(K)$ is also equal to $(\mathbb{Z}/d_i\mathbb{Z})\overline{\beta\mathfrak{a}_i^j}$ for any $\beta \in K$ and any j coprime to d_i, so we may replace $\mathfrak{a}_i$ by any ideal of the form $\beta\mathfrak{a}_i^j$ with $(j, d_i) = 1$. More generally, we may replace the *set* of $\mathfrak{a}_i$ by any set of ideals that generate the class group, since for our purposes we do not really need the Smith normal form.

In addition, since we already have the primes dividing $\mathfrak{m}$ in S, we need only to take a set of ideals whose classes generate $Cl(K)/ < S_{\mathfrak{m}} >$, where $S_{\mathfrak{m}}$ is the set of (all) prime ideals dividing $\mathfrak{m}$.

According to a weak form of Tchebotarev's density theorem (see, for example, [Lan3]), in any ideal class there exists an infinite number of prime ideals, hence if desired we may choose the $\mathfrak{a}_i$ to be prime ideals, but this is not always a good idea.

5.4.3 Construction of the Extension L/K by Kummer Theory

Let $u = r_1 + r_2 + |S|$, and let $(\varepsilon_1, \ldots, \varepsilon_u)$ be S-units whose classes modulo $U_S(K)^n$ form a $\mathbb{Z}/n\mathbb{Z}$-basis of $U_S(K)/U_S(K)^n$. According to Proposition

5.4.4, the desired α can be taken of the form

$$\alpha = \prod_{1 \le j \le u} \varepsilon_j^{x_j}$$

with $0 \le x_j < n$.

Consider the large field $N = K(\theta_1, \ldots, \theta_u)$, where $\theta_j^n = \varepsilon_j$, as well as the subfields $N_j = K(\theta_j)$. These fields are Kummer extensions of K. We will not need to work algorithmically in these fields, but it is necessary to introduce them in order to justify the final algorithm. We note that N is an Abelian extension of K equal to the compositum of the N_j, and by Proposition 5.4.4 the desired field extension L/K is a subextension of N/K. We will determine L/K explicitly as a subextension of N/K.

Since the classes of the ε_j form a $\mathbb{Z}/n\mathbb{Z}$-basis of $U_S(K)/U_S(K)^n$, it follows from Kummer theory (Theorem 10.2.5) that $\mathrm{Gal}(N/K) \simeq (\mathbb{Z}/n\mathbb{Z})^u$ and that $\mathrm{Gal}(N_j/K) \simeq \mathbb{Z}/n\mathbb{Z}$. More precisely, we can give an element $\sigma \in \mathrm{Gal}(N/K)$ by specifying the images of the θ_j by σ, hence by setting

$$\sigma(\theta_j) = \zeta_n^{s_{\sigma,j}} \theta_j \ ,$$

and all the possible σ correspond to all the possible choices of $s_{\sigma,j}$ modulo n.

Proposition 5.4.5. *Let $\mathfrak{a}$ be an integral ideal of K coprime to the prime ideals of S, and let $\sigma_{\mathfrak{a}} = \mathrm{Art}_{N/K}(\mathfrak{a}) \in \mathrm{Gal}(N/K)$ be the automorphism corresponding to $\mathfrak{a}$ by the Artin map in $\mathrm{Gal}(N/K)$. Then $\sigma_{\mathfrak{a}}$ is determined by $\sigma_{\mathfrak{a}}(\theta_j) = \zeta_n^{s_{\mathfrak{a},j}} \theta_j$ with*

$$s_{\mathfrak{a},j} = \sum_{\mathfrak{p}} y_{\mathfrak{p},j} v_{\mathfrak{p}}(\mathfrak{a}) z_{\mathfrak{p}}^{-1} \ ,$$

where if for each prime ideal $\mathfrak{p}$ we denote by $g_{\mathfrak{p}}$ a generator of $(\mathbb{Z}_K/\mathfrak{p})^*$, we set in this last group

$$\overline{\zeta_n} = g_{\mathfrak{p}}^{z_{\mathfrak{p}}(\mathcal{N}(\mathfrak{p})-1)/n}$$

and

$$\overline{\theta_j} = g_{\mathfrak{p}}^{y_{\mathfrak{p},j}} \ .$$

Proof. By Proposition 3.5.6, the computation of $\sigma_{\mathfrak{a}}(\theta_j)$ can be done in the subextension N_j/K. Proposition 5.4.1 then tells us that $\sigma_{\mathfrak{a}}(\theta_j) = \zeta_n^{s_{\mathfrak{a},j}} \theta_j$ with

$$s_{\mathfrak{a},j} = \sum_{\mathfrak{p}} y_{\mathfrak{p},j} v_{\mathfrak{p}}(\mathfrak{a}) z_{\mathfrak{p}}^{-1} \ ,$$

where $y_{\mathfrak{p},j}$ and $z_{\mathfrak{p}}$ are as in the proposition. □

To be able to use class field theory, we must determine a modulus $\mathfrak{m}_N$ that is suitable for the extension N/K in the sense of Definition 3.4.2. Although it would be possible to determine the exact conductor of N/K, it is much easier to give such a modulus with essentially no computation, by using the

following theorem due to H. Hasse (see [Has], p. 232), which we state for a general Kummer extension, not necessarily a cyclic extension of prime power degree (see Exercise 6 of Chapter 10 for the proof of a slightly weaker result).

Theorem 5.4.6. *Let N/K be a Kummer extension of exponent n. Denote by R the set of prime ideals of K that are ramified in N/K and do not divide n, together with the infinite places of K ramified in N/K. Then the conductor of N/K divides*

$$\mathfrak{m}_N = \prod_{\mathfrak{p}\in R} \mathfrak{p} \prod_{\mathfrak{p}\mid n} \mathfrak{p}^{c_\mathfrak{p}} \ ,$$

with

$$c_\mathfrak{p} = \left(v_\ell(n) + \frac{1}{\ell - 1} \right) e(\mathfrak{p}/\ell) + 1 \ ,$$

where ℓ is the prime number below $\mathfrak{p}$.

The following lemma explicitly gives our desired Kummer extension L/K as a subextension of N/K.

Lemma 5.4.7. *Let $\overline{D}$ be the kernel of the canonical surjection s from $Cl_{\mathfrak{m}_N}(K)$ to $Cl_\mathfrak{m}(K)/\overline{C}$. Then*

$$Cl_{\mathfrak{m}_N}(K)^n \subset \mathrm{Ker}(\mathrm{Art}_{N/K}) \subset \overline{D} \ ,$$

and if we set

$$\mathcal{H} = \mathrm{Art}_{N/K}\left(\overline{D}\right) = \mathrm{Art}_{N/K}\left(\overline{D}/Cl_{\mathfrak{m}_N}(K)^n\right) \ ,$$

then $L = N^\mathcal{H}$.

Proof. Note that by abuse of notation, we denote also by $\mathrm{Art}_{N/K}$ the Artin map at the level of ideals or of ideal classes.

The Artin map $\mathrm{Art}_{N/K}$ is a surjective homomorphism from the ideals of K coprime to $\mathfrak{m}_N$ onto $\mathrm{Gal}(N/K)$ with kernel containing $P_{\mathfrak{m}_N}$, hence also from $Cl_{\mathfrak{m}_N}(K)$ onto $\mathrm{Gal}(N/K)$, which is isomorphic to $(\mathbb{Z}/n\mathbb{Z})^u$. Let us first show the two inclusions. By class field theory, we have

$$Cl_{\mathfrak{m}_N}(K)/\mathrm{Ker}(\mathrm{Art}_{N/K}) \simeq \mathrm{Gal}(N/K) \simeq (\mathbb{Z}/n\mathbb{Z})^u \ .$$

It follows that $Cl_{\mathfrak{m}_N}(K)/\mathrm{Ker}(\mathrm{Art}_{N/K})$ is of exponent n, hence $Cl_{\mathfrak{m}_N}(K)^n \subset \mathrm{Ker}(\mathrm{Art}_{N/K})$.

For the second inclusion, we note that since $\mathfrak{m}$ is the conductor of L/K we have $\mathfrak{m} \mid \mathfrak{m}_N$. Thus $\overline{D}$ is equal to the kernel of $\mathrm{Art}_{L/K}$ viewed as a map from $Cl_{\mathfrak{m}_N}(K)$ to $\mathrm{Gal}(L/K)$. By Proposition 3.5.6 it follows that

$$\mathrm{Ker}(\mathrm{Art}_{N/K}) \subset \mathrm{Ker}(\mathrm{Art}_{L/K}) = \overline{D} \ ,$$

proving the second inclusion. Furthermore, this also proves that

$$\mathrm{Gal}(N/L) = \mathrm{Art}_{N/K}\big(\overline{D}/\mathrm{Ker}(\mathrm{Art}_{N/K})\big) \ ,$$

so $L = N^{\mathcal{H}}$, where, with an evident abuse of notation, we can write indifferently

$$\mathcal{H} = \mathrm{Art}_{N/K}\big(\overline{D}/\mathrm{Ker}(\mathrm{Art}_{N/K})\big) = \mathrm{Art}_{N/K}\big(\overline{D}\big) = \mathrm{Art}_{N/K}\big(\overline{D}/Cl_{\mathfrak{m}_N}(K)^n\big) \ ,$$

finishing the proof of the lemma. □

Remarks

(1) For future reference, note that we do not use any special properties of N in this lemma, only the fact that N/K is an Abelian extension containing the extension L/K, and that $\mathfrak{m}_N$ is a suitable modulus.
(2) It is costly to determine $\mathrm{Ker}(\mathrm{Art}_{N/K})$, and this is why we prefer to define $\mathcal{H}$ simply as the image by $\mathrm{Art}_{N/K}$ of $\overline{D}/Cl_{\mathfrak{m}_N}(K)^n$, which is already a small group of exponent dividing n.

5.4.4 Picking the Correct α

The field L is now determined as the fixed field of the large field N by the subgroup $\mathcal{H}$ of $\mathrm{Gal}(N/K)$. Since the map f is explicitly known, its kernel $\overline{D}$ can be explicitly computed (using, for example, Algorithm 4.1.11), and computing $\overline{D}/Cl_{\mathfrak{m}_N}(K)^n$ is also a very simple matter. Thus, in principle the problem is solved. Giving the answer as $L = N^{\mathcal{H}}$ for a large field N and an explicit group $\mathcal{H}$ is, however, not satisfactory; we must really find a defining polynomial for L/K, in other words an $\alpha \in \mathbb{Z}_K$ such that $L = K(\theta)$ with $\theta^n = \alpha$.

Using the work done in the preceding sections, we know that we can write

$$\alpha = \prod_{1 \le j \le u} \varepsilon_j^{x_j} \ ,$$

where the classes of the ε_j form a $\mathbb{Z}/n\mathbb{Z}$-basis of $U_S(K)/U_S(K)^n$, and the x_j are defined modulo n. We must now find necessary and sufficient conditions on the x_j so that $L = K(\theta) = N^{\mathcal{H}}$ with the notation of Lemma 5.4.7. Since $\overline{D}/Cl_{\mathfrak{m}_N}(K)^n$ is of exponent dividing n, we can explicitly compute

$$\overline{D}/Cl_{\mathfrak{m}_N}(K)^n = \sum_{1 \le i \le d} (\mathbb{Z}/f_i\mathbb{Z})\overline{\mathfrak{f}_i} \ ,$$

where the f_i are divisors of n (of the form ℓ^k with $k \le r$ if $n = \ell^r$, ℓ prime, as we generally assume). Thus L is the fixed field of N by the $\mathrm{Art}_{N/K}(\mathfrak{f}_i)$ for $1 \le i \le d$.

Proposition 5.4.5 explicitly gives us the action of $\mathrm{Art}_{N/K}(\mathfrak{f}_i)$ on θ_j as

$$\mathrm{Art}_{N/K}(\mathfrak{f}_i)(\theta_j) = \zeta_n^{s_{i,j}}\theta_j$$

for an explicitly computed $s_{i,j}$. Thus, for

$$\theta = \prod_{1 \le j \le u} \theta_j^{x_j} ,$$

we have

$$\mathrm{Art}_{N/K}(\mathfrak{f}_i)(\theta) = \zeta_n^{r_i}\theta \quad \text{with} \quad r_i = \sum_{1 \le j \le u} s_{i,j}x_j .$$

Since $K(\theta)$ is stable by $\sigma \in \mathrm{Gal}(N/K)$ if and only if θ is stable, it follows that $K(\theta)$ is stable by the group $\mathcal{H}$ if and only if the following system of d linear congruences in the u unknowns x_j is satisfied:

$$\sum_{1 \le j \le u} s_{i,j}x_j \equiv 0 \pmod{n} \quad \text{for} \quad 1 \le i \le d .$$

Since by class field theory we know that the extension L/K exists and is unique, this system must have a solution. In addition, since L/K is a cyclic extension, by Kummer theory all subextensions of $K(\theta)/K$ are of the form $K(\theta^j)$ for some integer j.

It follows that if (x_j) is a solution corresponding to the desired extension L/K, then any solution to our system of congruences is of the form (λx_j) modulo n for some integer λ not necessarily prime to n. Thus the group of solutions to our system must be cyclic of order n and the desired extension L/K corresponds to any generator of this group. The solution to our system of congruences is found using Algorithm 4.1.22.

5.4.5 Algorithmic Kummer Theory When $\zeta_n \in K$ Using Artin

We are now ready to give a detailed algorithm for the construction of the class field L/K corresponding to a congruence subgroup $(\mathfrak{m}, C)$ of conductor $\mathfrak{m}$, cyclic of degree $n = \ell^r$, when $\zeta_n \in K$ (use Algorithm 5.1.2 if necessary to reduce to this case).

Algorithm 5.4.8 (Kummer Extension When $\zeta_n \in K$ Using Artin). Let $(\mathfrak{m}, C)$ be a congruence subgroup of a number field K such that $Cl_{\mathfrak{m}}(K)/\overline{C}$ is cyclic of order $n = \ell^r$ with ℓ prime and such that $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$. Assume that $\zeta_n \in K$. This algorithm outputs a defining polynomial for the Abelian extension L/K corresponding to $(\mathfrak{m}, C)$ under the Takagi correspondence. We assume computed the class group $Cl(K) = \bigoplus_i (\mathbb{Z}/c_i\mathbb{Z})\overline{\mathfrak{c}_i}$, the unit group $U(K)$, the ray class group $Cl_{\mathfrak{m}}(K) = \bigoplus_i (\mathbb{Z}/b_i\mathbb{Z})\overline{\mathfrak{b}_i}$, and the subgroup $\overline{C}$ as an HNF matrix H on the generators $\overline{\mathfrak{b}_i}$.

1. [Factor $\mathfrak{m}$ and ℓ] Using Algorithm 2.3.22 (in the absolute case), find the prime ideals dividing the finite part $\mathfrak{m}_0$ of $\mathfrak{m}$, and using [Coh0, Algorithm 6.2.9], compute the prime ideal factorization of $\ell\mathbb{Z}_K$.

2. [Compute the set S] Compute a set S_0 of prime ideals whose classes generate $Cl(K)$. For this, either search in the ideal class of each $\mathfrak{c}_i$ for a prime ideal (which will exist), or factor $\mathfrak{c}_i$ into prime ideals using Algorithm 2.3.22 and take all the prime ideal factors, or range through small prime ideals until they generate the class group. Then let S be the union of S_0 with the prime ideals divisors of $\mathfrak{m}_0$ and of $\ell\mathbb{Z}_K$ found in step 1.
3. [Compute $U_S(K)$] Set $u \leftarrow r_1 + r_2 + |S|$, where (r_1, r_2) is the signature of K. Using Algorithm 7.4.8, compute a generating set $(\eta_i)_{1 \le i \le u}$ of $U_S(K)$ (with the notation of Algorithm 7.4.8, we may take $\eta_i = \varepsilon_{i-1}$ for $1 \le i \le r_1 + r_2$ and $\eta_{r_1+r_2+i} = \gamma_i$ for $1 \le i \le |S|$).
4. [Compute $\mathfrak{m}_N$] Let $\mathfrak{m}_N$ be the modulus of K whose infinite part is that of $\mathfrak{m}$ and whose finite part is given by

$$\mathfrak{m}_{N,0} \leftarrow \prod_{\mathfrak{p} \in S,\, \mathfrak{p} \nmid \ell} \mathfrak{p} \prod_{\mathfrak{p} \mid \ell} \mathfrak{p}^{(r+1/(\ell-1))e(\mathfrak{p}/\ell)+1}$$

   (recall that $n = \ell^r$).
5. [Compute $Cl_{\mathfrak{m}_N}(K)/Cl_{\mathfrak{m}_N}(K)^n$] Using Algorithm 4.3.1, compute the SNF of $Cl_{\mathfrak{m}_N}(K)$ as $Cl_{\mathfrak{m}_N}(K) = \bigoplus_i (\mathbb{Z}/e_i\mathbb{Z})\overline{\mathfrak{e}_i}$. Let s be the largest index (0 if none exist) such that $(n, e_i) > 1$, and set $e_i \leftarrow (n, e_i)$ for $1 \le i \le s$. Finally, set $E = (\overline{\mathfrak{e}_1}, \ldots, \overline{\mathfrak{e}_s})$, $D_E = \mathrm{diag}(e_1, \ldots, e_s)$, so that

$$Cl_{\mathfrak{m}_N}(K)/Cl_{\mathfrak{m}_N}(K)^n = (E, D_E) = \bigoplus_{1 \le i \le s} (\mathbb{Z}/e_i\mathbb{Z})\overline{\mathfrak{e}_i}\ .$$

6. [Compute $\overline{D}/Cl_{\mathfrak{m}_N}(K)^n$] Using Algorithm 4.1.11, or more precisely Exercise 3 of Chapter 4, compute the kernel of the canonical surjection $\overline{s}$ from $Cl_{\mathfrak{m}_N}(K)/Cl_{\mathfrak{m}_N}(K)^n$ to $Cl_{\mathfrak{m}}(K)/\overline{C}$ (which is well-defined by Lemma 5.4.7) as a left HNF divisor H_E of D_E. Then let (F, D_F) be the SNF of $\overline{D}/Cl_{\mathfrak{m}_N}(K)^n$ obtained by applying Algorithm 4.1.3 to the system of generators and relations $(EH_E, H_E^{-1}D_E)$.
7. [Write linear system of congruences] Let $F = (\overline{\mathfrak{f}_1}, \ldots, \overline{\mathfrak{f}_d})$. For $1 \le i \le d$ and $1 \le j \le u$, using Algorithm 5.4.2 on the extensions $N_j/K = K(\theta_j)/K$, compute integers $s_{i,j}$ defined modulo n such that

$$\mathrm{Art}_{N_j/K}(\mathfrak{f}_i)(\theta_j) = \zeta_n^{s_{i,j}}\theta_j\ .$$

8. [Solve system] Using Algorithm 4.1.22 and the remark following it, compute the SNF of the group of solutions of the system $\sum_{1 \le j \le u} s_{i,j}x_j \equiv 0 \pmod{n}$ for $1 \le i \le d$. If this group is not cyclic of order n, there is a bug in the author's write-up of the algorithm or in its implementation. Otherwise, let $X = (x_1, \ldots, x_u)^t$ be a generator.
9. [Terminate] Set $\alpha \leftarrow \prod_{1 \le j \le u} \eta_j^{x_j}$, and try to reduce α by multiplying it by nth powers of the η_j or by replacing $(x_1, \ldots, x_u)^t$ by $\lambda(x_1, \ldots, x_u)^t$ modulo n for λ coprime to n. For future use, output the modulus $\mathfrak{m}_N$, output the defining polynomial $X^n - \alpha = 0$, and terminate the algorithm.

Remark. In practice it will be very costly to compute $s_{i,j}$ since the ideals $\mathfrak{f}_i$ have been obtained after a series of HNF or SNF transforms. Thus it is preferable to keep explicitly the HNF and SNF transformation matrices and to compute the image of the Artin map on the much simpler ideals $\mathfrak{e}_i$. In addition, for the same reason it may be useful to look for representatives of the class of $\overline{\mathfrak{e}_i}$ that are simpler, in particular, that are prime ideals. The details are left to the reader.

5.5 Explicit Use of the Artin Map When $\zeta_n \notin K$

We now consider the general case where we search for the class field L/K corresponding to a congruence subgroup $(\mathfrak{m}, C)$ such that $Cl_{\mathfrak{m}}(K)/\overline{C}$ is cyclic of order $n = \ell^r$, and where we do not assume that $\zeta_n \in K$.

As already explained, we proceed in three steps. In the first step, we construct the extension $K_z = K(\zeta_n)$ and compute all its necessary invariants such as its class group, unit group, and so on. This will often be a very costly part of the computation and is the main drawback of Kummer theory (unfortunately, no other completely general method is known).

In a second step we apply Algorithm 5.4.8 (using the Artin map), or for that matter Algorithm 5.2.14 when n is prime (using Hecke's theorem) to find a suitable extension L_z/K_z corresponding to the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K_z}, \mathcal{N}_{K_z/K}^{-1}(C))$ of K_z. Note that if we use Hecke's theorem, we must find the conductor of this congruence subgroup, while when using the Artin map, this is not necessary. In any case, at the end of this step we have found $\alpha \in K_z$ such that $L_z = K_z(\sqrt[n]{\alpha})$.

In a final step, we must construct the extension L/K, by going down from the extension L_z/K_z. We have already seen how to do this using Lagrange resolvents. In this section, we will see another method based once again on the use of the Artin map.

5.5.1 The Extension K_z/K

Let $K_z = K(\zeta_n)$. The following proposition generalizes Proposition 5.3.2.

Proposition 5.5.1. *There exists a subgroup G_n of $(\mathbb{Z}/n\mathbb{Z})^*$ such that the extension K_z/K is Abelian with Galois group G_z given by $G_z = \{\tau_a / \ a \in G_n\}$, where τ_a is the K-automorphism of K_z sending ζ_n to ζ_n^a. In particular, $[K_z : K]$ divides $\phi(n)$. Conversely, if $a \notin G_n$, there does not exist a K-automorphism τ_a of K_z such that $\tau_a(\zeta_n) = \zeta_n^a$.*

Proof. By Proposition 5.3.1, we deduce as for Proposition 5.3.2 that K_z/K is Abelian with Galois group isomorphic to a subgroup of $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$, hence to a subgroup of $(\mathbb{Z}/n\mathbb{Z})^*$, where $a \in (\mathbb{Z}/n\mathbb{Z})^*$ corresponds to $\tau_a \in \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ by $\tau_a(\zeta_n) = \zeta_n^a$. For the converse, we can either note that all

the K-automorphisms of K_z are accounted for by the τ_a for $a \in G_n$ or, more positively, notice that G_n is the set of $a \in (\mathbb{Z}/n\mathbb{Z})^*$ such that ζ_n^a is a root of $f(X) = 0$, where $f(X)$ is the minimal polynomial of ζ_n in $K[X]$ (see also Exercise 11). □

Once the extension K_z/K is understood, we must compute a number of invariants of K_z. As in Algorithm 5.3.17 (see the discussion just before that algorithm), to perform these computations in K_z, we use either Algorithm 2.1.9 combined with the analog of Theorem 2.1.14 proved in Exercise 9 of Chapter 2, or the relative methods explained just before Algorithm 5.3.17. Once again, for simplicity of exposition, we will assume that we use the first method, but a serious implementation should use the second method.

5.5.2 The Extensions L_z/K_z and L_z/K

• **Description of** $\mathrm{Gal}(L_z/K_z)$ **and** $\mathrm{Gal}(L_z/K)$

Once K_z is constructed, we must lift the class field problem from K to K_z. We have $L_z = L(\zeta_n) = K_zL$, and as above we know, thanks to Proposition 3.5.5, that the congruence subgroup $(\mathfrak{m}\mathbb{Z}_{K_z}, \mathcal{N}_{K_z/K}^{-1}(C))$ of K_z corresponds to L_z/K_z under the Takagi correspondence, where $\mathfrak{m}\mathbb{Z}_{K_z}$ is not necessarily its conductor. Note that an added benefit of Fieker's method described above compared to Hecke's is that we do not need to compute the conductor of L_z/K_z, since this would probably waste more time than would be gained by having a smaller modulus.

By Proposition 5.3.1 applied to the extensions L/K and K_z/K, we know that $\mathrm{Gal}(L_z/K_z) = \mathrm{Gal}(LK_z/K_z)$ can be identified with a subgroup of $\mathrm{Gal}(L/K)$. Since by assumption this group is cyclic of order n, it follows that $\mathrm{Gal}(L_z/K_z)$ is cyclic of order dividing n. Furthermore, the same proposition tells us that L_z/K is an Abelian extension. From this, we can easily prove the following generalization of (part of) Theorem 5.3.5.

Proposition 5.5.2. *Let K be a number field, $K_z = K(\zeta_n)$, let*

$$G_z = \mathrm{Gal}(K_z/K) = \{\tau_a / \ a \in G_n\}$$

for some subgroup G_n of $(\mathbb{Z}/n\mathbb{Z})^*$ as in Proposition 5.5.1, let L/K be a cyclic extension of degree n, and let $L_z = LK_z = L(\zeta_n)$. Then L_z/K_z is a cyclic extension of degree m dividing n, and if we write $L_z = K_z(\theta)$ with $\theta = \sqrt[m]{\alpha}$ for some $\alpha \in K_z$, the extension L_z/K is Abelian if and only if for each $a \in G_n$ there exists $\gamma_a \in K_z$ such that

$$\tau_a(\alpha) = \gamma_a^m \alpha^a \ .$$

If this condition is satisfied, we can choose an extension of τ_a to L_z (again denoted by τ_a) such that $\tau_a(\theta) = \gamma_a \theta^a$, and we have

$$\mathrm{Gal}(L_z/K) = \{\sigma^j \tau_a / \ 0 \le j < m, a \in G_n\} \ ,$$

where σ is the K_z-automorphism of L_z such that $\sigma(\theta) = \zeta_m \theta$.

Proof. We have already seen that L_z/K_z is a cyclic extension of degree m dividing n. Assume that the extension L_z/K is Abelian. Since the extension L_z/K is normal we must have $\tau_a(\theta) \in L_z$ for any extension of τ_a to L_z, which we denote again by τ_a by abuse of notation. Thus $K_z(\tau_a(\theta)) \subset L_z = K_z(\theta)$. Applying this to $\tau_a^{-1} = \tau_{a^{-1}}$, we obtain $K_z(\tau_a^{-1}(\theta)) \subset K_z(\theta)$, hence $K_z(\theta) \subset K_z(\tau_a(\theta))$, which finally shows that we have the equality $K_z(\tau_a(\theta)) = K_z(\theta)$. Since $\tau_a(\theta)^m = \tau_a(\theta^m) = \tau_a(\alpha)$, it follows from the uniqueness theorem of Kummer theory (Corollary 10.2.7 in our case) that there exist $\gamma_a \in K_z$ and b coprime to m such that $\tau_a(\alpha) = \gamma_a^m \alpha^b$, hence we can choose $\tau_a(\theta) = \gamma_a \theta^b$ (all the other extensions τ of τ_a to L_z are obtained from this one by setting $\tau(\theta) = \zeta_m^s \gamma_a \theta^b$ for $0 \le s < m$).

We now use the fact that L_z/K is not only normal but Abelian. Let σ be the generator of the cyclic group $\mathrm{Gal}(L_z/K_z)$, which sends θ to $\zeta_m \theta$ and leaves K_z fixed (note that $\zeta_m = \zeta_n^{n/m} \in K_z$). Then

$$\sigma(\tau_a(\theta)) = \sigma(\gamma_a \theta^b) = \gamma_a \zeta_m^b \theta^b \ ,$$

while

$$\tau_a(\sigma(\theta)) = \tau_a(\zeta_m \theta) = \zeta_m^a \gamma_a \theta^b \ .$$

Comparing, we obtain $b \equiv a \mod m$; in other words, $b = a$ since b is only defined modulo m anyway. This shows that $\tau_a(\alpha) = \gamma_a^m \alpha^a$ as claimed.

Conversely, if this condition is satisfied, then for all $a \in G_n$, τ_a and σ commute on the generating elements θ and ζ_n of L_z/K; hence L_z/K is Abelian. The last statement is clear. □

Corollary 5.5.3. *Keep the hypotheses and notation of the proposition, and assume that $L_z = K_z(\theta)$ is an Abelian extension of K, with $\theta^m = \alpha$. Let $\mathfrak{p}$ be a prime ideal of K not dividing n or α, and denote as usual by $\sigma_\mathfrak{p}$ the Frobenius automorphism associated to $\mathfrak{p}$ in the extension L_z/K. Then*

$$\sigma_\mathfrak{p}(\theta) = \zeta_m^{u_\mathfrak{p}} \gamma_{\mathcal{N}(\mathfrak{p})} \theta^{\mathcal{N}(\mathfrak{p})} \ ,$$

where $u_\mathfrak{p}$ is the unique integer modulo m (which exists) such that

$$\zeta_m^{u_\mathfrak{p}} \gamma_{\mathcal{N}(\mathfrak{p})} \equiv 1 \ (\mathrm{mod}\ \mathfrak{p}\mathbb{Z}_{K_z}) \ .$$

Proof. Since the extension K_z/K is ramified only at prime ideals dividing n and L_z/K_z is ramified only at prime ideals dividing n and α, it follows that $\mathfrak{p}$ is unramified in L_z/K; hence $\sigma_\mathfrak{p}$ is well-defined.

We have $\sigma_\mathfrak{p}(\zeta_m) \equiv \zeta_m^{\mathcal{N}(\mathfrak{p})} \ (\mathrm{mod}\ \mathfrak{p}\mathbb{Z}_{K_z})$, and since $\mathfrak{p} \nmid n$ we have in fact the equality $\sigma_\mathfrak{p}(\zeta_m) = \zeta_m^{\mathcal{N}(\mathfrak{p})}$. It follows that $\sigma_\mathfrak{p}|_{K_z} = \tau_{\mathcal{N}(\mathfrak{p})}$, hence by the proposition $\sigma_\mathfrak{p}(\alpha) = \gamma_{\mathcal{N}(\mathfrak{p})}^m \alpha^{\mathcal{N}(\mathfrak{p})}$, so

$$\sigma_{\mathfrak{p}}(\theta) = \zeta_m^u \gamma_{\mathcal{N}(\mathfrak{p})} \theta^{\mathcal{N}(\mathfrak{p})}$$

for some u. We conclude by using the congruence $\sigma_{\mathfrak{p}}(\theta) \equiv \theta^{N(\mathfrak{p})} \pmod{\mathfrak{p}\mathbb{Z}_{L_z}}$ and the fact that $\mathfrak{p} \nmid \alpha$. □

• **Computation of the γ_a**

It will be essential to compute the γ_a explicitly for $a \in G_n$. Since G_n is a subgroup of $(\mathbb{Z}/n\mathbb{Z})^*$ and $n = \ell^r$, it follows that G_n is cyclic if $\ell > 2$ and is generated by at most two elements if $\ell = 2$. By Exercise 12, to compute all the γ_a it is sufficient to compute γ_a for the generator(s) of G_n, hence at most two computations.

To compute such a γ_a, we can proceed as follows. Let $\beta \leftarrow \dot{\tau}_a(\alpha)/\alpha^m \in K_z$. Then the roots in K_z of the equation $X^m - \beta = 0$ are the numbers $\zeta_m^j \gamma_a$, and any one is suitable for our purposes. Thus, we have a polynomial of degree m with coefficients in K_z which is known to be a product of m distinct linear factors, and we want to find one of them. We may of course use a general factoring algorithm over number fields, but in this special case it is more efficient to write a special-purpose algorithm to solve this problem. The details are left to the reader (Exercise 13).

• **Finding a Suitable Modulus for L_z/K**

Recall from Definition 3.4.2 that a modulus $\mathfrak{m}_L$ is *suitable* for the extension L_z/K if $\mathfrak{m}_L$ is a multiple of the conductor of L_z/K or, equivalently, if L_z/K is a subextension of the ray class field $K(\mathfrak{m}_L)$. Since we know that the modulus $\mathfrak{m}_N$ is suitable for the extension L_z/K_z and since L_z/K is Abelian, it follows from Proposition 3.5.6 that $\mathcal{N}_{K_z/K}(\mathfrak{m}_N)$ is suitable for the extension L_z/K.

On the other hand, by definition $\mathfrak{m}$ is suitable for L/K, and $n\mathbb{Z}_K$ is suitable for K_z/K, hence $\mathrm{lcm}(\mathfrak{m}, n\mathbb{Z}_K)$ is suitable for $LK_z/K = L_z/K$.

Thus, we can take as suitable modulus for L_z/K the modulus

$$\mathfrak{m}_L = \gcd(\mathrm{lcm}(\mathfrak{m}, n\mathbb{Z}_K), \mathcal{N}_{K_z/K}(\mathfrak{m}_N)) \ .$$

Note that this will in general be simpler than $\mathcal{N}_{K_z/K}(\mathfrak{m}_N)$ since the prime ideals generating the class group are not necessary.

5.5.3 Going Down to the Extension L/K

Once $\mathfrak{m}_L$ is chosen, we use Lemma 5.4.7 applied to the Abelian extension L_z/K. Thus, we have $L = L_z^{\mathcal{H}}$ with

$$\mathcal{H} = \mathrm{Art}_{L_z/K}(\overline{D}) = \mathrm{Art}_{L_z/K}(\overline{D}/Cl_{\mathfrak{m}_L}(K)^n) \ ,$$

where $\overline{D}$ is the kernel of the canonical surjection from $Cl_{\mathfrak{m}_L}(K)$ to $Cl_{\mathfrak{m}}(K)/\overline{C}$. As in the case $\zeta_n \in K$, but in a different context, we must make this more explicit.

When the extension L/K is cyclic of prime degree ℓ, we have seen in Theorem 5.3.5 that we can come down from L_z to L by taking $L = K(\alpha)$ with $\alpha = \mathrm{Tr}_{L_z/L}(\theta)$. In our more general case where L/K is cyclic of prime power order, this is not necessarily true. We have, however, the following easy lemma.

Lemma 5.5.4. *Assume that η is such that $L_z = K(\eta)$, and let*

$$P_\eta(X) = \sum_{i=0}^{d} (-1)^i t_i X^{d-i}$$

be the characteristic polynomial of η in $L[X]$ with $d = [L_z : L]$. Then $L = K(t_i)$ for at least one value of i (note that $K(t_i) \subset L$ for all i).

Proof. Since L/K is cyclic of prime power degree ℓ^r, there exists a maximal nontrivial subextension L_1/K of degree ℓ^{r-1}. So assume the conclusion of the lemma is false. Then $t_i \in L_1$ for all i, so that $[L_z : L_1] \leq d$, which is absurd since $[L_z : L_1] = d\ell$. □

Remarks

(1) To apply this lemma we need η such that $L_z = K(\eta)$. Since $L_z = K(\zeta_n, \theta)$, the primitive element theorem tells us that we can choose $\eta = \theta + q\zeta_n$ for some small integer q. Note that here it is useless to consider also $\eta = \theta\zeta_n + q\zeta_n = (\theta + q)\zeta_n$ since θ and $\theta\zeta_n$ are both nth roots of $\alpha = \theta^n \in K_z$, so they play exactly the same role.

(2) In practice, we can hope that $\mathrm{Tr}_{L_z/L}(\eta) = t_1$ already satisfies $L = K(t_1)$, so we try it first, and then we try the other coefficients until a suitable one is found.

To obtain the conjugates of η over L, we must apply the elements of $\mathrm{Gal}(L_z/L)$, which by the remark made at the beginning of this section are simply the elements of the subgroup $\mathcal{H}$ of $Cl_{\mathfrak{m}_L}(K)/Cl_{\mathfrak{m}_L}(K)^n$ given above. The explicit action of the Artin map on η, in other words on ζ_n and on θ, is given by Corollary 5.5.3.

5.5.4 Algorithmic Kummer Theory When $\zeta_n \notin K$ Using Artin

We are now ready to describe the complete algorithm for computing the ray class field associated to a congruence subgroup $(\mathfrak{m}, C)$ using the Artin map. As in the rest of this chapter, we assume that we have done the preliminary reduction to the case where $Cl_{\mathfrak{m}}(K)/\overline{C}$ is cyclic of prime power degree $n = \ell^r$ by using Algorithm 5.1.2.

Algorithm 5.5.5 (Kummer Extension When $\zeta_n \notin K$ Using Artin). Let $K = \mathbb{Q}(\alpha_K)$ be a number field and let $(\mathfrak{m}, C)$ be a congruence subgroup of K such

that $Cl_{\mathfrak{m}}(K)/\overline{C}$ is cyclic of order $n = \ell^r$ for some prime ℓ. This algorithm outputs a defining polynomial for the Abelian extension L/K corresponding to the congruence subgroup $(\mathfrak{m}, C)$ under the Takagi correspondence. We denote by $T_K(X)$ the minimal monic polynomial of α_K in $\mathbb{Q}[X]$, assumed to be in $\mathbb{Z}[X]$.

1. [Adjoin ζ_n] Using Algorithm 2.1.9, compute a compositum K_z of K with $\mathbb{Q}(\zeta_n)$ given by the defining polynomial $\Phi_n(X) = 0$, where

$$\Phi_n(X) = \frac{X^n - 1}{X^{n/\ell} - 1} = X^{\ell^{r-1}(\ell-1)} + X^{\ell^{r-1}(\ell-2)} + \cdots + 1$$

   is the nth cyclotomic polynomial (there may be several factors in the compositum, but they all define isomorphic number fields so any one can be taken). The algorithm outputs an integer k, an irreducible monic polynomial $R(X) \in \mathbb{Z}[X]$ having $\alpha_z = \zeta_n(\alpha_K + k)$ as a root such that $K_z = \mathbb{Q}(\alpha_z)$, and polynomials $A_1(X)$ and $A_2(X)$ such that $\alpha_K = A_1(\alpha_z)$ and $\zeta_n = A_2(\alpha_z)$.

2. [Compute the action of the τ_a on α_z] For each $a \in (\mathbb{Z}/n\mathbb{Z})^*$, compute $U_a(X) \leftarrow XA_2(X)^{a-1} \bmod R(X)$ (we will have $U_a(\alpha_z) = \tau_a(\alpha_z)$).

3. [Compute the group G_n] Set $d_z \leftarrow \deg(R(X))/\deg(T_K(X))$. By simple enumeration, compute the set G_n of $a \in (Z/n\mathbb{Z})^*$ such that $R(X) \mid R(U_a(X))$ (we must have $|G_n| = [K_z : K] = d_z$).

4. [Compute data for K_z] Using the standard algorithms for the absolute case given in [Coh0], or the algorithm for the relative case described in Section 7.3.3, compute an integral basis, the unit group $U(K_z)$, and the class group $Cl(K_z)$ (see Remark (1) below).

5. [Lift congruence subgroup] Set $\mathfrak{m}_z \leftarrow \mathfrak{m}\mathbb{Z}_{K_z}$, and compute the SNF of $Cl_{\mathfrak{m}_z}(K_z)$. Using Algorithm 4.1.11, compute $C_z \leftarrow \mathcal{N}^{-1}_{K_z/K}(C)$ as a subgroup of $Cl_{\mathfrak{m}_z}(K_z)$ given by an HNF matrix H_{C_z}. Set $m_z \leftarrow \det(H_{C_z})$ (this will be the degree of the desired extension L_z/K_z, hence of the form ℓ^{r_z} with $r_z \le r$).

6. [Apply Kummer] Apply Algorithm 5.4.8 to the congruence subgroup $(\mathfrak{m}_z, C_z)$ of K_z (we already know that $Cl_{\mathfrak{m}_z}(K_z)/\overline{C_z}$ is cyclic of prime power order). Let $\mathfrak{m}_N$ be the modulus of K_z and $X^{m_z} - \alpha$ be the defining polynomial for L_z/K_z output by this algorithm. Set $\theta \leftarrow \sqrt[m_z]{\alpha}$.

7. [Compute suitable modulus for L_z/K] Compute

$$\mathfrak{m}_L \leftarrow \gcd(\mathrm{lcm}(\mathfrak{m}, n\mathbb{Z}_K), \mathcal{N}_{K_z/K}(\mathfrak{m}_N))$$

   as a modulus of K (see Remark (2) below).

8. [Compute $Cl_{\mathfrak{m}_L}(K)/Cl_{\mathfrak{m}_L}(K)^n$] Using Algorithm 4.3.1, compute the SNF of $Cl_{\mathfrak{m}_L}(K)$ as $Cl_{\mathfrak{m}_L}(K) = \bigoplus_i (\mathbb{Z}/e_i\mathbb{Z})\overline{\mathfrak{e}_i}$. Let s be the largest index (0 if none exist) such that $(n, e_i) > 1$, and set $e_i \leftarrow (n, e_i)$ for $1 \le i \le s$. Finally, set $E = (\overline{\mathfrak{e}_1}, \ldots, \overline{\mathfrak{e}_s})$, $D_E = \mathrm{diag}(e_1, \ldots, e_s)$, so that

$$Cl_{\mathfrak{m}_L}(K)/Cl_{\mathfrak{m}_L}(K)^n = (E, D_E) = \bigoplus_{1 \le i \le s} (\mathbb{Z}/e_i\mathbb{Z})\overline{\mathfrak{e}_i} \ .$$

9. [Compute $\overline{D}/Cl_{\mathfrak{m}_L}(K)^n$] Using Algorithm 4.1.11, or more precisely Exercise 3 of Chapter 4, compute the kernel of the canonical surjection $\overline{s}$ from $Cl_{\mathfrak{m}_L}(K)/Cl_{\mathfrak{m}_L}(K)^n$ to $Cl_{\mathfrak{m}}(K)/\overline{C}$ (which is well-defined by Lemma 5.4.7) as a left HNF divisor H_E of D_E. Then let (F, D_F) be the SNF of $\overline{D}/Cl_{\mathfrak{m}_N}(K)^n$ obtained by applying Algorithm 4.1.3 to the system of generators and relations $(EH_E, H_E^{-1}D_E)$.
10. [Compute the γ_a] Using the method explained above and the action of τ_a on α_z computed in step 2, compute $\gamma_a \leftarrow \sqrt[m]{\tau_a(\alpha)/\alpha^a}$ for the (at most two) generators of G_n, and using the cocycle condition given in Exercise 12 compute the γ_a for all $a \in G_n$.
11. [Compute a suitable η] For $q = 0, \pm 1$, and so on, set $\eta \leftarrow \theta + q\zeta_n$ until the elements $\gamma_a \theta^a \zeta_m^{aj} + q\zeta_n^a$ are distinct for $0 \leq j < m$ and $a \in G_n$ (we now have $L_z = K(\eta)$).
12. [Compute $P_\eta(X)$] Compute

$$P_\eta(X) \leftarrow \prod_{\mathfrak{a} \in \overline{D}/Cl_{\mathfrak{m}_L}(K)^n} (X - \mathrm{Art}_{L_z/K}(\mathfrak{a})(\eta)) \in L_z[X] \ ,$$

where $\mathrm{Art}_{L_z/K}(\mathfrak{a})(\eta)$ is computed as explained above, using the action of the τ_a on α_z computed in step 2 and Corollary 5.5.3 (see Exercise 14). Write $P_\eta(X) = \sum_{i=0}^{d} (-1)^i t_i X^{d-i}$ with $t_i \in L_z$ represented as explained in Remark (4) below, and set $j \leftarrow 0$.

13. [Loop on coefficients] Set $j \leftarrow j + 1$. If $j > d$, there is a bug in the implementation or in the algorithm, and terminate. Otherwise set $t \leftarrow t_j$.
14. [Compute minimal polynomial] Using the explicit description of $\mathrm{Gal}(L_z/K)$ given by Proposition 5.5.2, compute all the conjugates of t in L_z. If there are not exactly n distinct conjugates, go to step 13. Otherwise, let $t^{(j)}$ be the n distinct conjugates of t and set $m_t(X) \leftarrow \prod_j (X - t^{(j)})$, which will be the minimal polynomial of t over K (see Remark (5) below for an alternate and often better way to compute $m_t(X)$).
15. [Compute $\mathbb{Z}_L$] Compute an integral pseudo-basis of $\mathbb{Z}_{L_z}/\mathbb{Z}_{K_z}$ (which is easy since this is a Kummer extension); then using the integral pseudo-basis of $\mathbb{Z}_{K_z}/\mathbb{Z}_K$, compute an integral pseudo-basis of $\mathbb{Z}_{L_z}/\mathbb{Z}_K$; finally, using Exercise 35 of Chapter 2, compute an integral pseudo-basis of $\mathbb{Z}_L/\mathbb{Z}_K$.
16. [Terminate] Using a relative polynomial reduction algorithm such as Algorithm 2.4.12 and the integral pseudo-basis of $\mathbb{Z}_L$ computed in the preceding step, replace the polynomial $m_t(X)$ by a polynomial that is as reduced as possible, output $m_t(X)$ as a defining polynomial for L/K and terminate the algorithm.

Remarks

(1) Note that it is easy to compute the integral basis of K_z using the relative round 2 algorithm (Algorithm 2.4.9) since the defining polynomial

of K_z/K is a divisor of $\Phi_n(X)$, hence its discriminant is divisible only by prime ideals above ℓ. It would, in general, be much more difficult to compute this integral basis using absolute algorithms since the factorization of the discriminant of the absolute defining polynomial could be difficult.

(2) Recall that the GCD of two moduli is obtained by adding the finite parts and intersecting the infinite parts, while the LCM of two moduli is obtained by intersecting the finite parts and making the union of the infinite parts.

(3) Steps 8 and 9 are identical (with a different modulus) to steps 5 and 6 of Algorithm 5.4.8.

(4) After step 11, we know that $L_z = K(\eta)$ for $\eta = \theta + q\zeta_n$. It is, however, much more convenient to keep the representation $L_z = K(\theta, \zeta_n)$, in other words to consider the elements of L_z as polynomials in the two variables θ and ζ_n, since the action of the Galois group of L_z/K is much simpler to write in this representation. In effect, we do not need η to represent the field L_z/K by a primitive element, but only to write an explicit characteristic polynomial $P_\eta(X)$.

(5) As already mentioned, in step 12, instead of computing the complete polynomial $P_\eta(X)$, it is in general faster to simply compute its coefficient t_1 using

$$t_1 \leftarrow \mathrm{Tr}_{L_z/L}(\eta) = \sum_{\overline{\mathfrak{a}} \in \overline{D}/Cl_{\mathfrak{m}_L}(K)^n} \mathrm{Art}_{L_z/K}(\mathfrak{a})(\eta) \ ,$$

and to test in step 14 if it is suitable to generate the extension L/K. Only in the case where it is not suitable (which happens very rarely in practice), we compute the complete polynomial $P_\eta(X)$.

(6) In step 14, an alternate way to compute $m_t(X)$ is the following. For $0 \le j \le n$ compute t^j in L_z (expressed as polynomials in ζ_n and θ as explained in remark (4)). Then solve the linear system $t^n = \sum_{0 \le j \le n-1} x_j t^j$, which is a system of $[L_z : K]$ equations in n unknowns with coefficients in K to find the polynomial $m_t(X)$. We may of course apply the standard Gaussian pivoting methods to solve this system, but considering the size and the complexity of the coefficients, it is preferable in an actual implementation to use modular techniques to solve the system.

(7) In step 16, we could directly try to reduce the polynomial $m_t(X)$ using Algorithm 2.4.12. The direct computation of the relative integral pseudo-basis of $\mathbb{Z}_L/\mathbb{Z}_K$ could involve expensive discriminant factorization, hence it is more efficient to proceed as explained in the algorithm. We have already mentioned this in the Hecke case.

5.5.5 Comparison of the Methods

We have now described in great detail several methods for computing class fields. It is necessary to give some practical advice on which methods to use.

For evident compatibility reasons, if the necessary roots of unity are in the base field K, we set $K_z = K$ and $L_z = L$.

(1) A common point to all the methods is the necessity and usefulness to reduce the problem to several problems involving the computation of class fields which are cyclic of prime power order $n = \ell^r$. This is easily done using Algorithm 5.1.2. We assume this reduction made.

(2) If $r > 1$, in other words if n is not a prime number, then we must use the method using the Artin map for the construction of the extension L_z/K_z. Indeed, Hecke's theorem is not directly applicable in that case, and it would be necessary to construct the extension as a tower of r cyclic extensions of order ℓ. This is clearly not practical if $r > 1$, except perhaps in very small cases such as $n = 4 = 2^2$ (see Exercise 18).

(3) If $r = 1$, in other words if $n = \ell$ is prime, then both the method using the Artin map and the method using Hecke's theorem may be considered. The comparison between the two methods depends probably on the efficiency of the implementation of the underlying algorithms for computing more basic objects such as class groups or the structure of $(\mathbb{Z}_K/\mathfrak{m})^*$ for the necessary moduli $\mathfrak{m}$. In addition, in both methods special-purpose algorithms can be written to speed up the necessary tasks without using general methods. Thus, for a serious and efficient implementation, I advise to implement both methods and to compare, doing some detailed profiling of both programs, to see where most of the time is spent.

(4) To go down from L_z to L (when the necessary roots of unity did not belong to the base field), we again have two methods. Although we have not explained the method using Lagrange resolvents in the case of nonprime degree, the theory can be generalized, at least when G_n is cyclic, which is always the case when $\ell > 2$ or when $n = 4$ (see Exercise 17 for this last case). As can already be seen from the expressions given in Section 5.3.2, the formulas become quite complicated, although they need not be written explicitly but computed when necessary. Thus as practical advice, to go down from L_z to L, I suggest using Lagrange resolvents for $n = 3$, 4, 5, and 7, and the method using the Artin map for larger values of n.

5.6 Two Detailed Examples

To illustrate the above results, we give two examples coming from the numerical results given in Section 12.2.2. The first example is that of a quadratic extension, for which it is not necessary to adjoin roots of unity. The second example is that of a cubic extension. For both these examples, we will use Hecke's theorem, but we advise the reader to perform similar computations using the Artin map and then compare (Exercise 19).

Warning. The quantities that we will compute (for example, a system of fundamental units) depend on the precise implementation, hence the results

that you will obtain will almost certainly be different from those presented here. Only the final reduced polynomial should essentially be the same.

5.6.1 Example 1

In the first example, we take $K = \mathbb{Q}(z)$ as base field, where z is a root of the polynomial $X^6 - X^5 + 2X^3 - 2X^2 + 1$. In this field, the prime number 41 splits as the product of three prime ideals of degree 1 and one prime ideal of degree 3. One of the prime ideals of degree 1 is equal to $\mathfrak{P}_{41} = 41\mathbb{Z}_K + (z+4)\mathbb{Z}_K$. The number field K has two real places ∞_1 and ∞_2. We take as modulus $\mathfrak{m} = \mathfrak{P}_{41}\infty_1\infty_2$ and $C = P_{\mathfrak{m}}$ as congruence subgroup. Using Algorithm 4.3.1, we find that the ray class group is of order 2, that $\mathfrak{m}$ is the conductor, hence that there exists a quadratic extension L of K ramified only at primes above $\mathfrak{m}$, hence totally complex and ramified only at the finite prime $\mathfrak{P}_{41}$. Corollary 3.5.12 tells us that its relative discriminant ideal is equal to $\mathfrak{P}_{41}$ itself, and we now want a defining polynomial for L/K.

Since L/K is quadratic, we have $\zeta_\ell = -1 \in K$, hence $L = K(\sqrt{\alpha})$ for some $\alpha \in \mathbb{Z}_K$. We use Algorithm 5.2.14. Since 41 is prime to $\ell = 2$, we have $S_{\mathfrak{m},\ell,i} = \varnothing$ for $i = 1$, 2, and 3. The set $S_{\mathfrak{m}}$ has the unique element $\mathfrak{P}_{41}$, and the set S_ℓ has the unique element $2\mathbb{Z}_K$, since it is easily checked that 2 is inert in K. Thus, the tests in steps 2 and 3 of the algorithm are satisfied.

To continue, we need the virtual units v_j. Using the general class group and unit group algorithms, we find that the class group is trivial and that we can choose as generators of the unit group $\varepsilon_0 = -1$, $\varepsilon_1 = z$, $\varepsilon_2 = z^3 + 1$, and $\varepsilon_3 = z^4 - z^3 + z^2 + z - 1$.

Using the principal ideal algorithm in K ([Coh0, Algorithm 6.5.10]), we compute that $\mathfrak{P}_{41} = \beta_{\mathfrak{p}}\mathbb{Z}_K$ with $\beta_{\mathfrak{p}} = z^5 + 2z^2 - 2z$; hence according to step 5, we set $v_i = \varepsilon_{i-1}$ for $1 \le i \le 4$ and $v_5 = \beta_{\mathfrak{p}}$.

Finally, since $S_{\mathfrak{m},\ell,2} = \mathfrak{m}'_\infty = \varnothing$, the only modulus to consider in step 5 is $\mathfrak{m}_1 = (2\mathbb{Z}_K)^{z(2\mathbb{Z}_K,2)-1} = 4\mathbb{Z}_K$ since $z(2\mathbb{Z}_K, 2) = 2e(2\mathbb{Z}_K/2) + 1 = 3$.

Applying Algorithm 4.2.21, we find that

$$(\mathbb{Z}_K/4\mathbb{Z}_K)^* \simeq (\mathbb{Z}/126\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})^5$$

with generators, for example, equal to z, $1 - 2z^2$, $1 - 2z$, $1 - 2z^4$, $1 - 2z^5$, and -1 (your own generators may, of course, be different).

Using Algorithm 4.2.24, we find that the matrix M of discrete logarithms of the v_j for $1 \le j \le 5$ is equal to

$$M = \begin{pmatrix} 0 & 1 & 97 & 11 & 68 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 \end{pmatrix},$$

By Gaussian elimination in $\mathbb{F}_2$, we find that the column vector $K_1 = (0,0,1,1,1)^t$ is a generator of the kernel modulo 2. Thus, if we set

$$\alpha = \varepsilon_2\varepsilon_3\beta_{\mathfrak{p}} = -z^5 - z^4 + z^3 - 3z^2 \ ,$$

the ramification conditions at all the finite primes are satisfied. Since the dimension of our kernel is equal to 1, it is not necessary to do any backtracking or to compute more discrete logarithms. We simply check that the real places are ramified, in other words that $\sigma(\alpha) < 0$ for both real places σ (if this was not the case, there would be no Abelian extension of conductor $\mathfrak{m}$ and degree 2).

Thus L can be defined over K by the relative defining polynomial

$$X^2 - (-z^5 - z^4 + z^3 - 3z^2) = 0 \ .$$

To get the absolute defining polynomial of L over $\mathbb{Q}$, we use Algorithm 2.1.11, and we obtain

$$X^{12} + 10X^{10} + 41X^8 + 121X^6 + 196X^4 + 147X^2 + 41$$

as the defining polynomial for our number field L.

Using polynomial reduction techniques (see [Coh0, Section 4.4], [Coh-Dia] or Algorithm 2.4.12), we obtain the final reduced polynomial

$$X^{12} - 2X^{11} + 2X^{10} - X^9 + 2X^8 - 5X^7 + 8X^6 - 7X^5 + 4X^4 - 3X^3 + 4X^2 - 3X + 1$$

given in Chapter 12.

5.6.2 Example 2

As a second example, we take $K = \mathbb{Q}(z)$ as base field, where z is a root of the polynomial $X^6 - 2X^5 + 3X^4 + X^2 + 3X + 1$. This is a totally complex number field in which 2 is inert, and we choose $\mathfrak{m} = 2\mathbb{Z}_K$ and $C = P_{\mathfrak{m}}$.

Using Algorithm 4.3.1, we find that the ray class group is of order 3, that $\mathfrak{m}$ is the conductor, hence that there exists a cubic extension L of K ramified only at $\mathfrak{m}$. Since $\mathfrak{m}$ is the conductor and $\ell = 3$ is prime, its relative discriminant ideal is equal to $\mathfrak{m}^2 = 4\mathbb{Z}_K$. We now want a defining polynomial for L/K. Since $\zeta_3 \notin K$, we must begin by adjoining ζ_3 to K. Set $K_z = K(\zeta_3)$ and $L_z = L(\zeta_3)$.

Thanks to Algorithm 2.1.8, we know that we can choose $k = -1$, in other words that $y_1 = \zeta_3 - z$ is a primitive element of K_z, and this algorithm also gives us a polynomial $P_1(X)$ and a polynomial $A_1(X)$ such that $z = A_1(y_1)$ (as explained above, it would be preferable at this stage to use Algorithm 2.1.9, but for such a small example it does not make much difference).

As indicated after Algorithm 2.1.8, we then use a polynomial reduction algorithm that outputs the reduced polynomial

$$X^{12}-2X^{11}+X^{10}-6X^9+8X^8+7X^7+5X^6-20X^5-2X^4+3X^3+8X^2+3X+1$$

(with root y, say) and a polynomial $B(X)$ such that $y = B(y_1)$. Using Algorithm 2.1.12, we can compute a polynomial $B^{-1}(X)$ such that $y_1 = B^{-1}(y)$. Replacing this value of y_1, we obtain $z = A_1(B^{-1}(y))$ and $\zeta_3 = z + B^{-1}(y)$ as polynomials in y. It is now easy to determine the action of τ on y. We have $\tau(\zeta_3) = \zeta_3^2 = -1 - \zeta_3$, hence $z + B^{-1}(\tau(y)) = -1 - z - B^{-1}(y)$; in other words,

$$B^{-1}(\tau(y)) = -1 - 2z - B^{-1}(y) = -1 - 2A_1(B^{-1}(y)) - B^{-1}(y) \ .$$

It follows finally that

$$\tau(y) = B(-1 - 2A_1(B^{-1}(y)) - B^{-1}(y))$$

gives the action of τ on y. We did not use the method given in step 2 of Algorithm 5.3.17 as written, since it is easier to compute $-1 - \zeta_3$ than ζ_3^2 (of course, the result is the same).

I do not explicitly give the formulas for most of the polynomials since they are rather complicated. As a correctness check, we verify that $\tau(\tau(y)) = y$.

Continuing to loosely follow Algorithm 5.3.17, we compute the class group and units of K_z. We find that the class group is trivial, and we find a generator ε_0 of the torsion units as well as a system of fundamental units ε_i for $1 \le i \le 6$.

Using the action of τ computed above, we find that

$$T_v = \begin{pmatrix} -1 & 2 & 0 & -2 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix} ,$$

and hence the kernel of $\overline{T_v - 2I_2}$ modulo 3 is an $\mathbb{F}_3$-vector space of dimension 4 generated by the vectors $(1,0,0,0,0,0)^t$, $(0,1,0,1,0,0)^t$, $(0,0,1,0,1,0)^t$, and $(0,0,1,0,0,1)^t$. Thus, we let $w_1 = \varepsilon_0$, $w_2 = \varepsilon_2\varepsilon_4$, $w_3 = \varepsilon_3\varepsilon_5$, $w_4 = \varepsilon_3\varepsilon_6$, and the classes of the w_i form an $\mathbb{F}_3$-basis of $e_1 V_3(K_z)/K_z^{*3}$.

Since $\mathfrak{m} = 2\mathbb{Z}_K$ is prime to 3, the conductor $\mathfrak{f}$ of L_z/K_z is equal to $2K_z$. In $\mathbb{Z}_{K_z}$, we find that $\mathfrak{f}$ splits into a product of two prime ideals $\mathfrak{p}_2$ and $\mathfrak{p}_2'$ of degree 6 that are exchanged by τ, while $\ell = 3$ splits as $3\mathbb{Z}_{K_z} = \mathfrak{p}_3^2\mathfrak{p}_3'^2$, where $\mathfrak{p}_3$ and $\mathfrak{p}_3'$ are prime ideals of degree 3 which are fixed by τ. Since $\mathfrak{f}$ is coprime to ℓ, we have $S_{\mathfrak{f},\ell,i} = \varnothing$ for $i = 1, 2, 3$, while $S_{\mathfrak{f}}/\langle\tau\rangle = \mathfrak{p}_2$ and $S_\ell/\langle\tau\rangle = \{\mathfrak{p}_3, \mathfrak{p}_3'\}$.

Hence, the conditions involving the conductor alone (step 11 of Algorithm 5.3.17) are trivially satisfied.

Theorem 5.3.15 tells us that there exists β of a precise form such that $L_z = K_z(\sqrt[3]{\alpha})$ with $\alpha = \beta^2\tau(\beta)$. Since K_z has class number equal to 1, using

[Coh0, Algorithm 6.5.10], we compute $\beta_{\mathfrak{p}}$ such that $\mathfrak{p}_2 = \beta_{\mathfrak{p}}\mathbb{Z}_{K_z}$. This $\beta_{\mathfrak{p}}$ is defined only up to units, and already at this stage it is important to reduce it by multiplying it with suitable units.

We now set up the matrix M as explained in step 15 of the algorithm. The only congruences to be satisfied are modulo $\mathfrak{m}_1 = \mathfrak{p}_3^{z(\mathfrak{p}_3,3)-1} = \mathfrak{p}_3^3$ and $\mathfrak{m}_2 = \mathfrak{p}_3'^{\,z(\mathfrak{p}_3',3)-1} = \mathfrak{p}_3'^{\,3}$. Thus, we compute

$$(\mathbb{Z}_K/\mathfrak{p}_3^3)^* \simeq (\mathbb{Z}_K/\mathfrak{p}_3'^{\,3})^* \simeq (\mathbb{Z}/78\mathbb{Z}) \times (\mathbb{Z}/3\mathbb{Z})^5 \ ,$$

and we find

$$M = \begin{pmatrix} 13 & 25 & 35 & 61 & 70 \\ 0 & 2 & 1 & 0 & 0 \\ 0 & 0 & 2 & 2 & 0 \\ 2 & 2 & 2 & 0 & 2 \\ 1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 2 & 0 \\ 65 & 11 & 10 & 75 & 28 \\ 1 & 1 & 1 & 2 & 2 \\ 0 & 1 & 2 & 1 & 2 \\ 1 & 2 & 0 & 0 & 1 \\ 2 & 2 & 1 & 0 & 1 \\ 2 & 2 & 0 & 2 & 1 \end{pmatrix} .$$

As already mentioned, the precise numerical values depend on the chosen generators, so this matrix varies with the implementation.

We then compute the kernel of $\overline{M}$ considered as a matrix with entries in $\mathbb{F}_3$, and we find that this kernel is one-dimensional, generated by $K_1 = (1,-1,-1,1,1)^t$. Since $S_{\mathfrak{f},\ell,2} = \varnothing$, there are no extra discrete logarithms to compute. Since the fifth component (in other words, the exponent $x_{\mathfrak{p}}$) of K_1 is nonzero, $X = K_1$ is the unique suitable solution. Hence, up to Kummer-equivalence, we can take $\beta = \beta_{\mathfrak{p}} w_1 w_4/(w_2 w_3)$.

If we directly use the value of β obtained in this way, we will obtain a polynomial with large coefficients. At this stage, as indicated in step 3 of Subalgorithm 5.3.18, it is essential to reduce β as much as possible.

Using the case $\ell = 3$ of Theorem 5.3.5 and Proposition 5.3.9 given above, we know that the defining polynomial of L/K is given by $X^3 - 3eX - eu$, with $e = \beta\tau(\beta)$ and $u = \beta + \tau(\beta)$. These elements are known to be in K and not only in K_z; hence if we use the expression $z = A_1(B^{-1}(y))$, ordinary linear algebra allows us to express them as polynomials in z instead of y, so we finally obtain the desired defining polynomial.

Using Algorithm 2.1.11, we can now easily compute an absolute defining polynomial for $L/\mathbb{Q}$. The polynomial we will obtain will, however, have rather large coefficients (typically 15 decimal digits) and a very large discriminant (typically 2000 decimal digits). Of course, we want to reduce this polynomial. For this, we use Algorithm 2.4.12, which in the present situation works very well. We finally find the polynomial

$$X^{18} - X^{17} + 3X^{16} + 2X^{15} - X^{14} + 11X^{13} + 3X^{12} + 3X^{11} + 28X^{10}$$
$$- 18X^9 + 47X^8 - 27X^7 + 45X^6 - 23X^5 + 27X^4 - 11X^3 + 9X^2 - 2X + 1 \ .$$

As already remarked, the explicit values for the elements involved in this computation depend on the implementation, so the reader will certainly have different values than ours. Only the final reduced polynomial should be similar.

5.7 Exercises for Chapter 5

1. If $d = \prod_{1 \le i \le k} d_i$ with the d_i pairwise coprime, show that, as claimed in the text, we have
$$(\mathbb{Z}/d\mathbb{Z})g = \bigoplus_{1 \le i \le k} (\mathbb{Z}/d_i\mathbb{Z})g^{d/d_i} \ .$$

2. As suggested in the text, write an algorithm similar to Algorithm 5.1.2 but one that uses only the SNF splitting of the ray class group instead of the p-Sylow splitting.

3. Prove the validity of the exact sequence involving the ℓ-Selmer group given in the text.

4. Let K be a number field of signature (r_1, r_2), let L/K be a quadratic extension, denote by τ the nontrivial K-automorphism of L, and as usual denote by $U(K)$ and $U(L)$ the unit groups of K and L, respectively.
   a) Assume that $\eta \in U(L)$ is such that there exists $a \in K$ and $x \in L$ such that $\eta = ax^3$. Show that there exists $\varepsilon \in U(K)$ and $y \in L$ such that $\eta = \varepsilon y^3$.
   b) Denote by $(U(L)/U(L)^3)[\tau - 2]$ the kernel of the map $\overline{\eta} \mapsto \overline{\tau(\eta)\eta^{-2}}$ from $U(L)/U(L)^3$ to itself. Assume that L is totally complex. Show that the dimension of $(U(L)/U(L)^3)[\tau - 2]$ as an $\mathbb{F}_3$-vector space is equal to $r_2 + 1$.

5. Show that, as claimed in the text, the definition $\lambda = \sum_{0 \le a < d} r_{d-1-a}\tau^a$ does not depend on the choice of the primitive root g.

6. With the notation of Section 5.3.2, let ℓ be an odd prime number and assume that $d = [K_z : K] = 2$, so that $m = (\ell - 1)/2$. Generalizing some of the formulas given in the text, show that if $e = \beta^{1+\tau} = \mathcal{N}_{K_z/K}(\beta)$, then
$$P(X) = \sum_{k=0}^{(\ell-1)/2} (-1)^k \frac{\ell}{\ell - k} \binom{\ell - k}{k} e^k X^{\ell - 2k} - e \operatorname{Tr}_{K_z/K}\left(\beta^{\ell-2}\right) \ .$$

7. Using a computer algebra package, find formulas analogous to those given at the end of Section 5.3.2, for $\ell = 7$ and $d = 2$, $d = 3$, and $d = 6$, and for $\ell = 11$ and $d = 2$, $d = 5$, and $d = 10$ (some of the formulas are completely unwieldy, but this is just an exercise to help the reader understand the use of Theorem 5.3.5 and Proposition 5.3.9).

8. Prove the validity of Algorithm 5.3.11.

9. Let ℓ be a prime number, and let $\Phi_{\ell^r}(X)$ be the ℓ^rth cyclotomic polynomial.
   a) Show that
$$\operatorname{disc}(\Phi_{\ell^r}(X)) = (-1)^{(\ell-1)/2} \ell^{\ell^{r-1}(r\ell - r - 1)} \ .$$

b) More generally, show that $\mathrm{disc}(\Phi_n(X))$ is a divisor of n^n, hence is divisible only by primes dividing n.
c) More precisely, show that if $n > 2$, then

$$\mathrm{disc}(\Phi_n(X)) = (-1)^{\phi(n)/2} \frac{n^{\phi(n)}}{\prod_{p|n} p^{\phi(n)/(p-1)}} .$$

10. Let ℓ be a prime, $n = \ell^r$, L/K a cyclic Kummer extension of degree n with $\zeta_n \in K$, so that $L = K(\theta)$ for some θ such that $\theta^n = \alpha \in \mathbb{Z}_K$. Let $\mathfrak{p}$ be a prime ideal of K unramified in the extension L/K.
    a) Show that $v_{\mathfrak{p}}(\alpha) \equiv 0 \pmod{n}$.
    b) Deduce from this that it is possible to modify α so that $v_{\mathfrak{p}}(\alpha) = 0$.

11. Let $\Phi_n(X)$ be the nth cyclotomic polynomial and K a number field. Show that the factorization of $\Phi_n(X)$ in $K[X]$ is of the form $\Phi_n(X) = \prod_{1 \le i \le d} f_i(X)$, where the polynomials f_i all have the same degree. If we denote by $f_1(X)$ the minimal polynomial of ζ_n in $K[X]$, and G_n is the set of $a \in (\mathbb{Z}/n\mathbb{Z})^*$ such that ζ_n^a is a root of $f_1(X)$ as in Proposition 5.5.1, show that the roots of $f_i(X)$ are the ζ_n^a for a belonging to a fixed coset of $(\mathbb{Z}/n\mathbb{Z})^*$ modulo G_n.

12. With the notation of the proof of Proposition 5.5.2, show that the γ_a must satisfy the cocycle condition $\gamma_{ab} = \tau_a(\gamma_b)\gamma_a^b$.

13. Let $K_z = K(\zeta_n)$, where $n = \ell^r$ is a prime power, and let $m \mid n$ as in Proposition 5.5.2. Finally, let $\beta \in \mathbb{Z}_{K_z}$ be an element known to be an mth power in $\mathbb{Z}_{K_z}$. Write a special-purpose algorithm (in other words, without use general factoring algorithms over number fields) that computes one root of the equation $X^m - \beta = 0$ in K_z. C. Fieker suggests the following method: first choose a suitable prime ideal $\mathfrak{p}$, and factor $X^m - \beta$ modulo $\mathfrak{p}$. Then use a Hensel–Newton method to lift this factorization modulo $\mathfrak{p}^k$ for a sufficiently large k. Finally, reconstruct the factors in K_z by using the integral LLL algorithm ([Coh0, Algorithm 2.6.7]) on a suitable lattice.

14.
    a) Write an algorithm for computing the integer $u_{\mathfrak{p}}$ given by Corollary 5.5.3.
    b) Deduce from this an algorithm for computing the quantities $\mathrm{Art}_{L_z/K}(\mathfrak{a})(\eta)$ that are used in step 10 of Algorithm 5.5.5.

15. Let $K = \mathbb{Q}(\theta_1)$ be a number field. Show that a necessary and sufficient condition for $\zeta_n\theta_1$ to be a primitive element of $K_z = K(\zeta_n)$ is that the ratio of two conjugates of θ_1 should never be an nth root of unity. Give an example where $\zeta_n \notin K$ but where this condition is not satisfied.

16. Let $K = \mathbb{Q}(\theta)$ be the number field defined by a root of the polynomial $T(X) = X^4 - X - 1$. Compute the reduced absolute polynomial of the ray class field over K corresponding to the modulus $\mathfrak{m} = (\theta+5)\mathbb{Z}_K\infty_1\infty_2$, where ∞_1 and ∞_2 are the two real places of K (note that $(\theta+5)\mathbb{Z}_K = \mathfrak{P}_{17}\mathfrak{P}_{37}$ using the notation of the text).

17. Using the techniques of this chapter, write the general defining polynomial for a cyclic quartic extension of a number field. More precisely, in the style of the examples given for Theorem 5.3.5 and Proposition 5.3.9, show that if $\alpha = \theta^4$, we can take $\alpha = \beta^{3+\tau}$, $\tau(\theta) = \theta^3/\beta^2$, $\eta = \theta + \theta^3/\beta^2$, and $P(X) = X^4 - 4eX^2 + e(4e - u^2)$ with $e = \beta^{1+\tau} = \mathcal{N}_{K_z/K}(\beta)$ and $u = \beta + \tau(\beta) = \mathrm{Tr}_{K_z/K}(\beta)$.

18. Write a complete algorithm to compute the ray class field corresponding to a congruence subgroup $(\mathfrak{m}, C)$ using Hecke's theorem (and not the Artin map)

when $Cl_{\mathfrak{m}}(K)/\overline{C}$ is a cyclic group of order 4 (use the preceding exercise to come down to L).

19. Compute defining polynomials for Examples 1 and 2 given in the text, but by using the Artin map instead of Hecke's theorem; in other words, by using Algorithms 5.4.8 and 5.5.5. Compare the efficiency of the computations.

6. Computing Defining Polynomials Using Analytic Methods

In the preceding chapter we saw how to construct Abelian extensions of a number field K corresponding to a given congruence subgroup $(\mathfrak{m}, C)$ by using Kummer theory. The main advantage of this method is that it is completely general (and hence is the only method used for the *proofs* of the existence results in class field theory), but its main drawback is the necessity of working over a larger field $K_z = K(\zeta_n)$.

The aim of the present chapter is to describe quite different methods based on analytic techniques that construct the desired extension L/K if we accept certain restrictions on the base field. When applicable, these methods are much more efficient than the methods using Kummer theory, at least when the necessary roots of unity are not in the base field.

One such method is the use of Stark units and Stark's conjecture, which can be used only if the base field satisfies certain conditions, which include, in particular, all totally real base fields. The other method is the use of complex multiplication techniques, and it is even more restrictive since it applies only when the base field is an imaginary quadratic field. In this case, however, it is much more efficient than the other methods.

The possibility of using analytic techniques for obtaining purely algebraic constructions is certainly one of the fascinating aspects of this part of number theory. Of course, it is nothing new and is already illustrated in the theory of cyclotomic fields by the use of the exponential function. One of the dreams of many number theorists, starting with Kronecker and illustrated by Hilbert's twelfth problem, is to generalize this to arbitrary number fields. We are far from reaching this goal (if it can be reached at all), but the techniques presented in this chapter represent a step in this direction. Generalizing complex multiplication, one can also consider the use of values of Siegel modular functions, but this has not been systematized or put into algorithmic form yet (see [Shim] for complete information on this subject).

6.1 The Use of Stark Units and Stark's Conjecture

One of the most remarkable methods for obtaining defining polynomials for Hilbert and ray class fields is the use of an important conjecture due to Harold Stark. Even though we will be using *conjectural* statements, the final

defining polynomial that we obtain can easily be checked to be correct by showing that it defines an Abelian extension having the correct conductor and congruence subgroup using Algorithm 4.4.6.

The basic reference concerning Stark's conjecture is Tate's book [Tat]. For the material presented below I refer to Xavier Roblot's thesis [Rob1] as well as the papers [Rob2], [Rob3] and [Coh-Rob], and I am indebted to him for the explanation of many technical details.

6.1.1 Stark's Conjecture

To state the conjecture (in fact, a special case of it) we need several analytic definitions.

Let L/K be a finite Abelian extension of number fields of conductor $\mathfrak{f}$, let S be a finite set of places of K containing the places at infinity and all the places of K that ramify in L (in other words, that divide $\mathfrak{f}$), and set $G = \mathrm{Gal}(L/K)$. An *S-unit* of L is an element of L whose valuation is equal to 0 at all prime ideals of L that are not above an element of S (see Definition 7.4.1).

Recall that we denote by Art the Artin reciprocity map from $I_{\mathfrak{f}}$ to G. If $\sigma \in \mathrm{Gal}(L/K)$, we define the *partial Dedekind zeta function* by

$$\zeta_{K,S}(s,\sigma) = \sum_{(\mathfrak{a},S)=1,\ \mathrm{Art}(\mathfrak{a})=\sigma} \mathcal{N}(\mathfrak{a})^{-s} ,$$

where, as indicated, the sum is over all integral ideals of K coprime to S whose image by Art is equal to σ.

A special case of Stark's conjecture is the following.

Conjecture 6.1.1. *Let K be a totally real number field, let L/K be a finite Abelian extension of K with Galois group G, and let S be as above. Assume that there exists a unique real embedding τ of K which is unramified in L/K, so that all the extensions of τ to L are real. Then there exists an S-unit $\varepsilon \in L$ with the following properties.*

(1) *For all real embeddings σ of L extending τ, we have*

$$\sigma(\varepsilon) = e^{-2\zeta'_{K,S}(0,\sigma)} .$$

(2) *For all embeddings σ of L into $\mathbb{C}$ which do not extend τ (hence which are nonreal), we have*

$$|\sigma(\varepsilon)| = 1 .$$

(3) *The extension $L(\sqrt{\varepsilon})/K$ is Abelian.*

From now on, we assume that this conjecture is true.

Stark's conjecture gives us numerical values for all the conjugates $\sigma(\varepsilon)$ for σ above τ, hence for the image by τ of the characteristic polynomial of ε.

We can hope to recover this characteristic polynomial exactly by a suitable algorithm. Hence, if ε generates L/K, this conjecture can allow us to build explicitly the extension L/K.

Before doing this, several problems need to be solved, but the most important is certainly the restriction on the existence of a unique unramified real embedding. To solve this, we use the following proposition.

Proposition 6.1.2. *Let K be a totally real number field distinct from $\mathbb{Q}$, and let L/K be a finite Abelian extension, where L is also a totally real field. Assume that N is a* quadratic *extension of L satisfying the following conditions.*

(1) *N/K is Abelian.*
(2) *There exists a unique real embedding τ of K which is unramified in N/K.*
(3) *Any prime ideal of L that is above a prime ideal of K ramified in L/K is inert or ramified in N/L.*

Let S be the set of infinite places of K together with the places ramified in N/K, and let $\varepsilon \in N$ be an S-unit given by Stark's Conjecture 6.1.1 for the extension N/K. Then ε is in fact a unit and

$$N = K(\varepsilon) = \mathbb{Q}(\varepsilon) \quad \textit{and} \quad L = K(\varepsilon + \varepsilon^{-1}) = \mathbb{Q}(\varepsilon + \varepsilon^{-1}) \ .$$

See [Rob1] for the proof. Note that it is not known whether such an extension N always exists. However, this is not a problem for computational purposes since we can always reduce to cases where N is known to exist (see Section 6.2.1).

6.1.2 Computation of $\zeta'_{K,S}(0,\sigma)$

Let N be as in Proposition 6.1.2, and let $\mathfrak{f}$ be the conductor of the Abelian extension N/K. To be able to use Stark's Conjecture 6.1.1, the main algorithmic problem is the computation of $\zeta'_{K,S}(0,\sigma)$ for $\sigma \in G = \mathrm{Gal}(N/K)$. For this, we introduce the *Hecke L-functions* defined as follows. Let χ be a character from G to $\mathbb{C}^*$. By the Artin reciprocity map from $Cl_{\mathfrak{f}}(K)$ to G, χ can also be viewed as a character on $Cl_{\mathfrak{f}}(K)$, hence on $I_{\mathfrak{f}}(K)$. We define

$$L_S(s,\chi) = \prod_{\mathfrak{p} \nmid \mathfrak{f}} (1 - \chi(\mathfrak{p})\,\mathcal{N}(\mathfrak{p})^{-s})^{-1} = \sum_{(\mathfrak{a},S)=1} \chi(\mathfrak{a})\,\mathcal{N}(\mathfrak{a})^{-s} \ .$$

(Note that $\mathfrak{p} \nmid \mathfrak{f}$ is equivalent to $\mathfrak{p} \notin S$ and $(\mathfrak{a},S) = 1$ is equivalent to $a \in I_{\mathfrak{f}}(K)$.)

These are Abelian L-functions, hence when the character χ is nontrivial, they can be analytically continued in the whole complex plane to holomorphic functions having a functional equation. By the orthogonality formula for characters, if $\hat{G}$ denotes the group of characters of G, we clearly have

$$\zeta_{K,S}(s,\sigma) = \frac{1}{[N:K]} \sum_{\chi \in \hat{G}} \overline{\chi}(\sigma) L_S(s,\chi) \ .$$

Thus, to compute $\zeta'_{K,S}(0,\sigma)$ it is enough to compute $L'_S(0,\chi)$ for all characters χ.

Let τ be the generator of the subgroup $\mathrm{Gal}(N/L)$ of G. Since $\mathrm{Gal}(N/L)$, hence τ, is of order 2, we have $\chi(\tau) = \pm 1$. We will say that χ is an *even* character if $\chi(\tau) = 1$ and is an *odd* character if $\chi(\tau) = -1$. It is easy to show that if χ is an even character we have $L_S(0,\chi) = L'_S(0,\chi) = 0$, while if χ is an odd character we have $L_S(0,\chi) = 0$ but not necessarily $L'_S(0,\chi) = 0$ (see Exercise 1; in fact in our situation, we will always have $L'_S(0,\chi) \neq 0$). Thus, we may assume that χ is odd.

The character χ may not be *primitive* (see Section 3.3.3), in which case although the L-function that we have defined is the one that must be computed, it is not the one with nice properties. Let $\mathfrak{f}(\chi)$ be the conductor of χ (see Definition 3.3.15), so that $\mathfrak{f}(\chi)$ is a divisor of $\mathfrak{f}$. Then χ is induced from a character of $Cl_{\mathfrak{f}(\chi)}(K)$ (which we still denote by χ by abuse of notation). We define

$$L(s,\chi) = \prod_{\mathfrak{p} \nmid \mathfrak{f}(\chi)} (1 - \chi(\mathfrak{p}) \mathcal{N}(\mathfrak{p})^{-s})^{-1} = \sum_{(\mathfrak{a},\mathfrak{f}(\chi))=1} \chi(\mathfrak{a}) \mathcal{N}(\mathfrak{a})^{-s} \ .$$

It is clear that L and L_S differ only by a finite number of Euler factors and, more precisely, that

$$L_S(s,\chi) = L(s,\chi) \prod_{\mathfrak{p} \mid \mathfrak{f}, \mathfrak{p} \nmid \mathfrak{f}(\chi)} (1 - \chi(\mathfrak{p}) \mathcal{N}(\mathfrak{p})^{-s}) \ .$$

Set $m = [K:\mathbb{Q}]$, $C(\chi) = (\pi^{-m} d(K) \mathcal{N}(\mathfrak{f}(\chi)))^{1/2}$, and

$$\Lambda(s,\chi) = C(\chi)^s \Gamma\left(\frac{s+1}{2}\right)^{m-1} \Gamma\left(\frac{s}{2}\right) L(s,\chi) \ .$$

We then have the functional equation (see, for example, [Mart1]):

$$\Lambda(1-s,\chi) = W(\chi) \Lambda(s,\overline{\chi}) \ ,$$

where $W(\chi)$ is a complex number of modulus equal to 1 (called the Artin root number; see [Mart1] for a definition) whose computation will be described in Section 6.2.3. Note that the form of the factors at infinity in the definition of Λ comes from assumption (2) of Proposition 6.1.2 and the fact that χ is an odd character.

From this, it is easy to obtain the following result.

Lemma 6.1.3. *Let χ be an odd character of G as above, and set $A(\chi) = \prod_{\mathfrak{p} \mid \mathfrak{f}, \mathfrak{p} \nmid \mathfrak{f}(\chi)} (1 - \chi(\mathfrak{p}))$. Then*

$$L_S'(0,\chi) = A(\chi)L'(0,\chi) = \frac{A(\chi)W(\chi)}{2}\frac{\Lambda(1,\overline{\chi})}{\pi^{(m-1)/2}} .$$

It remains to compute the value $\Lambda(1,\overline{\chi})$. For this, we use the following formula due to A. F. Lavrik and E. Friedman (see [Lav], [Fri], and Section 10.3).

Theorem 6.1.4. *For $x > 0$, $s \in \mathbb{C}$, and $\delta > \max(\mathrm{Re}(s), 0)$, set*

$$f_m(s,x) = \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} x^z \frac{\Gamma((z+1)/2)^{m-1}\Gamma(z/2)}{z-s}\,dz .$$

Then, if $L(s,\chi) = \sum_{n\geq 1} a_n(\chi)n^{-s}$, we have

$$\Lambda(1,\chi) = \sum_{n\geq 1}\left(a_n(\chi)f_m\left(1,\frac{C(\chi)}{n}\right) + W(\chi)\overline{a_n(\chi)}f_m\left(0,\frac{C(\chi)}{n}\right)\right) .$$

In other words,

$$W(\chi)\Lambda(1,\overline{\chi}) = \sum_{n\geq 1}\left(a_n(\chi)f_m\left(0,\frac{C(\chi)}{n}\right) + W(\chi)\overline{a_n(\chi)}f_m\left(1,\frac{C(\chi)}{n}\right)\right) .$$

The computation of $f_m(s,x)$ and the use of the above series for computing $\Lambda(1,\overline{\chi})$ is explained in detail in [Tol] and summarized in [Rob1], and we refer to both as well as to Section 10.3. We will now restrict to an important special case where this computation is much easier: the case of real quadratic fields.

6.1.3 Real Class Fields of Real Quadratic Fields

In this section, we assume that K is a real quadratic field, and we let $(\mathfrak{m}, C)$ be a congruence subgroup modulo $\mathfrak{m}$, where $\mathfrak{m}$ is assumed to be an integral ideal of K; in other words, we assume that $\mathfrak{m}_\infty = \varnothing$. Let L be the (totally real) ray class field of K defined by $(\mathfrak{m}, C)$, where without loss of generality we may assume that $\mathfrak{m}$ is the conductor of L/K. The field L will be called a *real ray class field* of K. We want to compute a defining polynomial for L/K using Stark's conjecture.

The main simplification is that the function $f_2(s,x)$ occurring in Theorem 6.1.4 can be expressed quite simply.

Proposition 6.1.5. *For $x > 0$ and $s \in \mathbb{C}$, we have*

$$f_2(s,x) = 2\sqrt{\pi}(x/2)^s \int_{2/x}^{\infty} t^{s-1}e^{-t}\,dt .$$

In particular,

$$f_2(1,x) = x\sqrt{\pi}e^{-2/x}$$

and

$$f_2(0,x) = 2\sqrt{\pi}\int_{2/x}^{\infty} \frac{e^{-t}}{t}\,dt = 2\sqrt{\pi}E_1\left(\frac{2}{x}\right) ,$$

where $E_1(x)$ is the exponential integral function.

Proof. By the duplication formula for the gamma function we have

$$\left(\frac{x}{2}\right)^{-s} f_2(s,x) = 2\sqrt{\pi}\frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} \left(\frac{x}{2}\right)^{z-s} \frac{\Gamma(z)}{z-s}\,dz .$$

Replacing x by $2/x$ and differentiating with respect to x, we obtain

$$\frac{d}{dx}\left(x^s f_2\left(s,\frac{2}{x}\right)\right) = -2\sqrt{\pi}\frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} x^{-(z-s+1)}\Gamma(z)\,dz .$$

On the other hand, by the formula giving the inverse Mellin transform, from $\Gamma(z) = \int_0^\infty x^{z-1}e^{-x}\,dx$ we deduce that

$$\frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} x^{-z}\Gamma(z)\,dz = e^{-x} ,$$

from which it follows that

$$\frac{d}{dx}\left(x^s f_2\left(s,\frac{2}{x}\right)\right) = -2\sqrt{\pi}x^{s-1}e^{-x} ,$$

hence that

$$x^s f_2\left(s,\frac{2}{x}\right) = C + 2\sqrt{\pi}\int_x^\infty t^{s-1}e^{-t}\,dt$$

for some constant C possibly depending on s but not on x. Coming back to the definition, we see that

$$x^s f_2\left(s,\frac{2}{x}\right) = \frac{2^z}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} x^{-(z-s)}\frac{\Gamma(z/2)\Gamma((z+1)/2)}{z-s}\,dz ,$$

and since $\delta > \mathrm{Re}(s)$, an easy analytic argument using the fact that the gamma function is bounded in vertical strips shows that the integral tends to 0 when x tends to infinity. It follows that the constant C is equal to 0, giving the first formula of the proposition. The other formulas are immediate consequences. □

The function $E_1(x)$ can be computed as explained in [Coh0, Proposition 5.6.12]. However, for our applications to the computation of $L(1,\overline{\chi})$, it is much more efficient to use the method explained in Exercise 2 (see also [Coh-Rob]). This is due to the fact that we need many values of the form $E_1(cn)$ for some constant c and consecutive integral values of n. This is also true for the computation of $L'(E,1)$ for an elliptic curve of odd rank (see [Coh0, Proposition 7.5.9] and Exercise 24 of Chapter 10).

The second simplification is in the computation of the coefficients $a_n(\chi)$.

Proposition 6.1.6. *Let χ be a character of G and let $n \geq 1$ be an integer. Set $\chi(\mathfrak{p}) = 0$ if $\mathfrak{p}$ divides the conductor of χ. Then $a_1(\chi) = 1$, and for $n > 1$, if $n = p_1^{m_1} \dots p_k^{m_k}$ is the prime factorization of n, then*

$$a_n(\chi) = a_{p_1^{m_1}}(\chi) \dots a_{p_k^{m_k}}(\chi) \ ,$$

where the coefficients $a_{p^m}(\chi)$ are given by one of the following formulas:

(1) *if p is inert, $a_{p^m}(\chi) = 0$ if m is odd and $a_{p^m}(\chi) = \chi(p)^{m/2}$ if m is even;*
(2) *if $p\mathbb{Z}_K = \mathfrak{p}^2$ is ramified, $a_{p^m}(\chi) = \chi(\mathfrak{p})^m$;*
(3) *if $p\mathbb{Z}_K = \mathfrak{p}\mathfrak{p}'$ splits, $a_{p^m}(\chi) = (m+1)\chi(\mathfrak{p})^m$ if $\chi(\mathfrak{p}) = \chi(\mathfrak{p}')$ and otherwise*

$$a_{p^m}(\chi) = \frac{\chi(\mathfrak{p})^{m+1} - \chi(\mathfrak{p}')^{m+1}}{\chi(\mathfrak{p}) - \chi(\mathfrak{p}')} \ .$$

Proof. The first assertion is a translation of the fact that the map $n \mapsto a_n(\chi)$ is multiplicative. The formulas for $n = p^m$ are clear when p is inert or ramified and are easily proved by induction when p is split. □

We will see in the next section how to find the characters χ and how to use them. Using a sieving procedure, we will then simultaneously compute the coefficients $a_n(\chi)$ for all characters χ.

6.2 Algorithms for Real Class Fields of Real Quadratic Fields

We can now explain in detail the construction of the class field L over the real quadratic field K. Although the principle is completely explained above, many technical details and algorithms have to be given before obtaining a complete implementation, and it is the purpose of this section to give them. Thus, this section can be skipped by readers not interested in the technical aspects of the algorithmic implementation of the ideas explained above.

Thus, let $K = \mathbb{Q}(\sqrt{D})$ be a real quadratic field of discriminant D and let C be a congruence subgroup of conductor $\mathfrak{m}$, where $\mathfrak{m}$ is an integral ideal of K, assumed to be the conductor of $(\mathfrak{m}, C)$. We want to compute the real ray class field L corresponding to $(\mathfrak{m}, C)$.

6.2.1 Finding a Suitable Extension N/K

Let $\Gamma = \mathrm{Gal}(L/K)$ be the Galois group of L/K (isomorphic to $Cl_{\mathfrak{m}}(K)/\overline{C}$). If Γ is not cyclic, we can build L/K as the compositum of cyclic extensions L_i/K corresponding to the cyclic components of Γ (see step 1 of Algorithm 6.2.6 below). Thus, we may reduce to the case where Γ is cyclic (this

is, of course, a completely general reduction that is valid in all class field constructions and not only here). We therefore assume from now on that $\Gamma = \mathrm{Gal}(L/K)$ is cyclic.

In this case, it easy to prove that N always exists and furthermore that one can construct such an extension by setting $N = L(\sqrt{\alpha})$, where α is a suitable element of the base field K (more generally, the existence of N is guaranteed under much weaker conditions; see [Rob1] for details).

There are two methods for constructing N. The first one is to choose at random elements $\alpha \in K$ of small norm until an α is found of negative norm. Indeed, if $\mathcal{N}(\alpha) < 0$, we clearly have $\tau_1(\alpha) < 0$ and $\tau_2(\alpha) > 0$ by suitably ordering the embeddings τ_1 and τ_2. As we shall see later, it is essential to choose α such that the conductor of the extension $K(\sqrt{\alpha})$ has the smallest possible norm. Since the norm of this conductor is closely linked to the norm of α, if we choose α of small norm, we can hope for a small conductor. In particular, if the fundamental unit of K has norm -1, we may choose α equal to this fundamental unit. However, if an extension constructed by this method is easily seen to satisfy conditions (1) and (2) of Proposition 6.1.2, it will still be necessary to check condition (3) (see below).

The second method, which is, in general, preferable, is as follows. Using Algorithm 2.3.23, compute the list of all integral ideals of K of norm less than or equal to a given bound B. For each such ideal $\mathfrak{a}$, set $\mathfrak{f}_0 = \mathfrak{m}\mathfrak{a}$ and compute the ray class number modulo $\mathfrak{f} = \mathfrak{f}_0\tau_1$. If this class number is an even multiple of the ray class number $h_{\mathfrak{m},C}(K)$, compute the kernel $\overline{\mathcal{K}}$ of the natural map from $Cl_{\mathfrak{f}}(K)$ to $Cl_{\mathfrak{m}}(K)/\overline{C}$. For every subgroup D of $\mathcal{K}$ of index 2, check whether τ_1 divides the conductor of $(\mathfrak{f}, D)$. If this is the case, the fixed field by $\overline{D}$ of the ray class field of conductor $\mathfrak{f}$ is a quadratic extension of L that satisfies conditions (1) and (2) of Proposition 6.1.2. We then use the method described below to check condition (3). If the bound B is not sufficient to find N, we increase it and continue. Note that an extension N computed by this method is not necessarily of the form $N = L(\sqrt{\alpha})$ with $\alpha \in K$, hence in general this method gives better results than the first method, and we are certain to obtain a modulus of smallest norm. The disadvantage is that the ray class group computations may take some time.

In any case we need to be able to test whether or not a given quadratic extension N/L also satisfies condition (3) of Proposition 6.1.2. Hence, for each prime ideal $\mathfrak{p}$ dividing $\mathfrak{m}$ (in other words, ramified in L/K), we must compute the number $g_{\mathfrak{p}}$ of prime ideals in L above $\mathfrak{p}$, and the number $G_{\mathfrak{p}}$ of prime ideals in N above $\mathfrak{p}$. Condition (3) is satisfied if and only if $G_{\mathfrak{p}} = g_{\mathfrak{p}}$ for each such prime ideal. These numbers are easily computed using Theorem 3.5.3.

We now give a formal write-up of the second method.

Subalgorithm 6.2.1 (Compute Splitting of a Prime Ideal). Let N/K be an Abelian extension defined by $(\mathfrak{f}, D)$, where D is a congruence subgroup of K of conductor $\mathfrak{f}$. Let $\mathfrak{p}$ be a prime ideal of K. This algorithm computes the number

of prime ideals in N above $\mathfrak{p}$. We assume that the ray class group $Cl_{\mathfrak{f}}(K)$ is given by its SNF (A, D_A) with $A = (\overline{\mathfrak{a}_1}, \ldots, \overline{\mathfrak{a}_r})$ and the subgroup $\overline{D}$ by an HNF matrix H.

1. [Get prime to $\mathfrak{p}$-part] Set $v \leftarrow v_{\mathfrak{p}}(\mathfrak{f})$ and set $\mathfrak{n} \leftarrow \mathfrak{p}^{-v}\mathfrak{f}$. Using Algorithm 4.3.1, compute the SNF (B, D_B) of the ray class group $Cl_{\mathfrak{n}}(K)$.
2. [Compute $DP_{\mathfrak{n}}/P_{\mathfrak{n}}$] Using Algorithm 4.3.2, compute the matrix M whose columns are given by the discrete logarithms of the ideals $\mathfrak{a}_i$ on the generators B of the ray class group $Cl_{\mathfrak{n}}(K)$. Let M_2 be the HNF of the matrix $(MH|D_B)$.
3. [Compute $G_{\mathfrak{p}}$] Apply Algorithm 4.1.3 to the system of generators and relations (B, M_2), thus obtaining the SNF (Q, D_Q) of the quotient group $Cl_{\mathfrak{n}}(K)/\overline{DP_{\mathfrak{n}}}$. Let f be the order of the ideal class of $\mathfrak{p}$ in this group, output $\det(M_2)/f$, and terminate the algorithm.

Subalgorithm 6.2.2 (Find Suitable $(\mathfrak{f}, D)$). Let K be a real quadratic field. Denote by τ_1 one of the embeddings of K in $\mathbb{R}$. This algorithm computes a small modulus $\mathfrak{f} = \mathfrak{f}_0\tau_1$ (where $\mathfrak{f}_0$ is a multiple of $\mathfrak{m}$) and a congruence subgroup D modulo $\mathfrak{f}$ such that $\mathfrak{f}$ is the conductor of $(\mathfrak{f}, D)$ and the corresponding field N satisfies the properties of Proposition 6.1.2.

1. [Initialize] Set $h \leftarrow h_{\mathfrak{m},C}(K)$, $B \leftarrow 50$, $b \leftarrow 1$ (we will look at all ideals of norm between $b+1$ and B). Let $\mathfrak{p}_1, \ldots, \mathfrak{p}_t$ be the prime ideals dividing $\mathfrak{m}$. Using Subalgorithm 6.2.1, for $1 \le i \le t$ compute the number g_i of ideals above $\mathfrak{p}_i$ in the Abelian extension defined by the congruence subgroup $(\mathfrak{m}, C)$.
2. [Compute ideal list] Using Algorithm 2.3.23, compute the list $\mathcal{L}$ of integral ideals of absolute norm less than or equal to B (where $\mathcal{L}_n$ contains the list of ideals of norm equal to n), and set $n \leftarrow b$.
3. [Next ideal norm] Set $n \leftarrow n+1$. If $n > B$, set $b \leftarrow B$, set $B \leftarrow 2B$, and go to step 2. Otherwise, set $\mathcal{I} \leftarrow \mathcal{L}_n$, let k be the number of elements of $\mathcal{I}$, and set $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{I}$).
4. [Next ideal] Set $j \leftarrow j+1$. If $j > k$, go to step 3. Otherwise, let $\mathfrak{c}$ be the jth element of $\mathcal{I}$, set $\mathfrak{f}_0 \leftarrow \mathfrak{m}\mathfrak{c}$, and set $\mathfrak{f} \leftarrow \mathfrak{f}_0\tau_1$. Using Algorithm 4.3.1, compute the SNF (B, D_B) of the ray class group $Cl_{\mathfrak{f}}(K)$ and let $D_B = \operatorname{diag}(b_1, \ldots, b_s)$.
5. [Ideal possibly suitable?] If the ray class number (the product of the b_i) is not a multiple of $2h$, go to step 4.
6. [Compute kernel] Using Algorithm 4.1.11, compute the HNF matrix H corresponding to the kernel of the natural map from $Cl_{\mathfrak{f}}(K)$ to $Cl_{\mathfrak{m}}(K)/\overline{C}$. Let S be the SNF of $H^{-1}D_B$ and U and V be unimodular matrices such that $UH^{-1}D_BV = S$ (we can discard the matrix V). Finally, let R be the largest index i such that $2 \mid S_{i,i}$, and set $r \leftarrow 0$.
7. [Next row] Set $r \leftarrow r+1$. If $r > R$, go to step 4. Otherwise, set $c \leftarrow -1$.

8. [Next subgroup] Set $c \leftarrow c+1$. If $c \geq 2^{R-r}$, go to step 7. Otherwise, let $H_Q = (h_{i,j})$ be the $s \times s$ upper-triangular matrix built as follows. The diagonal entries $h_{i,i}$ are all equal to 1 except for $h_{r,r}$, which is equal to 2. The off-diagonal entries $h_{i,j}$ are all equal to 0 except perhaps those where $i = r$ and $j > r$, for which $h_{i,j}$ is equal to 1 if and only if bit number $j-r-1$ of c is equal to 1.
9. [Is subgroup suitable?] Let H_D be the HNF of the matrix $(HU^{-1}H_Q|D_B)$. Using Algorithm 4.4.2, test if the conductor of the subgroup of $Cl_{\mathfrak{f}}(K)$ corresponding to the matrix H_D is indeed equal to $\mathfrak{f}$. If this is not the case, then go to step 8. Otherwise, using Subalgorithm 6.2.1, for $1 \leq i \leq t$ compute the number G_i of ideals above $\mathfrak{p}_i$ in the field corresponding to $(\mathfrak{f}, H_D)$. If for all i, $g_i = G_i$, output the conductor $\mathfrak{f}$ and the congruence subgroup D of $I_{\mathfrak{f}}$ (represented by the matrix H_D), and terminate the algorithm. Otherwise, go to step 8.

Proof. We want to find $(\mathfrak{f}, D)$ such that the corresponding extension N/K satisfies the properties of Proposition 6.1.2 and such that $\mathfrak{f}$ has minimum norm. The field N must be a quadratic extension of L, so by class field theory this means that the congruence subgroup modulo $\mathfrak{f}$ corresponding to the extension N/K must have index 2 in the congruence subgroup modulo $\mathfrak{f}$ corresponding to the extension L/K. But this group is also the kernel of the canonical surjection from $Cl_{\mathfrak{f}}(K)$ to $Cl_{\mathfrak{m}}(K)/\overline{C}$ that is computed in step 6 of the algorithm.

Steps 7 and 8 are the exact algorithmic translation of Proposition 4.1.19 and give a systematic list of congruence subgroups of index 2 using the binary bits of c. They are, in fact, a summary of Algorithm 4.1.20 for $\ell = 2$.

Finally, step 9 checks whether or not $\mathfrak{f}$ is the conductor and if condition (3) is also satisfied. If $\mathfrak{f}$ is indeed the conductor, then τ_1 is ramified in N/L and, since we have taken the moduli in order of increasing norm, this condition is also sufficient for $\mathfrak{f}$ to be minimal, since otherwise we would have found a suitable $\mathfrak{f}$ earlier. □

6.2.2 Computing the Character Values

We now assume that a suitable extension N/K of conductor $\mathfrak{f}$ corresponding to a congruence subgroup $(\mathfrak{f}, D)$ has been chosen.

Let χ be a character of $G = \mathrm{Gal}(N/K)$ and let τ be the unique nontrivial element of $\mathrm{Gal}(N/L)$. Recall that χ is odd if $\chi(\tau) = -1$; otherwise χ is even. Even characters are the characters of G induced by the characters of $\mathrm{Gal}(L/K)$. On the contrary, it is easy to show that odd characters are exactly the characters whose conductor has a nontrivial infinite part. Since $L'_S(0,\chi) = 0$ for an even character, we will only be concerned by odd characters.

Let $Cl_{\mathfrak{f}}(K) = \bigoplus_{1 \leq i \leq g} (\mathbb{Z}/e_i\mathbb{Z})h_i$ be the HNF of $Cl_{\mathfrak{f}}(K)$, and let H_D be the HNF matrix representing $\overline{D}$ as a subgroup of $Cl_{\mathfrak{f}}(K)$ (as obtained, for example, in Subalgorithm 6.2.2). Let (Q, D_Q) denote the SNF of the

quotient group $Cl_{\mathfrak{f}}(K)/\overline{D}$ computed using Algorithm 4.1.7 (thus (Q, D_Q) is a presentation of the group G). Recall from Section 4.4.3 that if $Q = (g_1, \ldots, g_k)$ and $D_Q = \mathrm{diag}(d_1, \ldots, d_k)$, then a character χ of G is defined by a vector $(a_1, \ldots, a_k)$, where $a_i \in \mathbb{Z}/d_i\mathbb{Z}$, through the formula

$$\chi\Big(\prod_{1 \le i \le k} g_i^{x_i}\Big) = \zeta^{\sum_i (d_1/d_i) a_i x_i} \ ,$$

where $\zeta = \exp(2i\pi/d_1)$.

Thus, we can compute all the characters of G by using the above representation, and we use Algorithm 4.4.7 to compute the conductor of these characters and hence to check which are odd characters.

Let χ be an odd character of G of conductor $\mathfrak{f}(\chi)$ and let $\mathfrak{a}$ be a fractional ideal of K coprime to $\mathfrak{f}(\chi)$. We must be able to compute $\chi(\mathfrak{a})$ even for $\mathfrak{a}$ coprime to $\mathfrak{f}(\chi)$ but not necessarily to $\mathfrak{f}$. For this, we can multiply $\mathfrak{a}$ by an element of $P_{\mathfrak{f}(\chi)}$ (which does not change the value of $\chi(\mathfrak{a})$) so that it becomes coprime to $\mathfrak{f}$. That this is possible follows from Lemma 3.3.1 and can be made algorithmic if desired.

An equivalent and usually preferable method is to compute exponents x_i such that $\chi(\mathfrak{a}) = \chi(g)$, where g is the element of G defined by $g = \prod_{i=1}^{k} g_i^{x_i}$. For this, we proceed as follows. First we compute the canonical surjection from $Cl_{\mathfrak{f}(\chi)}(K)$ to $Cl_{\mathfrak{f}}(K)/\overline{D}$ given by a matrix M, then we use Algorithm 4.3.2 to compute exponents y_i such that $\overline{\mathfrak{a}} = \prod_{1 \le i \le g} h_i^{y_i}$ in $Cl_{\mathfrak{f}(\chi)}(K)$. The exponents x_i are then given by the first k entries of the column vector MY, where Y is the column vector of the y_i.

6.2.3 Computation of $W(\chi)$

The Artin root number $W(\chi)$ is a mathematical quantity that arises in many fields of algebraic number theory, hence it is worthwhile to give an algorithm to compute it in a more general situation. In this section, we will assume only that χ is an Abelian character defined over some ray class group $Cl_{\mathfrak{f}}(K)$ of a number field K, where $\mathfrak{f}$ is the conductor of χ. We will essentially follow the method given in [Dum-Tan] with a slightly different computational approach.

We will say that an algebraic number $\theta \in K$ is $\mathfrak{f}_\infty$-positive if $\sigma(\theta) > 0$ for all $\sigma \in \mathfrak{f}_\infty$. The following result, due to Landau, gives an explicit formula for $W(\chi)$.

Proposition 6.2.3. *Let $\mathfrak{D}$ denote the different of $K/\mathbb{Q}$. Choose an $\mathfrak{f}_\infty$-positive element $\lambda \in \mathfrak{D}\mathfrak{f}_0$ such that the integral ideal $\mathfrak{g} = \lambda\mathfrak{D}^{-1}\mathfrak{f}_0^{-1}$ is coprime to $\mathfrak{f}_0$, and choose an $\mathfrak{f}_\infty$-positive element $\mu \in \mathfrak{g}$ such that the integral ideal $\mathfrak{h} = \mu\mathfrak{g}^{-1}$ is coprime to $\mathfrak{f}_0$. Define the following Gauss sum:*

$$G(\chi) = \chi(\mathfrak{h}) \sum_{\beta} \chi(\beta) e^{2i\pi \,\mathrm{Tr}(\beta\mu/\lambda)} \ ,$$

where Tr *denotes the absolute trace of* $K/\mathbb{Q}$ *and* β *runs through a complete residue system of* $(\mathbb{Z}_K/\mathfrak{f}_0)^*$ *chosen to be* $\mathfrak{f}_\infty$*-positive elements. Then*

$$W(\chi) = \frac{(-i)^{f_\infty} G(\chi)}{\sqrt{N\mathfrak{f}_0}}.$$

This yields the following algorithm.

Algorithm 6.2.4 (Computation of $W(\chi)$). Let χ be a character of conductor $\mathfrak{f}$. This algorithm computes the Artin root number $W(\chi)$ attached to this character.

1. [Compute λ] Using Proposition 1.3.8, compute an element $\lambda' \in \mathfrak{D}\mathfrak{f}_0$ such that $v_\mathfrak{p}(\lambda') = v_\mathfrak{p}(\mathfrak{D}\mathfrak{f}_0)$ for all prime ideals $\mathfrak{p}$ dividing $\mathfrak{f}_0$. Then, using Algorithm 4.2.20, compute an $\mathfrak{f}_\infty$-positive element λ such that $\lambda \equiv \lambda' \pmod{\mathfrak{f}_0}$, and set $\mathfrak{g} \leftarrow \lambda\mathfrak{D}^{-1}\mathfrak{f}_0^{-1}$.
2. [Compute μ] Using Algorithm 1.3.2, compute two elements $\mu' \in \mathfrak{g}$ and $\nu \in \mathfrak{f}_0$ such that $\mu' + \nu = 1$ (note that $\mathfrak{g}$ and $\mathfrak{f}_0$ are coprime by construction). Then, using Algorithm 4.2.20, compute an $\mathfrak{f}_\infty$-positive element μ such that $\mu \equiv \mu' \pmod{\mathfrak{f}_0}$, and set $\mathfrak{h} \leftarrow \mu\mathfrak{g}^{-1}$.
3. [Initialize $(\mathbb{Z}_K/\mathfrak{f}_0)^*$] Let (A, D_A) be the SNF of the group $(\mathbb{Z}_K/\mathfrak{f}_0)^*$ as computed by Algorithm 4.2.21. Write $A = (\alpha_1, \ldots, \alpha_r)$, $D_A = \mathrm{diag}(a_1, \ldots, a_r)$, and let m denote the cardinality of $(\mathbb{Z}_K/\mathfrak{f}_0)^*$ (hence m is the product of the a_i). Set $G \leftarrow 0$ and $c \leftarrow 0$.
4. [Compute Gauss sum] Set $d \leftarrow c$ and $\beta' \leftarrow 1$. For $i = 1, \ldots, i = r$ (in this order), set $e \leftarrow d \mod a_i$, $\beta' \leftarrow \beta'\alpha_i^e$ (reduced modulo $\mathfrak{f}_0$ using Algorithm 1.4.13), and $d \leftarrow (d-e)/a_i$. Then using Algorithm 4.2.20, compute an $\mathfrak{f}_\infty$-positive element β such that $\beta \equiv \beta' \pmod{\mathfrak{f}_0}$. Set $G \leftarrow G + \chi(\beta)e^{2i\pi\,\mathrm{Tr}(\beta\mu/\lambda)}$ and $c \leftarrow c+1$. If $c < m$, go to step 4.
5. [Output result] Set $W \leftarrow (-i)^{|\mathfrak{f}_\infty|}\chi(\mathfrak{h})G/\sqrt{N\mathfrak{f}_0}$. Output W and terminate the algorithm.

Remark. It is possible to improve this algorithm in several ways. First, we can compute at the beginning the complex values $\chi(\alpha_i)$ and use the multiplicativity of χ to obtain the value of $\chi(\beta)$ in step 4. Another improvement would be to choose the generators α_i to be $\mathfrak{f}_\infty$-positive, since this would avoid correcting the sign of β'. On the other hand, the powers α_i^e in step 4 must also be reduced modulo $\mathfrak{f}_0$ using Algorithm 1.4.13 since we are only interested in the class modulo $\mathfrak{f}_0$. But when we do so, we will generally not obtain an $\mathfrak{f}_\infty$-positive element, so these last two improvements are incompatible. Numerical experiments show, however, that the latter improvement is far more important than the former. Thus, it is preferable not to choose $\mathfrak{f}_\infty$-positive generators α_i but to reduce all the powering operations modulo $\mathfrak{f}_0$ and to adjust the sign of the result only at the end.

6.2.4 Recognizing an Element of $\mathbb{Z}_K$

Let ε be the unit given by Stark's conjecture, let $\alpha = \varepsilon + \varepsilon^{-1} \in L$, and let $P(X) = \sum_{0 \le i \le h} \beta_i X^i$ be the characteristic and minimal polynomial of α over K. Stark's conjecture gives us two items of information on the coefficients $\beta_i \in \mathbb{Z}_K$. The first one is a good numerical approximation to $\tau_2(\beta_i)$. The second one is an upper bound for $|\tau_1(\beta_i)|$. Indeed, for all σ above τ_1 we know that $|\sigma(\varepsilon)| = 1$; hence $|\sigma(\alpha)| \le 2$, so

$$|\tau_1(\beta_i)| \le \binom{h}{i} 2^i \ .$$

Thus we have the following problem to solve: find $\gamma \in \mathbb{Z}_K$ knowing a good numerical approximation β to $\tau_2(\gamma)$, say $|\tau_2(\gamma) - \beta| \le \varepsilon$ for some small $\varepsilon > 0$, and an upper bound of the form $|\tau_1(\gamma)| \le B$.

To solve this problem, there are essentially two methods. The first one, which we can call the naive method although it gives reasonably good results, is to perform an exhaustive search in the following way: let $(1, \omega)$ be a $\mathbb{Z}$-basis of $\mathbb{Z}_K$ and let $\omega_i = \tau_i(\omega)$ for $i = 1$ and 2. If $\gamma = a + b\omega$, then we want $|a + b\omega_2 - \beta| \le \varepsilon$ and $|a + b\omega_1| \le B$. Combining these two conditions easily gives an upper bound for b, and for each b the first inequality gives at most one possible value of a which is tested. The details are left to the reader (Exercise 3). This method has the advantage of being simple, but it needs $O(B/(\omega_2 - \omega_1))$ steps. Hence it becomes impractical when h, hence B, is large.

A second method is to use the LLL algorithm in a well-targeted manner. Consider the lattice $\Lambda = \mathbb{Z}^3$, and the positive definite quadratic form q defined on this lattice by

$$q(x, y, z) = \left(\frac{B}{\varepsilon}\right)^2 (x + y\omega_2 - \beta z)^2 + (x + y\omega_1)^2 + B^2 z^2 \ .$$

We have the following proposition.

Proposition 6.2.5. *Keep the above notation. If $\gamma = a + b\omega \in \mathbb{Z}_K$ satisfies $|\tau_2(\gamma) - \beta| \le \varepsilon$ and $|\tau_1(\gamma)| \le B$, then $q(a, b, 1) \le 3B^2$. Conversely, assume in addition that $\varepsilon < 1/(3(B+1)(\sqrt{D}+1))$. Then if $q(x, y, z) \le 3B^2$ and $(x, y, z) \ne (0, 0, 0)$, we have $z = \pm 1$ and $\gamma = z(x + y\omega)$ is a solution to the slightly weaker problem $|\tau_2(\gamma) - \beta| \le \varepsilon\sqrt{3}$ and $|\tau_1(\gamma)| \le B\sqrt{3}$.*

Proof. The first statement is clear. Conversely, assume that $q(x, y, z) \le 3B^2$ and that $(x, y, z) \ne (0, 0, 0)$. It is clear that $|z| \le 1$, and Exercise 4 (which is an excellent exercise on the properties of continued fractions) shows that if we assume the given inequality for ε, then $z \ne 0$, so $z = \pm 1$ and the rest of the proposition follows. □

Thanks to this proposition, we see that we must find the nonzero solutions to $q(x,y,z) \leq 3B^2$. This can be done using the Fincke–Pohst algorithm ([Coh0, Algorithm 2.7.7]). Among the solutions found, we could, if desired, keep only those satisfying our two inequalities. Note, however, that using ε or $\varepsilon\sqrt{3}$ is in practice equivalent and, moreover, that it would be very surprising if such a good approximation could be obtained with $B < |x + y\omega_1| \leq B\sqrt{3}$, since the chosen bound B is in fact very pessimistic, so finding nonzero solutions to $q(x,y,z) \leq 3B^2$ can in practice be considered to be almost equivalent to our initial problem.

We leave the details of this algorithm to the reader (Exercise 5).

6.2.5 Sketch of the Complete Algorithm

We end this chapter by giving an overview of the complete algorithm used to compute a real ray class field of a real quadratic field.

Algorithm 6.2.6 (Computation of a Real Ray Class Field Using Stark Units). Let K be a real quadratic field and let $(\mathfrak{m}, C)$ be a congruence subgroup of conductor $\mathfrak{m}$, where $\mathfrak{m}$ is an integral ideal of K. Let L denote the class field corresponding to $(\mathfrak{m}, C)$. This algorithm computes a defining polynomial for the field extension L/K by using Stark's conjecture.

1. [Split the Galois group] Using Algorithm 5.1.2, compute s congruence subgroups $(\mathfrak{m}_j, C_j)$ such that the compositum of the Abelian extensions L_j/K corresponding to these congruence subgroups is equal to the desired extension L/K and such that L_j/K is cyclic (see Remark (2) below).
2. [Compute class field] For $1 \leq j \leq s$, compute the field L_j using Subalgorithm 6.2.7 and set L to be the compositum of all the fields L_j.
3. [Check the result] Using Algorithm 4.4.5, compute the norm group of L and check whether it is equal to C. If yes, output L and a message saying that L is the class field corresponding to $(\mathfrak{m}, C)$ under GRH (recall that Algorithm 4.4.5 assumes GRH); otherwise, output a message saying that the algorithm fails. Terminate the algorithm.

Subalgorithm 6.2.7 (Compute the Cyclic Field L/K). This algorithm computes the real ray class field L over the real quadratic field K assuming that L/K is cyclic (note that this algorithm usually works even if L/K is not cyclic; this is just a sufficient hypothesis to make sure that the field N exists; see Section 6.2.1).

1. [Find N] Using Algorithm 6.2.2, find $(\mathfrak{f}, D)$ such that $\mathfrak{f}$ is the conductor of the congruence subgroup D and the corresponding field N satisfies the properties of Proposition 6.1.2.
2. [Compute $L'_S(0,\chi)$] Set $G \leftarrow Cl_{\mathfrak{f}}(K)/\overline{D}$ and let S be the set of prime ideals dividing $\mathfrak{f}$. Using the results of Section 6.2.2, compute the odd characters

χ_i. For each such character, compute the Artin root number $W(\chi_i)$ using Algorithm 6.2.4, and use the results of Section 6.1.3 to obtain accurate values of $L'_S(0,\chi_i)$.

3. [Compute approximation to $P(X)$] Let σ_j be a system of representatives of $G/\langle\tau\rangle$. Using the methods described in the preceding sections, for each j compute

$$\zeta'_S(0,\sigma_j) \leftarrow \frac{1}{[N:K]}\sum_i \overline{\chi_i}(\sigma_j)L'_S(0,\chi_i) \ ,$$

then set $\varepsilon_j \leftarrow e^{-2\zeta'_S(0,\sigma_j)}$ and $\alpha_j \leftarrow \varepsilon_j + \varepsilon_j^{-1}$. Finally, set $P(X) \leftarrow \prod_j (X - \alpha_j)$.

4. [Round to algebraic and terminate] Write $P(X) = \sum_{0\le i\le h} \beta_i X^i$, where the β_i are real approximations to algebraic integers. For each i, use the algorithm mentioned in Exercise 3 or in Exercise 5 to compute $\gamma_i \in \mathbb{Z}_K$ such that $\tau_2(\gamma_i)$ closely approximates β_i and such that $\tau_1(\gamma_i)$ is not too large (see Section 6.2.4 above). If that algorithm fails for some i, the accuracy used in the present algorithm was not sufficient, so terminate the algorithm with an error message, or start it again using a higher accuracy. Otherwise, set $P(X) \leftarrow \sum_{0\le i\le h} \gamma_i X^i$, output $P(X)$, and terminate the subalgorithm.

Remarks

(1) Assuming that Stark's conjecture and the GRH are both correct, failure of Algorithm 6.2.6 can happen only if the computations (essentially that of the $\zeta'_{K,S}(0,\sigma_j)$) have not been done with sufficient accuracy. In that case we must start again with a higher accuracy.

(2) Contrary to the case of Kummer theory, where it is essential to split the construction of L into a number of much simpler constructions, this is not so useful here and, in fact, is usually a bad idea. Indeed, practice shows that the algorithm is fastest with no splitting at all, omitting step 1 entirely. The reason we do split at least into cyclic extensions (not necessarily of prime power degree) is that otherwise the existence of a suitable quadratic extension N/L is not guaranteed. Thus, a good strategy is to directly apply Subalgorithm 6.2.7 and, if Algorithm 6.2.2 does not succeed in finding a suitable N in a reasonable amount of time, to split the problem into smaller (for example, cyclic) subproblems.

6.2.6 The Special Case of Hilbert Class Fields

In the special case of Hilbert class fields, we can modify the defining polynomial $P(X)$ so that in fact $P(X) \in \mathbb{Z}[X]$. We begin with the following theorem (see [Cor-Ros]).

Theorem 6.2.8. *Let K be a number field such that $K/\mathbb{Q}$ is a* cyclic *extension, and let $L = K(1)$ be its Hilbert class field. There exists a number field L_K (called a* splitting field *for L) such that $K \cap L_K = \mathbb{Q}$ and such that*

$L = KL_K$. *In other words, there exists a relative defining polynomial for the Hilbert class field that belongs to* $\mathbb{Q}[X]$.

In particular, to give the Hilbert class field of a quadratic field, it is enough to give the field extension $L_K/\mathbb{Q}$. Note that this theorem is trivially false if $K/\mathbb{Q}$ is noncyclic (see Exercise 6).

In the case of imaginary quadratic fields, complex multiplication methods such as the standard use of the j-function or Schertz's improved functions (see Theorem 6.3.7 below) directly give the field extension $L_K/\mathbb{Q}$, since the defining polynomial of $K(1)/K$ belongs to $\mathbb{Z}[X]$.

In the real quadratic case, however, Algorithm 6.2.6 really gives a defining polynomial $P_2(X) \in K[X]$ and not in $\mathbb{Z}[X]$ in general.

To find a defining polynomial in $\mathbb{Z}[X]$, we use the following simple-minded yet efficient algorithm.

Algorithm 6.2.9 (Computation of $P(X) \in \mathbb{Z}[X]$ for Hilbert Class Fields). Given an irreducible polynomial $P_2(X) \in K[X]$ defining the Hilbert class field $K(1)$ of a real quadratic field K of discriminant D, this algorithm computes a polynomial $P(X) \in \mathbb{Z}[X]$ that is irreducible in $K[X]$, a root of which also defines $K(1)/K$.

1. [Compute absolute defining polynomial] Using Algorithm 2.1.11, compute an absolute defining polynomial $Q(X)$ for $K(1)/\mathbb{Q}$, set $h \leftarrow \deg(Q)/2 = [K(1) : K] = |Cl(K)|$, and set $d(L) \leftarrow D^h$.
2. [First use Polred] Using a polynomial reduction algorithm ([Coh0, Algorithm 4.4.11]) on the polynomial Q, find a list $\mathcal{L}_1$ of polynomials defining some subfields of $K(1)/\mathbb{Q}$, and set $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{L}_1$).
3. [Next element of $\mathcal{L}_1$] Set $j \leftarrow j + 1$. If $j > |\mathcal{L}_1|$, go to step 5. Otherwise, let $P(X)$ be the jth polynomial in the list $\mathcal{L}_1$.
4. [Test if suitable] Using Subalgorithm 6.2.10 below, test if the polynomial $P(X)$ is suitable. If it is, output $P(X)$ and terminate the algorithm. Otherwise go to step 3.
5. [Find subfields] Using [Klu], find a list $\mathcal{L}_2$ of polynomials defining all the subfields of $K(1)$ of degree h (see below), and set $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{L}_2$).
6. [Next element of $\mathcal{L}_2$] Set $j \leftarrow j + 1$. If $j > |\mathcal{L}_2|$, output an error message saying that there is a bug in the algorithm and terminate. Otherwise, let $P(X)$ be the jth polynomial in the list $\mathcal{L}_2$.
7. [Test if suitable] Using Subalgorithm 6.2.10 below, test if the polynomial $P(X)$ is suitable. If it is, output $P(X)$ and terminate the algorithm. Otherwise, go to step 6.

Subalgorithm 6.2.10 (Test if $P(X)$ is Suitable). Given a monic polynomial $P(X) \in \mathbb{Z}[X]$ irreducible in $\mathbb{Z}[X]$, this subalgorithm tests whether P defines a

number field L_K over $\mathbb{Q}$ such that $L_K \cap K = \mathbb{Q}$ and $KL_K = K(1)$. We use all the quantities computed in the main algorithm.

1. [Easy case I] If $\deg(P) \neq h$, P is not suitable and terminate.
2. [Easy case II] If h is odd, $P(X)$ is suitable and terminate. Otherwise, compute the discriminant $d(L_K)$ of the number field defined by $P(X)$. If $d(L) \neq d(L_K)^2$, $P(X)$ is suitable and terminate.
3. [Harder case] (Here $P(X)$ is of degree h and $d(L_K)^2 = d_L$.) Using [Coh0, Algorithm 3.6.4], check if $P(X)$ is irreducible in $K[X]$. If it is, $P(X)$ is suitable; otherwise, $P(X)$ is not suitable. Terminate the subalgorithm.

Proof. (1). Main algorithm. We must find a subfield L_K of $L = K(1)$ such that $L_K \cap \mathbb{Q} = \mathbb{Q}$ and $KL_K = L$; hence, in particular, $[L_K : \mathbb{Q}] = [L : K] = h$. The Polred algorithm is a fast algorithm that gives h subfields of L (including the trivial subfields $\mathbb{Q}$ and L), hence we may hope to find the desired subfield L_K among those given by Polred. This is the reason for which we begin by using it.

If none of the number fields thus found is suitable, we have to use a more systematic procedure. One such procedure is to generate more subfields than h in the Polred algorithm (see Exercise 7), but this may be very costly. Another probably preferable procedure is to use an algorithm for finding subfields of given degree. We have not given any such algorithm in this book or in [Coh0] since it is quite technical, so I refer to [Klu] or to [Klu-Poh] for a detailed description.

(2). Subalgorithm. As above, call L_K the number field defined by a root of $P(X)$ (note that the polynomials given by Polred and by the subfield algorithm are all irreducible over $\mathbb{Q}$, but not necessarily over K). Clearly a necessary condition for $P(X)$ to be suitable is that $\deg(P) = h$. On the other hand, it is clear that if $\deg(P) = h$, then $P(X)$ is suitable if and only if $K \not\subset L_K$. If h is odd, this is trivially the case. If h is even and $K \subset L_K$, then L/L_K is a subextension of L/K, hence is unramified (since $L = K(1)$), so $d(L) = d(L_K)^2$. Hence if this condition is not satisfied, $P(X)$ is suitable.

Finally, if none of these simple criteria suffices to determine whether or not $P(X)$ is suitable, we factor $P(X)$ in $K[X]$, and clearly $K \not\subset L_K$ if and only if $P(X)$ is irreducible in $K[X]$. □

6.3 The Use of Complex Multiplication

Warning. Due to an unfortunate oversight, it is necessary to exchange ω_1 and ω_2 almost everywhere in Chapter 7 of the first three printings of [Coh0]. For details, see the errata sheet at the URL

`ftp://megrez.math.u-bordeaux.fr/pub/cohenbook/errata4.tex`

All references to Chapter 7 assume that these corrections have been made.

6.3.1 Introduction

I follow here quite closely a number of papers written by R. Schertz ([Sch1], [Sch2], [Sch3], [Sch4], [Sch5]), and I thank the Kant group for code and references.

The basic principle of complex multiplication is as follows. Let f be a function defined on the upper half-plane $\mathcal{H}$ and which is modular for some congruence subgroup of $\mathrm{PSL}_2(\mathbb{Z})$ (see [Lan2] for definitions). Then, up to suitable normalizations, if τ is a quadratic number in $\mathcal{H}$, we can expect $f(\tau)$ to be an algebraic number with interesting arithmetic properties. This is of course very vague, but it will be made completely precise for the examples that we will need for Hilbert and ray class group computations.

Let K be an imaginary quadratic field of discriminant $D < 0$, and let $(\mathfrak{m}, C)$ be a congruence subgroup of K. As in the rest of this chapter, we want to give an explicit defining polynomial for the Abelian extension L of K corresponding to this congruence subgroup by Takagi's theorem, but this time by using algebraic values of modular forms and functions on elements of K.

Denote by s the canonical surjection from $Cl_{\mathfrak{m}}(K)$ to $Cl(K)$, and recall that we write $\overline{C} = C/P_{\mathfrak{m}}$. Denote by Z the kernel of the restriction of the map s to $\overline{C}$, and by T the kernel of the natural map from $Cl_{\mathfrak{m}}(K)/\overline{C}$ to $Cl(K)/s(\overline{C})$. We clearly have the following commutative diagram of exact sequences.

$$\begin{array}{ccccccccc}
 & & 1 & & 1 & & 1 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
1 & \longrightarrow & Z & \longrightarrow & \overline{C} & \longrightarrow & s(\overline{C}) & \longrightarrow & 1 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
1 & \longrightarrow & \dfrac{(\mathbb{Z}_K/\mathfrak{m})^*}{\mathrm{Im}(U(K))} & \longrightarrow & Cl_{\mathfrak{m}}(K) & \longrightarrow & Cl(K) & \longrightarrow & 1 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
1 & \longrightarrow & T & \longrightarrow & \dfrac{Cl_{\mathfrak{m}}(K)}{\overline{C}} & \longrightarrow & \dfrac{Cl(K)}{s(\overline{C})} & \longrightarrow & 1 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 1 & & 1 & & 1 & &
\end{array}$$

We will proceed in two almost independent steps. We will first construct the *unramified* Abelian extension $K(1)^{s(\overline{C})}$ of K (hence a subextension of the Hilbert class field $K(1)$ of K) corresponding to the subgroup $s(\overline{C})$ of $Cl(K)$. We will then construct the desired extension L as an extension of $K(1)^{s(\overline{C})}$

such that $\mathrm{Gal}(L/K(1)^{s(\overline{C})}) \simeq T$ by the Artin reciprocity map. Thus, the above commutative diagram corresponds to the following diagram of field extensions.

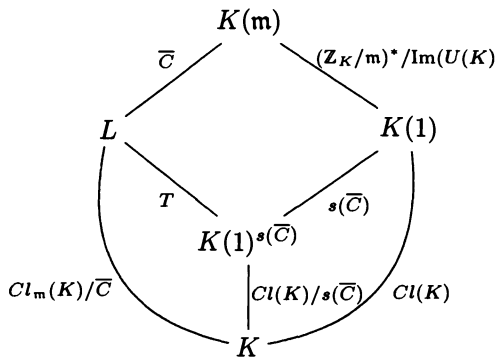

For ease of exposition, however, we will construct $K(\mathfrak{m})$ as an extension of $K(1)$, ignoring the congruence subgroup C, and take it into account only at the very end. In an actual implementation, this should be done differently as explained above.

6.3.2 Construction of Unramified Abelian Extensions

In [Coh0, Section 7.6], and in particular in Algorithm 7.6.1, we saw how to use values of the elliptic function $j(\tau)$ to construct the Hilbert class field of K. A similar construction leads to a construction of subextensions $K(1)^C$, where C is a congruence subgroup modulo $\mathbb{Z}_K$. Note that from now on, for simplicity of notation we write $K(1)^C$ instead of $K(1)^{\mathrm{Art}(C)}$.

The fundamental first step is to know precisely the action of the Galois group of $K(1)/K$ on the values of $j(\tau)$. Recall that by [Coh0, Theorem 5.2.4], we can identify quadratic numbers τ modulo the additive action of $\mathbb{Z}$ with equivalence classes of fractional ideals modulo the multiplicative action of $\mathbb{Q}^*$, or with equivalence classes of *positive definite* quadratic forms of discriminant D modulo the action of Γ_∞. Since $j(\tau)$ is $\mathbb{Z}$-periodic, it is thus permissible to write $j(\mathfrak{a})$ for an ideal $\mathfrak{a}$. Indeed, if (ω_1, ω_2) is a $\mathbb{Z}$-basis of $\mathfrak{a}$ oriented in such a way that $\mathrm{Im}(\omega_1/\omega_2) > 0$, then by definition $j(\mathfrak{a}) = j(\omega_1/\omega_2)$, and since j is a modular invariant, this is independent of the choice of basis. In addition, since for any $\alpha \in K^*$, $(\alpha\omega_1, \alpha\omega_2)$ is an oriented basis of $\alpha\mathfrak{a}$, $j(\mathfrak{a})$ clearly only depends on the ideal class of $\mathfrak{a}$.

The following proposition, known in a more general setting as Shimura's reciprocity law, gives the action of the Galois group of $K(1)/K$ on the values of j.

Proposition 6.3.1. *Let $\mathfrak{a}$ and $\mathfrak{c}$ be fractional ideals of K, and let* $\mathrm{Art}(\mathfrak{c})$ *be the element of* $\mathrm{Gal}(K(1)/K)$ *corresponding to the ideal $\mathfrak{c}$ by the Artin reciprocity map (since $K(1)/K$ is unramified, there are no ramification conditions*

on $\mathfrak{c}$). *Then* $j(\mathfrak{a}) \in \mathbb{Z}_{K(1)}$ (*in particular, it is an algebraic integer*) *and*

$$j(\mathfrak{a})^{\mathrm{Art}(\mathfrak{c})} = j(\mathfrak{a}\mathfrak{c}^{-1}) \ .$$

Since we know that $K(1) = K(j(\mathbb{Z}_K))$, it follows from this proposition that

$$\alpha = \mathrm{Tr}_{K(1)/K(1)^C}(j(\mathbb{Z}_K)) = \sum_{\overline{\mathfrak{c}} \in \overline{C}} j(\mathfrak{c}^{-1}) \in K(1)^C \ .$$

It can be shown (see [Sch3]) that α does not belong to any subfield of $K(1)^C$, in other words, that for any $\mathfrak{b} \notin C$ we have

$$\sum_{\overline{\mathfrak{c}} \in \overline{C}} j(\mathfrak{b}^{-1}\mathfrak{c}^{-1}) \neq \sum_{\overline{\mathfrak{c}} \in \overline{C}} j(\mathfrak{c}^{-1}) \ .$$

It follows that $K(1)^C = K(\alpha)$, hence the problem of the construction of $K(1)^C$ is in principle solved.

As already remarked in [Coh0, Section 7.6] for the case of $K(1)$ itself, the coefficients of the polynomial obtained in this way are huge and the result is not satisfactory. To improve on this, one approach already mentioned in [Coh0] is to use the Weber functions instead of the j-function. More generally, we can use any reasonable modular function.

Let f be a function defined on the upper half-plane $\mathcal{H}$ which transforms under $\gamma = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{PSL}_2(\mathbb{Z})$ by

$$f(\gamma(\tau)) = f\left(\frac{a\tau + b}{c\tau + d}\right) = (c\tau + d)^k f(\tau) \ ,$$

where k is called the *weight* of f, assumed to be integral.

If $\mathfrak{a}$ is a fractional ideal and (ω_1, ω_2) is a $\mathbb{Z}$-basis of $\mathfrak{a}$ ordered so that $\mathrm{Im}(\omega_1/\omega_2) > 0$, we will set by abuse of notation

$$f(\mathfrak{a}) = \left(\frac{2i\pi}{\omega_2}\right)^k f\left(\frac{\omega_1}{\omega_2}\right) \ .$$

It is immediately checked that $f(\mathfrak{a})$ is independent of the chosen (oriented) basis, hence its definition makes sense. In particular, for the function $f(\tau) = j(\tau)$ we have $k = 0$, so we recover the definition of $j(\mathfrak{a})$ that we have given.

We will use this definition mainly for products of the Dedekind η-function $\eta(\tau)$. Recall that $\eta(\tau)$ is defined by

$$\eta(\tau) = e^{2i\pi\tau/24} \prod_{n \geq 1} (1 - q^n) \ ,$$

where, as usual, $q = e^{2i\pi\tau}$. By abuse of notation, we will write $e^{2i\pi\tau/24} = q^{1/24}$, but it is understood that it is *this* specific 24th root.

We have the identity

$$\eta(\tau) = q^{1/24}\left(1 + \sum_{n\geq 1}(-1)^n\left(q^{n(3n-1)/2} + q^{n(3n+1)/2}\right)\right) ,$$

which gives a fast way to compute $\eta(\tau)$ since the exponents of q in the sum are quadratic in n (see Corollary 6.3.16 below for a proof). In addition, the function $\eta(\tau)$ is almost modular of weight $1/2$. More precisely, since $\mathrm{PSL}_2(\mathbb{Z})$ is generated by $\tau \mapsto \tau + 1$ and $\tau \mapsto -1/\tau$, the following formulas suffice to characterize the transformation formula:

$$\eta(\tau+1) = e^{2i\pi/24}\eta(\tau), \qquad \eta\left(\frac{-1}{\tau}\right) = \left(\frac{\tau}{i}\right)^{1/2}\eta(\tau) ,$$

where we must choose the determination of the square root having positive real part. The exact transformation formula for $\gamma = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{PSL}_2(\mathbb{Z})$ is as follows (see Exercise 15). Normalize the representative of $\mathrm{PSL}_2(\mathbb{Z})$ in $\mathrm{SL}_2(\mathbb{Z})$ so that $c \geq 0$, and in addition $d > 0$ if $c = 0$. Write $c = 2^v g$ with g odd, where we set $g = 1$ and $v = 0$ for $c = 0$. We then have $\eta(\gamma(\tau)) = v_\eta(\gamma)(c\tau+d)^{1/2}\eta(\tau)$ (again with the principal part of the square root), where the *multiplier system* $v_\eta(\gamma)$ is given by

$$v_\eta(\gamma) = \left(\frac{a}{g}\right)\exp\left(\frac{2i\pi}{24}\left(ab + cd(1-a^2) - ac + 3(a-1)g + \frac{3v(a^2-1)}{2}\right)\right) .$$

The standard reduction algorithm ([Coh0, Algorithm 7.4.2]) gives the following algorithm for computing $\eta(\tau)$.

Algorithm 6.3.2 (Computation of $\eta(\tau)$). Given τ in the upper half-plane $\mathcal{H}$, this algorithm computes the value of $\eta(\tau)$.

1. [Initializations] Set $\zeta \leftarrow \exp(2i\pi/24)$, $p \leftarrow 1$.
2. [Reduce real part] Set $n \leftarrow \lfloor \mathrm{Re}(\tau) \rceil$. If $n \neq 0$, set $\tau \leftarrow \tau - n$ and $p \leftarrow p\zeta^n$.
3. [Inverse?] Set $m \leftarrow \tau\bar{\tau}$. If $m \leq 0.999$, set (in this order) $p \leftarrow p\sqrt{i/\tau}$ (with $\mathrm{Re}(\sqrt{i/\tau}) > 0$), $\tau \leftarrow -\bar{\tau}/m$, and go to step 2.
4. [Start computation of η] Set $q_1 \leftarrow \exp(2i\pi\tau/24)$, $q \leftarrow q_1^{24}$, $s \leftarrow 1$, $q_s \leftarrow 1$, and $q_n \leftarrow 1$.
5. [Main loop] Set $t \leftarrow -qq_n^2q_s$, $q_n \leftarrow q_nq$, $q_s \leftarrow q_nt$, and $s \leftarrow s + t + q_s$.
6. [Finished?] If q_s is less than the desired relative accuracy, output pq_1s and terminate the algorithm. Otherwise, go to step 5.

Proof. The proof of this algorithm's validity follows from the transformation formula and the power series expansion of η. The variable p contains the accumulated products coming from the transformation formula, and in

the nth loop, q_n contains q^n, q_t contains $(-1)^n q^{n(3n-1)/2}$, and q_s contains $(-1)^n q^{n(3n+1)/2}$. The details are left to the reader (Exercise 9). □

Remarks

(1) It would be slightly nicer to have p contain the square of the accumulated products coming from the transformation formula, since this would avoid computing all square roots except one at the end. This is not possible, however, without explicitly using the complete transformation formula, since it is *essential* to multiply by the principal part of the square root at each step, and it is *not* true that $\sqrt{xy} = \sqrt{x}\sqrt{y}$, where $\sqrt{x}$ denotes the principal part of the square root of x. For simplicity, we have preferred to give the algorithm in this form.

(2) In the reduction process (step 3), we have written $m \leq 0.999$ instead of $m < 1$ to avoid roundoff errors. Indeed, this practically does not influence the speed of convergence of the series computed in step 5, and it avoids infinite loops that may occur, because with roundoff errors we may well simultaneously have $|\tau| < 1$ and $|-1/\tau| < 1$.

Since $\eta(\tau)$ possesses a multiplier system v_η under $\mathrm{PSL}_2(\mathbb{Z})$ transformations, it is not possible to define $\eta(\mathfrak{a})$ for an ideal $\mathfrak{a}$ without imposing some restrictions. On the other hand, multiplicative combinations of η often lead to a trivial multiplier system, hence the value at an ideal $\mathfrak{a}$ makes sense.

For example, the function $\Delta(\tau) = \eta(\tau)^{24}$ is modular of weight 12, and hence $\Delta(\mathfrak{a})$ makes sense as we have defined it above:

$$\Delta(\mathfrak{a}) = \left(\frac{2i\pi}{\omega_2}\right)^{24} \eta^{24}\left(\frac{\omega_1}{\omega_2}\right) ,$$

where (ω_1, ω_2) is any oriented $\mathbb{Z}$-basis of $\mathfrak{a}$.

More subtle, but more important for our applications, is the following example. Let p and q be two integers coprime to 6 but not necessarily prime, and set

$$g_{p,q}(\tau) = \frac{\eta(\tau/p)\eta(\tau/q)}{\eta(\tau/pq)\eta(\tau)} .$$

The following proposition is an immediate consequence of the complete transformation formula for the η-function under $\mathrm{PSL}_2(\mathbb{Z})$ given above.

Proposition 6.3.3. *Let $\Gamma^0(pq)$ be the group of matrices $\gamma = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{PSL}_2(\mathbb{Z})$ such that $pq \mid b$.*

(1) *For any $\gamma \in \Gamma^0(pq)$, we have $g_{p,q}(\gamma(\tau)) = v_g(\gamma) g_{p,q}(\tau)$, where the multiplier system v_g is given by*

$$v_g(\gamma) = \exp\left(-\frac{2i\pi}{24}(p-1)(q-1)\left(cd(1-a^2) - ac + 3(a-1) + a\frac{b}{pq}\right)\right)$$

(using the same normalization for γ as above).

(2) *If e is an integer such that $24 \mid e(p-1)(q-1)$, the function $g_{p,q}^e(\tau)$ is invariant under $\Gamma^0(pq)$.*

Proof. Left to the reader (Exercise 10). □

If $\mathfrak{a}$ is an ideal coprime to $6pq$ and (ω_1, ω_2) an ordered $\mathbb{Z}$-basis of $\mathfrak{a}$, we can define $g_{p,q,e}(\mathfrak{a}) = g_{p,q}(\omega_1/\omega_2)^e$ as long as we impose on ω_1 and ω_2 a normalization condition that compels the basis transformations to be in $\Gamma^0(pq)$. We will now see how to do this.

Recall that an ideal $\mathfrak{p}$ is *primitive* if $\mathfrak{p}$ is an integral ideal and if there does not exist a natural integer $n \geq 2$ such that $\mathfrak{p}/n$ is also integral. Recall also that if $\mathfrak{p}$ is a primitive ideal, then $\mathfrak{p}$ has a canonical HNF basis, which can be written

$$\mathfrak{p} = p\mathbb{Z} \oplus \frac{-u + \sqrt{D}}{2}\mathbb{Z} \ ,$$

with $p = \mathcal{N}(\mathfrak{p})$ the absolute norm of $\mathfrak{p}$.

The following easy proposition will be crucial for us in the sequel.

Proposition 6.3.4. *Let $\mathfrak{p}$ and $\mathfrak{q}$ be two primitive ideals of respective norms p and q such that the product $\mathfrak{pq}$ is also primitive. If*

$$\mathfrak{pq} = pq\mathbb{Z} \oplus \frac{-w + \sqrt{D}}{2}\mathbb{Z} \ ,$$

then $\mathfrak{p} = p\mathbb{Z} \oplus ((-w + \sqrt{D})/2)\mathbb{Z}$ and $\mathfrak{q} = q\mathbb{Z} \oplus ((-w + \sqrt{D})/2)\mathbb{Z}$.

Proof. Since $\mathfrak{pq} \subset \mathfrak{p}$, we have $w \equiv u \pmod{2p}$ and similarly for $\mathfrak{q}$, proving the proposition. □

Corollary 6.3.5. *Let $\mathfrak{a}$, $\mathfrak{p}$, $\mathfrak{q}$ be three primitive ideals such that $\mathfrak{apq}$ is a primitive ideal. Let $a = \mathcal{N}(\mathfrak{a})$, $p = \mathcal{N}(\mathfrak{p})$, $q = \mathcal{N}(\mathfrak{q})$, and assume that e is a positive integer such that $24 \mid e(p-1)(q-1)$.*

(1) *There exists an oriented basis (ω_1, ω_2) of $\mathfrak{a}$ such that $(\omega_1, p\omega_2)$ is a basis of $\mathfrak{ap}$, $(\omega_1, q\omega_2)$ is a basis of $\mathfrak{aq}$, and $(\omega_1, pq\omega_2)$ is a basis of $\mathfrak{apq}$.*
(2) *The quantity $g_{p,q,e}(\omega_1/\omega_2)$ is independent of the choice of oriented basis satisfying* (1).

Proof. For (1), we write $\mathfrak{apq} = apq\mathbb{Z} \oplus ((-w + \sqrt{D})/2)\mathbb{Z}$. It follows from the proposition that $\omega_1 = (-w + \sqrt{D})/2$ and $\omega_2 = a$ is a suitable basis.

For (2), we note that (ω_1', ω_2') is another suitable basis of $\mathfrak{a}$ if and only if there exists $\gamma = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{PSL}_2(\mathbb{Z})$ such that $a\omega_1 + b\omega_2 \in \mathfrak{apq}$ and $pq(c\omega_1 + d\omega_2) \in \mathfrak{apq}$. This is equivalent to the single condition

$$b\omega_2 \in \mathfrak{apq} = \omega_1\mathbb{Z} \oplus pq\omega_2\mathbb{Z} \ ,$$

hence to $pq \mid b$, so $\gamma \in \Gamma^0(pq)$. Therefore, $g_{p,q,e}(\omega_1'/\omega_2') = g_{p,q,e}(\omega_1/\omega_2)$ since $g_{p,q,e}$ is invariant under $\Gamma^0(pq)$, as was to be proved. □

This leads to the following definition.

Definition and Proposition 6.3.6. *Let* $\mathfrak{a}$ *be* any *fractional ideal of* K, *and let* $\mathfrak{p}$, $\mathfrak{q}$, *and* e *be as above.*

(1) *There exists* $\alpha \in K$ *such that* $\alpha\mathfrak{a}$ *and* $\alpha\mathfrak{a}\mathfrak{p}\mathfrak{q}$ *are primitive ideals.*
(2) *If* (ω_1, ω_2) *is a basis of* $\alpha\mathfrak{a}$ *satisfying the hypotheses of the corollary, the quantity* $g_{p,q,e}(\omega_1/\omega_2)$ *is independent of* ω_1, ω_2, *and* α *and will be denoted* $g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})$.

Proof. By Corollary 1.2.11, there exists α such that $\alpha\mathfrak{a}$ is an integral ideal coprime to pq, and of course we may assume that it is primitive. In particular, $\alpha\mathfrak{a}\mathfrak{p}\mathfrak{q}$ is primitive, proving (1).

The corollary implies that $g_{p,q,e}(\omega_1, \omega_2)$ is independent of the chosen basis (ω_1, ω_2). On the other hand, if α' is such that $\alpha'\mathfrak{a}$ is a primitive ideal such that $(\alpha'\mathfrak{a})\mathfrak{p}\mathfrak{q}$ is primitive, then $(\omega_1', \omega_2') = (\alpha'/\alpha)(\omega_1, \omega_2)$ is a basis of $\alpha'\mathfrak{a}$, which clearly also satisfies the hypothesis of the corollary, and obviously $g_{p,q,e}(\omega_1'/\omega_2') = g_{p,q,e}(\omega_1/\omega_2)$, proving (2). □

Remark. One can give more general (and a little more complicated) statements than those above, valid with the condition $\mathfrak{p}$ and $\mathfrak{q}$ coprime to 6 only, and not necessarily $24 \mid e(p-1)(q-1)$. We will not need this generality, and we refer to [Sch1], [Sch2] for details.

It can be proved that Proposition 6.3.1 is also valid for the quantities $g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})$, in other words that these are algebraic integers (they are, in fact, even *units*), and that for any ideal $\mathfrak{c}$ we have

$$g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})^{\mathrm{Art}(\mathfrak{c})} = g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}\mathfrak{c}^{-1}) \ .$$

The main theorem proven by Schertz using a clever but quite simple idea is that, under suitable hypotheses on $\mathfrak{p}$, $\mathfrak{q}$, and e, the function $g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})$ can replace the function $j(\mathfrak{a})$ in the construction of the Hilbert class field.

Theorem 6.3.7 (Schertz). *Let* $(\mathfrak{a}_i)_{1\le i\le h(K)}$ *be a system of representatives of the ideal classes of* $K = \mathbb{Q}(\sqrt{D})$, *chosen to be primitive. Let* $\mathfrak{p}$ *and* $\mathfrak{q}$ *be ideals of* K *of norm* p *and* q, *respectively. Assume that*

(1) *the ideals* $\mathfrak{p}$ *and* $\mathfrak{q}$ *are primitive ideals that are nonprincipal;*
(2) *if both classes of* $\mathfrak{p}$ *and* $\mathfrak{q}$ *are of order* 2 *in the class group, these classes are equal;*
(3) *for all* i, $\mathfrak{p}\mathfrak{q}\mathfrak{a}_i$ *is a primitive ideal;*
(4) e *is a positive integer chosen such that* $24 \mid e(p-1)(q-1)$.

Set

$$P_{\mathfrak{p},\mathfrak{q},e}(X) = \prod_{1\le i\le h(K)} (X - g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}_i)) = \prod_{1\le i\le h(K)} \left(X - \left(\frac{\eta(\tau_i/p)\eta(\tau_i/q)}{\eta(\tau_i/pq)\eta(\tau_i)}\right)^e\right) ,$$

where $\mathfrak{a}_i\mathfrak{p}\mathfrak{q} = a_i(pq\mathbb{Z} + \tau_i\mathbb{Z})$.

Then $P_{\mathfrak{p},\mathfrak{q},e}(X) \in \mathbb{Z}[X]$, *it is irreducible in* $\mathbb{Z}[X]$ *and in* $K[X]$, *its constant term is equal to* ± 1, *and the field obtained by adjoining to* K *a root of* $P_{\mathfrak{p},\mathfrak{q},e}$ *is the Hilbert class field* $K(1)$ *of* K.

We refer to [Sch1] for the proof. The statement (and the proof) given by Schertz is slightly incorrect because of the omission of condition (2) above on classes of order 2. An example is $K = \mathbb{Q}(\sqrt{-30})$, which has class group isomorphic to $(\mathbb{Z}/2\mathbb{Z})^2$, with $\mathfrak{p}$ an ideal above 11, $\mathfrak{q}$ an ideal above 37, and $e = 1$. The polynomial $P_{\mathfrak{p},\mathfrak{q},e}$ found in this case is the *square* of an irreducible polynomial in $\mathbb{Z}[X]$. However, it is easy to correct the statement and proof as above, as was remarked by the author and Schertz himself (see [Sch4]). □

Remark. Using the complete determination of imaginary quadratic fields of class number 2, one can prove that the exponent e can always be chosen equal to 1 except in the following cases:

- for $D = -20$, -52, and -148, where $e = 2$ is possible;
- for $D = -15$, -24, -51, -123, and -267, where $e = 3$ is possible.

These are exactly the discriminants $D < 0$ for which $h(D) = 2$ and $D \equiv 4 \pmod 8$ or $D \equiv 0 \pmod 3$, respectively.

If $\overline{C}$ is a subgroup of $Cl(K)$, it is easily shown that one of the symmetric functions from $K(1)$ to $K(1)^C$ of $g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}_i)$ (for any i) generates the field $K(1)^C$. In practice, the trace defined by

$$\mathrm{Tr}_{K(1)/K(1)^C}(g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}_i)) = \sum_{\mathfrak{c}\in\overline{C}} g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}_i\mathfrak{c}^{-1})$$

always works (no counterexample has been found), but I do not know if this can be proven. In the algorithm given below, we will assume that this is the case. If not, the necessary modifications are immediate and left to the reader.

We can now easily transform [Coh0, Algorithm 7.6.1], which computes the minimal polynomial of the j-function, into an algorithm that computes the minimal polynomial of the function $g_{\mathfrak{p},\mathfrak{q},e}$. Contrary to that algorithm, however, it is necessary to precompute all the reduced forms before choosing $\mathfrak{p}$ and $\mathfrak{q}$ since we must have $\mathfrak{a}_i\mathfrak{p}\mathfrak{q}$ primitive. Another possibility would be to choose $\mathfrak{p}$ and $\mathfrak{q}$ first and then multiply any reduced form by suitable elements to make it coprime to pq. This is probably more expensive. The algorithm is thus as follows.

Algorithm 6.3.8 (Unramified Abelian Extensions Using CM). Given an imaginary quadratic field $K = \mathbb{Q}(\sqrt{D})$ and a subgroup $\overline{C}$ of the class group $Cl(K)$ of cardinality equal to m, this algorithm returns a polynomial $P \in \mathbb{Z}[X]$, irreducible in $K[X]$, such that the extension of K defined by P is the fixed field by $\overline{C}$ of the Hilbert class field of K.

1. [Compute reduced forms] If $D \geq -11$, set $P(X) \leftarrow X$ and terminate. Otherwise, using Subalgorithm 6.3.9 below, compute the list $\mathcal{L}$ of reduced forms of discriminant D, as well as the auxiliary numbers h and z.
2. [Partition $\mathcal{L}$] Partition $\mathcal{L}$ as a disjoint union of cosets

$$\mathcal{L} = \bigcup_{1 \leq j \leq h/m} \mathcal{L}_j$$

under the multiplicative action of the group $\overline{C}$. If $h/m = 1$, output $P(X) \leftarrow X$ and terminate the algorithm.

3. [Find $\mathfrak{p}$ and $\mathfrak{q}$] Using Subalgorithm 6.3.10 below, find two suitable ideals $\mathfrak{p}$ and $\mathfrak{q}$ of norms p and q, respectively, and a positive exponent e. Let $pq\mathbb{Z} \oplus ((-u + \sqrt{D})/2)\mathbb{Z}$ be the HNF of the ideal $\mathfrak{pq}$.
4. [Initialize loop on cosets] Set $j \leftarrow 0$, $P(X) \leftarrow 1$.
5. [Loop on cosets] Set $j \leftarrow j + 1$. If $j > h/m$, round the coefficients of $P(X)$ to the nearest integer, output $P(X)$, and terminate the algorithm.
6. [Initialize loop on forms] Set $i \leftarrow 0$, $s \leftarrow 0$.
7. [Loop on forms] Set $i \leftarrow i + 1$. If $i > m$, set $P(X) \leftarrow (X - s)P(X)$ and go to step 5. Otherwise, let (a, b) be the ith element of $\mathcal{L}_j$. By the Chinese remainder theorem, find w such that $w \equiv -b \pmod{2a}$ and $w \equiv u \pmod{2pq}$, and set $\alpha \leftarrow (-w + \sqrt{D})/2a$.
8. [Compute $g_{\mathfrak{p},\mathfrak{q},e}(\alpha)$] Using Algorithm 6.3.2, compute

$$g_{\mathfrak{p},\mathfrak{q},e}(\alpha) \leftarrow \left(\frac{\eta(\alpha/p)\eta(\alpha/q)}{\eta(\alpha/pq)\eta(\alpha)} \right)^e ,$$

set $s \leftarrow s + g_{\mathfrak{p},\mathfrak{q},e}(\alpha)$, and go to step 7.

Remarks

(1) Contrary to the case of the j-function, where one can use $j(\mathfrak{a}^{-1}) = \overline{j(\mathfrak{a})}$, one cannot cut the work in half by computing only the values of $g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})$ for $b \geq 0$, since in general it is not true that $g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}^{-1}) = \overline{g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})}$ (see Exercise 11).

(2) The accuracy to which the computations must be made is not completely clear a priori. Since the coefficients of the polynomial $P_{\mathfrak{p},\mathfrak{q},e}(X)$ will be much smaller than those of the minimal polynomial of j, this should not be too much of a problem. To be perfectly rigorous, however, once $P(X)$ is obtained (and after checking that the rounding process is reasonable, for example, when the coefficients are at most 10^{-5}, say, from integers), we should check that $P(X)$ does indeed define the desired extension. For this, since we know a complex approximation to the roots of $P(X)$, it is easy to prove rigorously that we have an Abelian extension. Furthermore, the computation of the relative discriminant of this extension will show that

it is unramified. Finally, we check that the Galois group of $K(1)/K(1)^C$ is isomorphic to $\overline{C}$ under the Artin reciprocity map. In practice, these verifications are usually not necessary, and the closeness of the rounding process suffices to guarantee correctness.

(3) The size of the coefficients of the polynomial $P_{\mathfrak{p},\mathfrak{q},e}(X)$ computed by the algorithm is very much dependent on the size of e (since this governs the number of η-products in the function g), and not on the size of p and q since it is primarily the ideal *class* of the corresponding ideals $\mathfrak{p}$ and $\mathfrak{q}$ which matter. However, there is a great variability in the size of the coefficients that one obtains. Since the class group is finite and since $g^{24}_{\mathfrak{p},\mathfrak{q},1}$ depends only on the ideal class of $\mathfrak{p}$ and $\mathfrak{q}$, the number of possible polynomials $P_{\mathfrak{p},\mathfrak{q},e}(X)$ as $\mathfrak{p}$, $\mathfrak{q}$, and e vary (with D fixed) is finite. Thus, if we want small polynomials and if we are willing to waste some time, it is worthwhile to apply the algorithm for several suitable pairs $(\mathfrak{p},\mathfrak{q})$ and to take the best polynomial that one obtains.

Let us give an example with $D = -199$. The "best" polynomial in some sense is the polynomial

$$P(X) = x^9 - x^8 - 3x^6 + 3x^3 + 3x^2 + 5x + 1 \ ,$$

which is obtained, for example, for $p = 31$ and $q = 53$ (and an infinity of other pairs), while the "worst" polynomial is the polynomial

$$P(X) = x^9 - 10x^8 + 43x^7 - 106x^6 + 172x^5 - 189x^4 + 135x^3 - 58x^2 + 14x - 1 \ ,$$

obtained for $p = 29$ and $q = 157$. I believe that in this case, for $e = 1$ only 40 different polynomials $P_{\mathfrak{p},\mathfrak{q},1}$ are possible, for $e = 2$ only 20 different polynomials $P_{\mathfrak{p},\mathfrak{q},2}$ are possible, for $e = 3$ only 40 different polynomials $P_{\mathfrak{p},\mathfrak{q},3}$ are possible, and finally for $e = 6$ only 20 different polynomials $P_{\mathfrak{p},\mathfrak{q},6}$ are possible.

The following is a simple algorithm for making a list of reduced forms, essentially identical to [Coh0, Algorithm 5.3.5], which is needed in step 1 of the main algorithm.

Subalgorithm 6.3.9 (List of Reduced Forms). Given an imaginary quadratic field $K = \sqrt{D}$, this subalgorithm computes the list $\mathcal{L}$ of reduced forms of discriminant D, the number h of such forms, and the product z of all the norms a of the ideals corresponding to the reduced forms.

1. [Initialize] Set $\mathcal{L} \leftarrow \varnothing$, $b \leftarrow D \bmod 2$, $b_2 \leftarrow b$, $h \leftarrow 0$, and $z \leftarrow 1$.
2. [Initialize a] Set $t \leftarrow (b_2 - D)/4$ and $a \leftarrow \max(b, 1)$.
3. [Test] If $a \nmid t$, go to step 4. Otherwise, set $z \leftarrow az$; if $a = b$ or $a^2 = t$ or $b = 0$, set $\mathcal{L} \leftarrow \mathcal{L} \cup \{(a,b)\}$ and $h \leftarrow h + 1$; else set $\mathcal{L} \leftarrow \mathcal{L} \cup \{(a,b),(a,-b)\}$ and $h \leftarrow h + 2$.
4. [Loop on a] Set $a \leftarrow a + 1$. If $a^2 \leq t$, go to step 3.
5. [Loop on b] Set $b \leftarrow b + 2$ and $b_2 \leftarrow b^2$. If $3b_2 \leq |D|$, go to step 2. Otherwise, output the list $\mathcal{L}$, the numbers h, z, and terminate the subalgorithm.

To find the ideals $\mathfrak{p}$ and $\mathfrak{q}$, we can use the following subalgorithm.

Subalgorithm 6.3.10 (Find Suitable $\mathfrak{p}$ and $\mathfrak{q}$). Given D and the integer z as computed by the main algorithm, this subalgorithm finds two primitive non-principal ideals $\mathfrak{p}$ and $\mathfrak{q}$ of norms p and q, respectively, and a positive exponent e such that $\mathfrak{pq}$ is primitive, $\mathfrak{p}$ and $\mathfrak{q}$ are in the same ideal class if the classes of $\mathfrak{p}$ and $\mathfrak{q}$ are of order 2 in the class group, and such that $(p,z) = (q,z) = 1$, $24 \mid e(p-1)(q-1)$ and e is as small as possible.

1. [Make list of primes and forms] Set $\mathcal{P} \leftarrow \varnothing$ and $\mathcal{F} \leftarrow \varnothing$. For each prime ℓ such that $5 \le \ell \le 500$, do as follows. If $\ell \nmid z$ and $\left(\frac{D}{\ell}\right) = 1$, compute a square root u of D modulo 4ℓ, and let (a,b,c) be the reduced quadratic form obtained by reducing the form $(\ell, u, (u^2 - D)/(4\ell))$ using [Coh0, Algorithm 5.4.2]. Finally, if $a > 1$, set $\mathcal{P} \leftarrow \mathcal{P} \cup \{\ell\}$ and $\mathcal{F} \leftarrow \mathcal{F} \cup (a,b,c)$. The forms in $\mathcal{F}$ will be indexed by the prime numbers of $\mathcal{P}$, and the form (a,b,c) corresponding to ℓ will be denoted by F_ℓ.
2. [Find p] If in the list of primes $\mathcal{P}$ there exists an $\ell \equiv 1 \pmod 3$, let p be the smallest such ℓ; otherwise, let p be the smallest element of $\mathcal{P}$.
3. [Check order 2] Let $(a,b,c) \leftarrow F_p$. If (a,b,c) is of order 2, that is, if $b = 0$ or $|b| = a$ or $a = c$, then if $p \equiv 3 \pmod 4$ go to step 4a; otherwise, go to step 4b. If (a,b,c) is not of order 2, then if $p \equiv 3 \pmod 4$, go to step 4c; otherwise, go to step 4d.

4a. [$\mathfrak{p}$ is of order 2 and $p \equiv 3 \pmod 4$] If in the list of primes $\mathcal{P}$ there exists an ℓ such that $\ell \equiv 1 \pmod 4$, and F_ℓ either is not of order 2 or is equal to (a,b,c), then let q be the smallest such ℓ; otherwise, let q be the smallest element $\ell \in \mathcal{P}$ such that F_ℓ is not of order 2 or is equal to (a,b,c). Go to step 5.

4b. [$\mathfrak{p}$ is of order 2 and $p \equiv 1 \pmod 4$] Let q be the smallest $\ell \in \mathcal{P}$ such that F_ℓ is not of order 2 or is equal to (a,b,c). Go to step 5.

4c. [$\mathfrak{p}$ is not of order 2 and $p \equiv 3 \pmod 4$] If in the list of primes $\mathcal{P}$ there exists an ℓ such that $\ell \equiv 1 \pmod 4$, let q be the smallest such ℓ; otherwise, let q be the smallest element $\ell \in \mathcal{P}$. Go to step 5.

4d. [$\mathfrak{p}$ is not of order 2 and $p \equiv 1 \pmod 4$] Let q be the smallest element $\ell \in \mathcal{P}$.

5. [Terminate] Set $e \leftarrow 24/\gcd((p-1)(q-1), 24)$, output e, a prime ideal factor $\mathfrak{p}$ of p and $\mathfrak{q}$ of q (the same if $p = q$), and terminate the subalgorithm.

Remarks

(1) The condition $a > 1$ on the reduced form (a,b,c) used in step 1 of the subalgorithm is exactly the condition that the prime ideals above ℓ are nonprincipal.

(2) For simplicity we use only prime ideals $\mathfrak{p}$ and $\mathfrak{q}$. They are automatically primitive, and if $p \neq q$ or if $p = q$ and $\mathfrak{p} = \mathfrak{q}$, the ideal $\mathfrak{pq}$ is also primitive.

(3) The bound 500 is arbitrary but should be more than sufficient, since the probability that a given ℓ not dividing z satisfies the conditions is greater or equal to $1/4$ (more precisely, it is equal to $(h-1)/(2h)$).

(4) As already mentioned, the integer e output by the subalgorithm will be equal to 1 except for eight discriminants for which it will be at most equal to 3.

6.3.3 Quasi-Elliptic Functions

Our next goal will be to compute ray class fields of imaginary quadratic fields. For this purpose, we will need to use classical meromorphic functions, which are closely related to elliptic functions, whose properties we recall here. For more details on this beautiful and very classical theory, the reader is strongly urged to consult any standard textbook on elliptic functions such as [Lan1].

Let L be a complex lattice and let (ω_1, ω_2) be an oriented $\mathbb{Z}$-basis of L. To ease notation, we will set $L^* = L \setminus \{0\}$ (we will never use the dual lattice of L in this chapter, so there is no risk of confusion). Recall that an *elliptic function* f is a meromorphic function on the complex plane such that $f(z+\omega) = f(z)$ for all $\omega \in L$; in other words, it is a meromorphic doubly periodic function. A prototypical example of such a function is the Weierstrass $\wp$-function $\wp(z, L)$ defined by the usual formula

$$\wp(z, L) = \frac{1}{z^2} + \sum_{\omega \in L^*} \left(\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} \right) .$$

Note that for aesthetic reasons, we prefer to use $z - \omega$ instead of $z + \omega$ in the sum.

The field of elliptic functions is the field $\mathbb{C}(\wp, \wp')$, and $\wp$ and $\wp'$ are linked by the algebraic differential equation

$$\wp'(z, L)^2 = 4\wp(z, L)^3 - g_2(L)\wp(z, L) - g_3(L) ,$$

with

$$g_2(L) = 60 \sum_{\omega \in L^*} \frac{1}{\omega^4}, \quad g_3(L) = 140 \sum_{\omega \in L^*} \frac{1}{\omega^6} .$$

Proposition 6.3.11. *Let L be a complex lattice.*

(1) *There exists a unique meromorphic function $\zeta(z, L)$, called the* Weierstrass ζ-function, *such that $\zeta'(z, L) = -\wp(z, L)$ and such that $\zeta(z, L)$ is an odd function.*

(2) *We have the following expansion, valid for all $z \notin L$:*

$$\zeta(z, L) = \frac{1}{z} + \sum_{\omega \in L^*} \left(\frac{1}{(z-\omega)} + \frac{1}{\omega} + \frac{z}{\omega^2} \right) = \frac{1}{z} + z^2 \sum_{\omega \in L^*} \left(\frac{1}{\omega^2 (z-\omega)} \right) .$$

(3) *There exist complex constants η_1 and η_2, called the* quasi-periods *of ζ associated to the periods ω_1 and ω_2, such that for any integers m and n we have*

$$\zeta(z + m\omega_1 + n\omega_2, L) = \zeta(z, L) + m\eta_1 + n\eta_2 ,$$

and in particular we have $\eta_i = 2\zeta(\omega_i/2, L)$ for $i = 1$ and $i = 2$.

(4) *Set* $\tau = \omega_1/\omega_2$, $q = \exp(2i\pi\tau)$, $u = \exp(2i\pi z/\omega_2)$, *and*

$$E_2(\tau) = 1 - 24\sum_{n\geq 1}\frac{nq^n}{1-q^n} = 1 - 24\sum_{n\geq 1}\frac{q^n}{(1-q^n)^2} = 1 - 24\sum_{n\geq 1}\sigma(n)q^n \ ,$$

where $\sigma(n) = \sigma_1(n)$ *is the sum of the positive divisors of* n. *Then*

$$\zeta(z,L) = -\frac{2i\pi}{\omega_2}\left(\frac{2i\pi}{\omega_2}\frac{E_2(\tau)}{12}z + \frac{1}{2}\frac{1+u}{1-u} + \sum_{n\geq 1}q^n\left(\frac{u}{1-q^nu} + \frac{1}{q^n-u}\right)\right) \ .$$

(5) *The quasi-periods are given by the formulas*

$$\eta_1 = \frac{\pi^2}{3\omega_2}\tau E_2(\tau) - \frac{2i\pi}{\omega_2} \quad \textit{and} \quad \eta_2 = \frac{\pi^2}{3\omega_2}E_2(\tau) \ ,$$

and, in particular, $\omega_1\eta_2 - \omega_2\eta_1 = 2i\pi$.

Proof. Since the proofs are easy, we leave some details to the reader (Exercise 13).

It is clear that the series defining $\zeta(z,L)$ converges uniformly on any compact subset not containing points of L (this is the case as soon as the general term goes to zero faster than $1/|\omega|^\alpha$ for any $\alpha > 2$). Thus the series defines a meromorphic function on $\mathbb{C}$ with poles at points of L, and by differentiating termwise it is clear that $\zeta(z,L)' = -\wp(z,L)$. In addition,

$$-\zeta(-z,L) = z^{-1} + z^2\sum_{\omega\in L^*}(\omega^2(z+\omega))^{-1} = z^{-1} + z^2\sum_{\omega\in L^*}(\omega^2(z-\omega))^{-1} \ ,$$

so $\zeta(z,L)$ is an odd function, thus proving (1) and (2), since clearly the property of being odd makes ζ unique among all antiderivatives of $-\wp$.

(3). Let $\omega \in L$, and set $f(z) = \zeta(z+\omega,L) - \zeta(z,L)$. Since the derivative of ζ is an elliptic function, it follows that $f'(z) = 0$, hence that $f(z)$ is constant, since $\mathbb{C} \setminus L$ is connected. Thus we can set for $i = 1$, and $i = 2$, $\eta_i = \zeta(z+\omega_i,L) - \zeta(z,L)$, so in particular $\eta_i = \zeta(\omega_i/2,L) - \zeta(-\omega_i/2,L) = 2\zeta(\omega_i/2,L)$ since $\zeta(z,L)$ is odd, and (3) follows by induction on m and n.

(4). Setting $u = e^{2i\pi z/\omega_2}$ and slightly modifying [Coh0, Proposition 7.4.4] (including the exchange of ω_1 and ω_2 already mentioned), we know that

$$\wp(z,L) = \left(\frac{2i\pi}{\omega_2}\right)^2\left(\frac{E_2(\tau)}{12} + \sum_{n=-\infty}^{\infty}\frac{q^nu}{(1-q^nu)^2}\right) \ .$$

Integrating termwise with respect to z, and taking care to take suitable integration constants, we obtain the given formula for $\zeta(z,L)$, proving (4).

(5). Since u is unchanged when z is changed into $z+\omega_2$, it is clear that

$$\eta_2 = \zeta(z+\omega_2,L) - \zeta(z,L) = -\frac{(2i\pi)^2}{\omega_2}\frac{E_2(\tau)}{12} = \frac{\pi^2}{3\omega_2}E_2(\tau) \ .$$

When z is changed into $z+\omega_1$, u is changed into qu and it is easy to see that the series for $\zeta(z+\omega_1,L)-\zeta(z,L)$ almost cancel and give the result of the proposition. More elegantly, by integrating the function $\zeta(z,L)$ along a fundamental parallelogram not crossing L and using the residue theorem, we immediately find the relation $\omega_1\eta_2-\omega_2\eta_1=2i\pi$, giving the formula for η_1. □

Formula (4) gives a fast way to compute $\zeta(z,L)$, after reduction to the fundamental domain, in a manner analogous to the other algorithms for computing functions of this sort, such as Algorithm 6.3.2 or [Coh0, Algorithm 7.4.5]. Since we will not need this algorithm, we leave the details to the reader (Exercise 14). On the other hand, we will need to compute the quasi-periods η_1 and η_2. To be able to reduce to the fundamental domain, we need to know the behavior of $E_2(\tau)$ under the action of $\mathrm{PSL}_2(\mathbb{Z})$. This is given by the following corollary.

Corollary 6.3.12. *For any* $\left(\begin{smallmatrix}a&b\\c&d\end{smallmatrix}\right)\in\mathrm{PSL}_2(\mathbb{Z})$ *and* $\tau\in\mathcal{H}$, *we have*

$$E_2\left(\frac{a\tau+b}{c\tau+d}\right)=(c\tau+d)^2E_2(\tau)+\frac{12c(c\tau+d)}{2i\pi}\ .$$

Proof. Set $\omega_1'=a\omega_1+b\omega_2$, $\omega_2'=c\omega_1+d\omega_2$, and $\tau'=\omega_1'/\omega_2'$. By assumption, (ω_1',ω_2') is still an oriented basis of L, so Proposition 6.3.11 applied to this basis gives in particular

$$\zeta(z+\omega_2',L)=\zeta(z,L)+\frac{\pi^2}{3\omega_2'}E_2(\tau')=\zeta(z,L)+\frac{\pi^2}{3\omega_2(c\tau+d)}E_2\left(\frac{a\tau+b}{c\tau+d}\right)\ .$$

On the other hand, the same proposition applied to the basis (ω_1,ω_2) gives

$$\zeta(z+\omega_2',L)=\zeta(z,L)+c\eta_1+d\eta_2=\zeta(z,L)+\frac{\pi^2}{3\omega_2}(c\tau+d)E_2(\tau)-c\frac{2i\pi}{\omega_2}\ ,$$

which gives the corollary by identification. □

From this, it is immediate to obtain an efficient algorithm for computing the quasi-periods.

Algorithm 6.3.13 (Computation of Quasi-Periods). Given an oriented basis (ω_1,ω_2) generating a complex lattice, this algorithm computes the quasi-periods η_1 and η_2 associated to ω_1 and ω_2.

1. [Reduce to fundamental domain] Set $\tau\leftarrow\omega_1/\omega_2$. Using the reduction algorithm [Coh0, Algorithm 7.4.2] on τ, compute a matrix $A=\left(\begin{smallmatrix}a&b\\c&d\end{smallmatrix}\right)$ and $\tau'=A\tau$ such that τ' is in the standard fundamental domain for $\mathrm{SL}_2(\mathbb{Z})$.
2. [Compute series] Set $q\leftarrow\exp(2i\pi\tau')$ and compute to the desired accuracy the sum of the series $E\leftarrow 1-24\sum_{n\geq1}nq^n/(1-q^n)$ (because of the preliminary

reductions, we will have $|q| \leq \exp(-\pi\sqrt{3}) < 1/230$, so this series will converge very rapidly).

3. [Compute corrected value] Set $u \leftarrow 1/(c\tau + d)$ and $E \leftarrow u^2 E + 6iuc/\pi$ (this is now the value of $E_2(\tau)$).
4. [Terminate] Set $\eta_2 \leftarrow \pi^2 E/(3\omega_2)$, set $\eta_1 \leftarrow \tau\eta_2 - 2i\pi/\omega_2$, output η_1 and η_2, and terminate the subalgorithm.

The most interesting quasi-elliptic function, the Weierstrass σ-function, is obtained essentially by integrating one more time as follows.

Proposition 6.3.14. *Let L be a complex lattice.*

(1) *There exists a unique holomorphic function $\sigma(z, L)$, called the* Weierstrass σ-function, *such that $\sigma'(z, L)/\sigma(z, L) = \zeta(z, L)$ and $\lim_{z\to 0} \sigma(z, L)/z = 1$.*

(2) *The function $\sigma(z, L)$ is an odd function having simple zeros exactly at all points of L. More precisely, we have the expansion*

$$\sigma(z, L) = z \prod_{\omega \in L^*} \left(1 - \frac{z}{\omega}\right) e^{\frac{z}{\omega} + \frac{z^2}{2\omega^2}} .$$

(3) *For any integers m and n we have*

$$\sigma(z + m\omega_1 + n\omega_2, L) = \pm e^{\eta_\omega (z+\omega/2)} \sigma(z, L) ,$$

where $\omega = m\omega_1 + n\omega_2$, $\eta_\omega = m\eta_1 + n\eta_2$, and the $\pm$ sign is equal to 1 if m and n are both even, and to -1 otherwise.

(4) *We have the expansion*

$$\sigma(z, L) = \frac{\omega_2}{2i\pi} e^{\eta_2 z^2/(2\omega_2)} (u^{1/2} - u^{-1/2}) \prod_{n \geq 1} \frac{(1 - q^n u)(1 - q^n/u)}{(1 - q^n)^2} ,$$

where $u^{\pm 1/2}$ is interpreted as $\exp(\pm i\pi z/\omega_2)$.

Proof. Once again, we leave some details to the reader (Exercise 16). The general term of the given product expansion tends to 1 as $|\omega|^{-3}$, hence the product converges uniformly on any compact subset of $\mathbb{C}$ and so defines a holomorphic function. By definition, its logarithmic derivative is equal to $\zeta(z, L)$, and we have $\lim_{z\to 0} \sigma(z, L)/z = 1$, a condition that also ensures uniqueness of $\sigma(z, L)$. Since L is symmetrical with respect to the origin, $\sigma(z, L)$ is an odd function, proving (1) and (2).

For (3), set $f_\omega(z) = \sigma(z + \omega, L)/\sigma(z, L)$. By definition and by Proposition 6.3.11, we have $f'_\omega(z)/f_\omega(z) = \zeta(z + \omega, L) - \zeta(z, L) = \eta_\omega$. It follows that $f_\omega(z) = C_\omega e^{\eta_\omega z}$ for a suitable constant C_ω depending on ω.

Assume first that m and n are not both even, in other words that $z_0 = -\omega/2 \notin L$. Since σ is an odd function and z_0 is not a zero of σ, we have

$f_\omega(z_0) = -1$, which gives $C_\omega = -e^{-\eta_\omega z_0}$, so that $f_\omega(z) = -e^{\eta_\omega(z+\omega/2)}$. This proves (3) when $\omega \notin 2L$. If $\omega \in 2L$, we can, for example, write $\omega = \omega - \omega_1 + \omega_1$ and use what we have just proved for $\omega - \omega_1$ and ω_1, which do not belong to $2L$, giving the final result of (3).

For (4), a short computation using Proposition 6.3.11 (4) gives

$$\sigma(z, L) = Ce^{\eta_2 z^2/(2\omega_2)}(u^{1/2} - u^{-1/2}) \prod_{n\geq 1}(1 - q^n u)(1 - q^n/u)$$

for some constant C. To determine C, we use $\lim_{z\to 0} \sigma(z, L)/z = 1$, which gives $C(2i\pi/\omega_2)\prod_{n\geq 1}(1-q^n)^2 = 1$, finishing the proof of the proposition. □

As for the ζ-function, (4) allows us to compute $\sigma(z, L)$ efficiently after suitable reductions to the fundamental domain (see Exercise 16). The σ-function is, however, connected to other types of functions, the theta functions, and this connection gives an even more efficient way to compute $\sigma(z, L)$. The result is the *Jacobi triple-product identity* as follows.

Proposition 6.3.15 (Jacobi Triple-Product Identity). *We have the identity*

$$\prod_{n\geq 1}(1 - q^n u)\left(1 - \frac{q^n}{u}\right)(1 - q^n) = \sum_{k\geq 0}(-1)^k \frac{u^{2k+1} - 1}{u^k(u-1)} q^{\frac{k(k+1)}{2}} ,$$

both formally and as an identity between complex numbers when $|q| < 1$ and $u \neq 0$.

Proof. There are many proofs of this famous identity. The following one is perhaps the simplest. We first do computation on polynomials. If we set

$$P_N(u, q) = (1 - u)\prod_{n=1}^{N}(1 - q^n u)\left(1 - \frac{q^n}{u}\right) ,$$

we can write

$$P_N(u, q) = \sum_{-N\leq k\leq N+1} a_{k,N}(q)u^k$$

for some polynomials $a_{k,N}(q)$. We have

$$\begin{aligned} P_N(qu, q) &= (1 - qu)\prod_{n=1}^{N}(1 - q^{n+1}u)\prod_{n=1}^{N}\left(1 - \frac{q^{n-1}}{u}\right) \\ &= -\frac{1-u}{u}\prod_{n=1}^{N+1}(1 - q^n u)\prod_{n=1}^{N-1}\left(1 - \frac{q^n}{u}\right) \end{aligned}$$

so that $P_N(qu,q)/P_N(u,q) = (1-q^{N+1}u)/(q^N-u)$, or in other words $(u-q^N)P_N(qu,q) = (q^{N+1}u-1)P_N(u,q)$. Identifying the coefficients of u^{k+1} for $-N \le k \le N$ gives $a_{k+1,N}(q)/a_{k,N}(q) = -q^k(1-q^{N+1-k})/(1-q^{N+k+1})$. On the other hand, it is clear on the definition that $a_{-N,N}(q) = \prod_{1\le n\le N}(-q^n) = (-1)^N q^{N(N+1)/2}$. Thus by induction we obtain

$$a_{k,N}(q) = (-1)^k q^{k(k-1)/2} \frac{\prod_{N+2-k\le n\le 2N+1}(1-q^n)}{\prod_{1\le n\le N+k}(1-q^n)} .$$

After simple transformations, this formula is equivalent to the finite identity

$$(1-u)\prod_{n=1}^{N}(1-q^nu)\left(1-\frac{q^n}{u}\right) = \sum_{k=-N}^{N+1}(-1)^k u^k q^{k(k-1)/2}\prod_{n=1}^{k+N}\frac{1-q^{2N+2-n}}{1-q^n} .$$

Considering u and q as formal variables, as $N \to \infty$, k being fixed, we have

$$a_{k,N}(q) \to (-1)^k \frac{q^{k(k-1)/2}}{\prod_{n\ge 1}(1-q^n)} ,$$

thus giving the Jacobi identity as a formal identity between power series after grouping the terms corresponding to k and $1-k$. The identity between complex numbers follows from this and immediate convergence arguments. □

Corollary 6.3.16.

$$\eta(\tau) = q^{1/24}\left(1+\sum_{k\ge 1}(-1)^k(q^{k(3k-1)/2}+q^{k(3k+1)/2})\right) .$$

Proof. This follows from the Jacobi identity by replacing (u,q) by (q,q^3) and rearranging terms. The details are left to the reader (Exercise 17). Note that this is the identity we used in Algorithm 6.3.2. □

Corollary 6.3.17.

$$\eta^3(\tau) = q^{1/8}\sum_{k\ge 0}(-1)^k(2k+1)q^{k(k+1)/2} .$$

Proof. This follows from the Jacobi identity by making $u \to 1$. □

Corollary 6.3.18. *With the same notation as above, we have the formula*

$$\sigma(z,L) = \frac{\omega_2}{2i\pi}e^{\eta_2 z^2/(2\omega_2)}\frac{\sum_{k\ge 0}(-1)^k(u^{k+1/2}-u^{-(k+1/2)})q^{k(k+1)/2}}{q^{-1/8}\eta(\tau)^3} .$$

Proof. Clear from Proposition 6.3.14 (4). □

Remarks

(1) Although not apparent in the above formula or in the q-product expansion of Proposition 6.3.14 (4), the value of $\sigma(z, L)$ really depends only on z and on the lattice L, and not on the particular oriented basis (ω_1, ω_2) of L. This follows from the definition of the σ-function.

(2) The expression $q^{k(k+1)/2}/q^{-1/8}$ can be written more nicely as $q^{(k+1/2)^2/2}$, but this is of no use for numerical computation.

(3) Corollary 6.3.18 is exactly the formula used in [Coh0, Algorithm 7.5.7] for computing the height contribution at infinity of a rational point on an elliptic curve.

Using Corollary 6.3.18 and reductions done using Proposition 6.3.14, we can now write an efficient algorithm for computing the σ-function.

Algorithm 6.3.19 (Computation of $\sigma(z, L)$). Given an oriented basis (ω_1, ω_2) of a complex lattice L and a complex number z, this algorithm computes the value of the Weierstrass σ-function $\sigma(z, L)$ at z.

1. [Reduce to fundamental domain] Using the reduction algorithm [Coh0, Algorithm 7.4.2] on $\tau = \omega_1/\omega_2$, compute a matrix $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and $\tau' = A\tau$ such that τ' is in the standard fundamental domain for $\mathrm{SL}_2(\mathbb{Z})$. Then set (in this order) $\omega_2 \leftarrow c\omega_1 + d\omega_2$, $\tau \leftarrow \tau'$, $\omega_1 \leftarrow \tau\omega_2$, and $z_2 \leftarrow z/\omega_2$.
2. [Reduce z] Set (in this order) $m \leftarrow \lfloor \mathrm{Im}(z_2)/\mathrm{Im}(\tau) \rceil$, $z_2 \leftarrow z_2 - m\tau$, $n \leftarrow \lfloor \mathrm{Re}(z_2) \rceil$, $z_2 \leftarrow z_2 - n$, and $z \leftarrow z_2\omega_2$.
3. [Compute corrections] Using Algorithm 6.3.13, compute the quasi-periods η_1 and η_2 associated to (ω_1, ω_2). Set $s \leftarrow (m\eta_1 + n\eta_2)(z + (m\omega_1 + n\omega_2)/2)$, and if m or n is odd, set $s \leftarrow s + i\pi$.
4. [Initialize] Set $y \leftarrow (\omega_2/(2i\pi))\exp(s + z_2 z\eta_2/2)$, $q_8 \leftarrow \exp(2i\pi\tau/8)$, $q \leftarrow q_8^8$, and $v \leftarrow \exp(i\pi z_2)$.
5. [Compute series] Compute to the desired accuracy the sum of the series
$$S \leftarrow \sum_{k \geq 0} (-1)^k (v^{2k+1} - v^{-2k-1}) q^{k(k+1)/2} \ .$$
6. [Terminate] Using Algorithm 6.3.2, compute $e \leftarrow \eta(\tau)$, output $q_8 y S/e^3$ as the value of $\sigma(z, L)$, and terminate the algorithm.

Remarks

(1) We have not written in detail the computation of the series used in the above algorithms, as we did in Algorithm 6.3.2, but this is of course easily done. Note that, because of the preliminary reductions, in the series S of step 5 we have $|q| \leq \exp(-\pi\sqrt{3}) < 1/230$ and $\exp(-\pi/2) \leq v \leq \exp(\pi/2)$, so the series converges rapidly (see Exercise 20).

(2) To compute $e^3 = \eta(\tau)^3$, we can either use Algorithm 6.3.2 as stated or use Corollary 6.3.17 together with the same reductions to the fundamental domain done in Algorithm 6.3.2.
(3) As already mentioned, the above algorithm is superior to the algorithm using the product expansion of Proposition 6.3.14 for multi-precision computations, and these are almost always necessary for ray class field computations.
(4) For our applications, we will not exactly need the σ-function itself, but a ratio of two values of the σ-function for the same lattice and for different values of z. It is then worthwhile to write a specific algorithm for this purpose, since it avoids some unnecessary computations.
(5) For certain values of the arguments the result may be large and overflow the possibilities of the implementation. Indeed, the quantities computed in the algorithm are of reasonable size, except perhaps the exponential computed in step 4 of the algorithm. Thus, instead of computing $\sigma(z, L)$ itself or a ratio of such, it is usually safer to compute the logarithm of such a ratio.

Elliptic and quasi-elliptic functions possess an amazing number of properties that were intensively studied at the end of the 19th century, and numerous thick treatises are devoted to the subject. We simply note the following formula.

Proposition 6.3.20. *Let L be a complex lattice and $a \notin L$. Then*

$$\wp(z, L) - \wp(a, L) = -\frac{\sigma(z-a, L)\sigma(z+a, L)}{\sigma(a, L)^2\sigma(z, L)^2} \ .$$

Proof. Using Proposition 6.3.14, it is easy to check that the ratio of the left- and right-hand sides is an elliptic function with no zero or poles, hence it is constant by Liouville's theorem. Making $z \to 0$ and using the expansions of $\wp$ and σ around 0 gives the result. □

The function we will use for the construction of ramified Abelian extensions comes from a modification of the σ-function constructed as follows. Since ω_1 and ω_2 are $\mathbb{R}$-linearly independent, for any $z \in \mathbb{C}$ there exist real numbers x_1 and x_2 such that $z = x_1\omega_1 + x_2\omega_2$, in other words (x_1, x_2) are the coordinates of z on the basis (ω_1, ω_2). We then set $z^*(L) = x_1\eta_1 + x_2\eta_2$, where the η_i are as usual the quasi-periods associated to the basis (ω_1, ω_2). Since the quasi-periods behave linearly in terms of the ω_i, it is clear that $z^*(L)$ does not depend on the basis (ω_1, ω_2) but only on z and on the lattice L, whence the notation.

We will set

$$\phi^*(z, L) = e^{-zz^*(L)/2}\sigma(z, L) \ .$$

Proposition 6.3.21. *Let $\omega = m\omega_1 + n\omega_2 \in L$ with m and n in $\mathbb{Z}$, and let $z = x_1\omega_1 + x_2\omega_2$ with x_1 and x_2 in $\mathbb{R}$, as above. Then*

$$\phi^*(z+\omega, L) = s(\omega)e^{2i\pi(nx_1 - mx_2)}\phi^*(z, L) \ ,$$

where $s(\omega) = +1$ if $\omega/2 \in L$, $s(\omega) = -1$ otherwise. In addition, the function ϕ^* is bounded on $\mathbb{C}$.

Proof. The proof of the transformation formula follows immediately from the corresponding formula for the σ-function seen above. It follows in particular that $|\phi^*(z+\omega, L)| = |\phi^*(z, L)|$, and since ϕ^* is a continuous (although nonholomorphic) function, $\phi^*(z, L)$ is bounded in any compact set, hence in any fundamental parallelogram of the form $(a, a+\omega_2, a+\omega_1+\omega_2, a+\omega_1)$. It follows that it is bounded on all of $\mathbb{C}$. In fact, it is not difficult to give explicit bounds if desired, using the fact that we can reduce to $\mathrm{Im}(\omega_1/\omega_2) \geq \sqrt{3}/2$ and to $z/\omega_2 = x + iy$ with $|x| \leq 1/2$ and $|y| \leq 1/2$ (Exercise 22). □

Since $\phi^*(z, L)$ is bounded (in fact, by a rather small constant; see Exercise 22), we can compute its values without having to worry about overflow problems, by first computing the logarithm of $\sigma(z, L)$, subtracting $zz^*(L)/2$, and only then computing the exponential, which is sure not to overflow.

6.3.4 Construction of Ramified Abelian Extensions Using Complex Multiplication

With the tools of the preceding section, we can now study the problem of constructing ray class fields of imaginary quadratic fields using complex multiplication.

In the unramified case, we had the simple result $K(1) = K(j(\mathbb{Z}_K))$, which unfortunately is difficult to use in practice because of very large coefficients, and so we introduced the more subtle functions $g_{\mathfrak{p},\mathfrak{q},e}$ to overcome this problem. In the ramified case, the situation is quite similar.

Set

$$w(z, L) = \begin{cases} -2^7 3^5 \dfrac{g_2(L)g_3(L)}{\Delta(L)}\wp(z, L) & \text{if } D \neq -3, -4 \ , \\ 2^8 3^4 \dfrac{g_2(L)^2}{\Delta(L)}\wp(z, L)^2 & \text{if } D = -4 \ , \\ -2^9 3^6 \dfrac{g_3(L)}{\Delta(L)}\wp(z, L)^3 & \text{if } D = -3 \ . \end{cases}$$

Note that the function w is "of weight 0", similarly to the functions j and $g_{\mathfrak{p},\mathfrak{q},e}$. These functions were introduced by Weber.

The following theorem is the analog of the corresponding theorem for the j-function.

Theorem 6.3.22. *Let K be an imaginary quadratic field, let $K(1)$ be the Hilbert class field of K, and let $\mathfrak{f}$ be a conductor of K. Let $\alpha \in K^*$ be such that $\alpha\mathfrak{f}$ is an integral ideal coprime to $\mathfrak{f}$ (which exists by Corollary 1.2.11). Then*

(1) *The ray class field $K(\mathfrak{f})$ of conductor $\mathfrak{f}$ is given by*

$$K(\mathfrak{f}) = K(1)(w(\alpha, \mathbb{Z}_K)) \ ,$$

where w is the function defined above.

(2) *If $\mathfrak{c}$ is an integral ideal of K coprime to $\mathfrak{f}$, then for any ideal $\mathfrak{b}$ we have*

$$w(\alpha, \mathfrak{b})^{\mathrm{Art}(\mathfrak{c})} = w(\alpha, \mathfrak{b}\mathfrak{c}^{-1}) \ ,$$

where $\mathrm{Art}(\mathfrak{c})$ is the element of $\mathrm{Gal}(K(\mathfrak{f})/K)$ corresponding to $\bar{\mathfrak{c}} \in Cl_{\mathfrak{f}}(K)$ by the Artin reciprocity map.

As in the unramified case, this shows that the problem of constructing $K(\mathfrak{f})$ (and in a similar way its subfields if we have a congruence subgroup C) is solved in principle. Once again the coefficients will be very large, so this construction cannot be used except in very small cases. Thus, we need to find other elliptic or quasi-elliptic functions that give smaller coefficients. The main problem is the presence of the modular function g_2g_3/Δ of weight -2, which leads to large coefficients. A first idea, introduced by Schertz in [Sch2], is to replace it by the function η^{-4}, a sixth root of $1/\Delta$, which is also of weight -2. It can be expected that the coefficients that will be obtained are much smaller (more precisely, with 4 times fewer decimal digits), and this is indeed the case. We leave to the reader the detailed study of this method (see Exercise 23).

More recently, R. Schertz has introduced an even better method for computing ramified Abelian extensions using complex multiplication (see [Sch4]). The main advantage of this method is that the coefficients of the resulting polynomials are considerably smaller than those obtained using the above-mentioned method (approximately 3 times fewer decimal digits than Schertz's preceding method; hence, in all, 12 times fewer compared to the method using directly Theorem 6.3.22). A small disadvantage is that in the form we give, it relies on the validity of an unproved technical conjecture, although the result itself can be checked a posteriori without assuming any conjecture. At the expense of a more complicated (but not slower) algorithm, it is possible to completely suppress the assumption of this conjecture (see Exercise 26), but for simplicity of exposition we will assume this conjecture here.

The main idea is to remove the normalizing factor involving $\Delta(L)$ in $w(z, L)$ by using quotients of values the Weierstrass σ-function instead of the $\wp$-function.

Thus, let K be an imaginary quadratic field of discriminant D, and let $\mathfrak{f}$ be an integral ideal of K. We assume that $\mathfrak{f}$ is the conductor of $(\mathfrak{f}, P_{\mathfrak{f}})$ (otherwise

use Algorithm 4.4.2 to reduce to this case). We want to compute the ray class field of conductor $\mathfrak{f}$. From this construction it will be easy to extract the necessary information to construct the ray class field corresponding to all the congruence subgroups $(\mathfrak{f}, C)$ if desired, exactly as we did in Algorithm 6.3.8, and we leave the details to the reader (Exercise 24). The theorem of Schertz that will allow us to construct ramified Abelian extensions is the following (we state only the special case of the theorem that we will need; see [Sch4] and [Sch5] for the complete version).

Theorem 6.3.23. *Let K be an imaginary quadratic field, let $\mathfrak{f}$ be a conductor of K, let f be the positive generator of $\mathfrak{f} \cap \mathbb{Z}$, let e be a positive integer, and let $\lambda \in \mathbb{Z}_K \setminus \{0\}$ satisfying the following conditions.*

(1) $e(\mathcal{N}_{K/\mathbb{Q}}(\lambda) - 1) \equiv 0 \pmod{2f}$.
(2) *The class of the ideal $\lambda\mathbb{Z}_K$ is not of order 1 or 3 in $Cl_{\mathfrak{f}}(K)$.*

For any primitive ideal $\mathfrak{c}$ coprime to f, set

$$\theta_{\lambda,\mathfrak{c}} = \left(\frac{\phi^*(\lambda, \mathfrak{f}\mathfrak{c}^{-1})}{\phi^*(1, \mathfrak{f}\mathfrak{c}^{-1})}\right)^e ,$$

where ϕ^* is the function defined in the preceding section. Then

$$K(\mathfrak{f}) = K(1)(\theta_{\lambda,\mathbb{Z}_K}) ,$$

and for any integral ideal $\mathfrak{c}$ coprime to f we have

$$\mathrm{Art}_{K(\mathfrak{f})/K}(\mathfrak{c})(\theta_{\lambda,\mathbb{Z}_K}) = \theta_{\lambda,\mathfrak{c}} .$$

Remark. Note that we really ask that $\mathfrak{c}$ be coprime to f and not only to the ideal $\mathfrak{f}$.

Schertz conjectures that this theorem is still valid if we only assume that the class of $\lambda\mathbb{Z}_K$ is not of order 1 in the ray class group, and we will make this conjecture in the sequel. Note that if by any chance it was false, the defining polynomial we would find at the end of Algorithm 6.3.27 below either would not have coefficients in $\mathbb{Z}_K$ or would not define $K(\mathfrak{f})/K(1)$, and all this can easily be checked (see also Exercise 26). On the other hand, this theorem is definitely not always true if $\lambda\mathbb{Z}_K$ is of order 1 (see Exercise 25).

The main problem with this theorem is to find a suitable pair (e, λ) satisfying the given conditions. To take an example, if $K = \mathbb{Q}(\sqrt{-163})$ and $\mathfrak{f} = \sqrt{-163}\mathbb{Z}_K$, the least possible value of e is 163, which will produce extremely large coefficients. To avoid this problem, we use the following theorem, which is essential for an algorithmic use of the theorem.

Theorem 6.3.24. *Let K be an imaginary quadratic field of discriminant D, let h be its class number, and let $\mathfrak{f}$ be a conductor of K different from $\mathbb{Z}_K$ (see Proposition 3.3.20 for all the possible conductors). Denote as usual by ζ_m a primitive mth root of unity, and for any prime number ℓ dividing D, denote by $\mathfrak{p}_\ell$ the unique ramified prime ideal above ℓ.*

(1) *If* $D = -3$ *and* $\mathfrak{f}$ *is in the following list, we give a relative defining polynomial* $P(X) \in K[X]$.
   a) If $\mathfrak{f} = f\mathbb{Z}_K$ *with* $f = 4$, $f = 5$, *or* $f = 7$, *then* $P(X) = \Phi_f(X)$, *the* f*th cyclotomic polynomial.*
   b) If $\mathfrak{f} = 3\mathfrak{p}_3 = \mathfrak{p}_3^3$, *then* $P(X) = X^3 + \zeta_3$.

(2) *If* $D = -4$ *and* $\mathfrak{f}$ *is in the following list, we give a relative defining polynomial* $P(X) \in K[X]$.
   a) If $\mathfrak{f} = f\mathbb{Z}_K$ *with* $f = 3$ *or* $f = 5$, *then* $P(X) = \Phi_f(X)$.
   b) If $\mathfrak{f} = 4\mathbb{Z}_K = \mathfrak{p}_2^4$, *then* $P(X) = X^2 + \zeta_4$.

(3) *If the pair* $(\mathfrak{f}, D)$ *is in the following list, then the ray class field* $K(\mathfrak{f})$ *is the compositum of the Hilbert class field* $K(1)$ *with the cyclotomic extension* $K(\zeta_m)$, *where* m *is given as follows (for completeness, we also give* $h(\mathfrak{f}) = [K(\mathfrak{f}) : K] = |Cl_{\mathfrak{f}}(K)|$).
   a) If $\mathfrak{f} = 2\mathbb{Z}_K$ *and* $D \equiv 8 \pmod{16}$, *then* $m = 4$ *and* $h(\mathfrak{f}) = 2h$.
   b) If $\mathfrak{f} = 3\mathbb{Z}_K$ *and* $D \equiv 1 \pmod{3}$, *then* $m = 3$ *and* $h(\mathfrak{f}) = 2h$.
   c) If $\mathfrak{f} = 4\mathbb{Z}_K$ *and* $D \equiv 1 \pmod{8}$, *then* $m = 4$ *and* $h(\mathfrak{f}) = 2h$.
   d) If $\mathfrak{f} = 6\mathbb{Z}_K$ *and* $D \equiv -8 \pmod{48}$, *then* $m = 12$ *and* $h(\mathfrak{f}) = 4h$.
   e) If $\mathfrak{f} = \mathfrak{p}_\ell$ *and* $\ell \mid D$ *with* $\ell > 3$, *then* $m = \ell$ *and* $h(\mathfrak{f}) = ((\ell - 1)/2)h$.
   f) If $\mathfrak{f} = 2\mathfrak{p}_\ell$ *and* $\ell \mid D$ *with* $\ell > 3$ *and* $D \equiv 8 \pmod{16}$, *then* $m = 4\ell$ *and* $h(\mathfrak{f}) = (\ell - 1)h$.

(4) *In all other cases, there exists* $\lambda \in \mathbb{Z}_K \setminus \{0\}$ *such that* $\mathcal{N}_{K/\mathbb{Q}}(\lambda) \equiv 1 \pmod{2f}$ *and* $\lambda\mathbb{Z}_K$ *not of order* 1 *in* $Cl_{\mathfrak{f}}(K)$, *where as above* f *is the positive generator of* $\mathfrak{f} \cap \mathbb{Z}$.

Proof. The special cases $D = -3$, $\mathfrak{f} = \mathfrak{p}_3^3$ and $D = -4$, $\mathfrak{f} = 4\mathbb{Z}_K$ are easily treated directly, so we exclude these cases. We first prove a lemma.

Lemma 6.3.25. *For any integer* m, *denote by* $\mathfrak{f}(m)$ *the conductor of the Abelian extension* $K(\zeta_m)/K$.

(1) *We have* $\mathfrak{f}(m) \mid m\mathbb{Z}_K$.
(2) *If* $K = \mathbb{Q}(\sqrt{D})$ *with* $D \equiv 8 \pmod{16}$, *then* $\mathfrak{f}(4) = 2\mathbb{Z}_K$; *while if* $D \equiv 12 \pmod{16}$, *then* $\mathfrak{f}(4) = \mathbb{Z}_K$.

Proof. (1) is nothing else but Proposition 3.5.5. For (2), we can, for example, use Hecke's Theorem 10.2.9 from which we borrow the notation. Indeed, since $[K(i) : K] = 2$, we have $\mathfrak{f}(4) = \mathfrak{d}(K(i)/K)$. Denote by $\mathfrak{p}$ the unique prime ideal above 2, so that $z(\mathfrak{p}, 2) = 5$. It is easily checked by an explicit calculation that the congruence $x^2 \equiv -1 \pmod{4\mathbb{Z}_K}$ is not soluble if $D \equiv 8 \pmod{16}$ and that it is soluble if $D \equiv 12 \pmod{16}$, while of course $x^2 \equiv -1 \pmod{2\mathbb{Z}_K}$ is always soluble. Since the largest value of $a < z(\mathfrak{p}, 2)$ that occurs in Hecke's theorem is odd, we deduce that $a = 3$ if $D \equiv 8 \pmod{16}$ and $a \geq 4$ if $D \equiv 12 \pmod{16}$, proving the lemma. □

Resuming the proof of the theorem, recall that by Proposition 3.3.21, if $\mathfrak{p} \mid \mathfrak{f}$ is above a prime number p that does not divide $h(\mathfrak{f})$, then $v_{\mathfrak{p}}(\mathfrak{f}) = 1$.

It is easily checked that the lemma together with this property imply that in all the cases mentioned in (1), (2), and (3) we have $K(\zeta_m) \subset K(\mathfrak{f})$. Since trivially $K(1) \subset K(\mathfrak{f})$, it follows that $K(\mathfrak{f})$ contains the compositum of $K(1)$ and $K(\zeta_m)$. To finish the proof, it is thus sufficient to determine the degree of this compositum or, equivalently, to compute the intersection of $K(1)$ and $K(\zeta_m)$. For $D = -3$ and $D = -4$ we have $K(1) = K$, so the result is trivial. In cases a), c), and d) of (3), $\mathfrak{p}_2$ is ramified in $K(\zeta_m)$ (we have proved this in the lemma for $D \equiv 8 \pmod{16}$, and the other case is treated similarly), and in cases b) and d), $\mathfrak{p}_3$ is ramified in $K(\zeta_m)$ by a similar reasoning. Thus, in the first four cases of (3) we have $K(1) \cap K(\zeta_m) = K$. On the other hand, in those cases it is easily checked that $[K(\zeta_m) : K] = [\mathbb{Q}(\zeta_m) : \mathbb{Q}] = \phi(m)$ and that $h(\mathfrak{f}) = \phi(m)h$ (using, for example, Proposition 3.2.3) so the result follows. For the last two cases, we must distinguish two possibilities.

• If $D = -\ell$, then $\ell \equiv 3 \pmod 4$ so $K \subset \mathbb{Q}(\zeta_\ell)$, hence $K(\zeta_\ell) = \mathbb{Q}(\zeta_\ell)$ and $[K(\zeta_\ell) : K] = (\ell - 1)/2$. If $\mathfrak{P}_\ell$ denotes the unique prime ideal of $K(\zeta_\ell)$ above $\mathfrak{p}_\ell$ (or above ℓ), we have

$$\ell - 1 = e(\mathfrak{P}_\ell/\ell) = e(\mathfrak{P}_\ell/\mathfrak{p}_\ell)e(\mathfrak{p}_\ell/\ell) = 2e(\mathfrak{P}_\ell/\mathfrak{p}_\ell) \ ,$$

hence $e(\mathfrak{P}_\ell/\mathfrak{p}_\ell) = (\ell-1)/2 = [K(\zeta_\ell) : K]$ so $\mathfrak{p}_\ell$ is totally ramified in $K(\zeta_\ell)/K$, from which it immediately follows that $K(1) \cap K(\zeta_\ell) = K$.

• If $\ell \mid D$ but $D \neq -\ell$, then I claim that the extension $K(\sqrt{\ell^*})/K$ is a subextension of $K(\zeta_\ell)/K$ that is unramified, where $\ell^* = (-1)^{(\ell-1)/2}\ell$. Indeed, the fact that it is a subextension is an elementary result on cyclotomic fields (it follows, for example, from the explicit evaluation of Gauss sums; see [Coh0, Exercise 16 of Chapter 1]). Thus the conductor of the extension $K(\sqrt{\ell^*})/K$ divides ℓ. On the other hand, we have $K(\sqrt{\ell^*}) = K(\sqrt{D/\ell^*})$, so the conductor also divides $4D/\ell^*$. But since $\ell \neq 2$, ℓ and $4D/\ell$ are coprime, the conductor is equal to $\mathbb{Z}_K$, and so the extension is unramified, as claimed.

On the other hand, it is easily checked that the extension $K(\zeta_\ell)/K(\sqrt{\ell^*})$ is totally ramified, from which we deduce that $K(1) \cap K(\zeta_\ell) = K(\sqrt{\ell^*})$ and no larger, proving e), and f) is treated similarly or follows from a) and e). This finishes the proof of (3).

(4). The proof of (4) is rather tedious, but simple. We first prove a lemma.

Lemma 6.3.26. *Let f be an integer such that $f \nmid 12$.*

(1) *If D is any integer, there exists t such that $(t^2 - D, 2f) = 1$ and $f \nmid 2t$.*
(2) *If $D \equiv 0 \pmod 4$ or $D \equiv 5 \pmod 8$ (in particular, if D is the discriminant of a quadratic field not congruent to 1 modulo 8), there exists t such that $t \equiv D \pmod 2$, $((t^2 - D)/4, 2f) = 1$, and $f \nmid t$.*

Proof. Let E_f be the set of t such that $0 \leq t < 2f$ and $(t^2 - D, 2f) = 1$. This condition means that $t \not\equiv D \pmod 2$ and that for each prime $p > 2$ dividing f, t must not be congruent to the square roots of D modulo p if they exist. It follows that

$$|E_f| = f \prod_{p|f,\ p>2} \left(1 - \frac{1 + \left(\frac{D}{p}\right)}{p}\right) .$$

If we write $f = \prod_{p|f} p^{v_p}$, we thus have

$$|E_f| \geq 2^{v_2} \prod_{p|f,\ p>2} p^{v_p - 1}(p-2) .$$

On the other hand, let F_f be the set of t such that $0 \leq t < 2f$ and $f \mid 2t$. Clearly, $|F_f| = 2$ if f is odd and $|F_f| = 4$ if f is even.

If there exists a prime $p \geq 5$ such that $p \mid f$, then $|E_f| \geq 3 \cdot 2^{v_2}$, hence $|E_f| > 2$ if $v_2 = 0$ and $|E_f| > 4$ if $v_2 \geq 1$; hence $|E_f \smallsetminus F_f| > 0$. A similar reasoning shows that the same conclusion still holds if $v_3 \geq 2$ or if $v_2 \geq 3$. Thus the only f for which (1) may be false are numbers of the form $f = 2^{v_2}3^{v_3}$ with $0 \leq v_2 \leq 2$ and $0 \leq v_3 \leq 1$; in other words, the divisors of 12. It is easy to check that the conclusion of (1) (and of (2)) is false if f is a divisor of 12.

(2) is proved in a similar manner. Here we denote by E_f the set of t such that $0 \leq t < 4f$, $t \equiv D \pmod 2$, and $((t^2 - D)/4, 2f) = 1$ and by F_f the set of t such that $0 \leq t < 4f$ and $f \mid t$, in other words $F_f = \{0, f, 2f, 3f\}$. We have

$$|E_f| = c_2 f \prod_{p|f,\ p>2} \left(1 - \frac{1 + \left(\frac{D}{p}\right)}{p}\right) ,$$

with $c_2 = 2$ if $D \equiv 5 \pmod 8$, $c_2 = 1$ if $D \equiv 0 \pmod 4$, and $c_2 = 0$ if $D \equiv 1 \pmod 8$. On the other hand, by looking only at congruences modulo powers of 2 we can easily compute an upper bound for $E_f \cap F_f$ and in particular prove that $|E_f \cap F_f| \leq 2$ unless we have $D \equiv 12 \pmod{16}$ and $f \equiv 0 \pmod 4$, in which case we can only say that $|E_f \cap F_f| \leq 4$.

Thus

$$|E_f \smallsetminus F_f| \geq c_2 \cdot 2^{v_2} \prod_{p|f\ p>2} p^{v_p - 1}(p-2) - a ,$$

where $a = 2$ if we are not in the special case $D \equiv 12 \pmod{16}$ and $f \equiv 0 \pmod 4$, and $a = 4$ otherwise. As in (1), if there exists a prime $p \geq 5$ dividing f or if $v_3 \geq 2$ or if $v_2 \geq 3$, we find that $|E_f \smallsetminus F_f| > 0$ if $D \not\equiv 1 \pmod 8$, since the case $a = 4$ can happen only if $v_2 \geq 2$, proving (2). □

Let us now prove statement (4) of Theorem 6.3.24. We assume that $D < -4$, leaving the cases $D = -3$ and $D = -4$ to the reader (Exercise 28). Let $\mathfrak{f}$ be an arbitrary conductor of K, and let f be the positive generator of $\mathfrak{f}$. Let α be an element of $\mathbb{Z}_K$ coprime to $2f$. If we choose $\lambda = \overline{\alpha}/\alpha = \overline{\alpha}^2/\mathcal{N}_{K/\mathbb{Q}}(\alpha)$ modulo $2f$, it is clear that $\mathcal{N}_{K/Q}(\lambda) = \overline{\lambda}\lambda \equiv 1 \pmod{2f}$, so such a λ satisfies the first condition of (4). Since $D < -4$, the condition that the class of $\lambda\mathbb{Z}_K$ is not of order 1 in the ray class group means that $\lambda \not\equiv \pm 1 \pmod{\mathfrak{f}}$ or,

equivalently, since α is coprime to $\mathfrak{f}$, that $\mathfrak{f} \nmid \alpha \pm \overline{\alpha}$. If $\alpha = (a + b\sqrt{D})/2$ with $a \equiv bD \pmod 2$, this condition means that $f \nmid a$ and that $\mathfrak{f} \nmid b\sqrt{D}$.

Assume first that $f \nmid 12$, $\mathfrak{f} \nmid \sqrt{D}$, and $D \not\equiv 1 \pmod 8$. Then Lemma 6.3.26 (2) tells us that there exists an integer t such that $t \equiv D \pmod 2$, $((t^2 - D)/4, 2f) = 1$ and $f \nmid t$, and this means that $\alpha = (t + \sqrt{D})/2$ is coprime to $2f$, and since $f \nmid t$ and $\mathfrak{f} \nmid \sqrt{D}$, $\lambda = \overline{\alpha}/\alpha$ modulo $2f$ is a suitable value.

Assume still that $f \nmid 12$ and $\mathfrak{f} \nmid \sqrt{D}$, but that $D \equiv 1 \pmod 8$. Then 2 is split, hence by Proposition 3.3.18 we also have $\mathfrak{f} \nmid 2\sqrt{D}$, so Lemma 6.3.26 (1) tells us that there exists an integer t such that $(t^2 - D, 2f) = 1$ and $f \nmid 2t$, and this means that $\alpha = t + \sqrt{D}$ is coprime to $2f$, and since $f \nmid 2t$ and $\mathfrak{f} \nmid 2\sqrt{D}$, $\lambda = \overline{\alpha}/\alpha$ modulo $2f$ is a suitable value.

If $f = 12$ and $D \equiv 0 \pmod 4$, we can take $t = 1$ if $D \equiv 0$ or $8 \pmod{12}$ and $t = 3$ if $D \equiv 4 \pmod{12}$, and we have again $(t^2 - D, 2f) = 1$ and $(f/2) \nmid t$, hence $\lambda = \overline{\alpha}/\alpha$ modulo $2f$ is suitable with $\alpha = t + \sqrt{D}$ as above. On the other hand, if $f = 12$ and $D \equiv 1 \pmod 4$, we can choose $\alpha = 4 + 3\sqrt{D}$ and clearly $\mathcal{N}_{K/\mathbb{Q}}(\alpha) = 16 - 9D$ is coprime to 24 and $(f/2) \nmid 8$, so $\lambda = \overline{\alpha}/\alpha$ is suitable.

We have thus proved that there exists a suitable value of λ satisfying the conditions of Theorem 6.3.24 (4) whenever f is not equal to 1, 2, 3, 4, or 6, and $\mathfrak{f} \nmid \sqrt{D}$. We must now treat these remaining cases.

Proposition 3.3.20 tells us that the only possible conductors $\mathfrak{f}$ with $f \mid 12$, $1 < f < 12$, are $\mathfrak{f} = 2\mathbb{Z}_K$ and $\mathfrak{f} = 6\mathbb{Z}_K$ when $D \not\equiv 1 \pmod 8$, $\mathfrak{f} = 3\mathbb{Z}_K$, and $\mathfrak{f} = 4\mathbb{Z}_K$. In each of these cases it is easy to find a suitable λ when the given congruence conditions on D are satisfied. Specifically:

(1) If $f = 2$, $D \not\equiv 1 \pmod 8$ and $D \not\equiv 8 \pmod{16}$, we take $\lambda = (t + \sqrt{D})/2$, where $t = 3$, 0, or 1 when $D \equiv 5$, 12, or 13 modulo 16.

(2) If $f = 3$ and $D \not\equiv 1 \pmod 3$, we take $\lambda = t + \sqrt{D}$, where $t = 1$, 3, 2, or 0 when $D \equiv 0$, 2, 3, or 5 modulo 6.

(3) If $f = 4$ and $D \not\equiv 1 \pmod 8$, we take $\lambda = 1 + \sqrt{D}$ if $D \equiv 8 \pmod{16}$, and $\lambda = (t + \sqrt{D})/2$, where $t = 3$, 4, 7, 5, 0, or 1 when $D \equiv 5$, 12, 13, 21, 28, or 29 modulo 32.

(4) If $f = 6$, $D \not\equiv 1 \pmod 8$, $D \not\equiv 8 \pmod{16}$ and $D \not\equiv 1 \pmod 3$, we take $\lambda = (t + \sqrt{D})/2$, where $t = 3$, 8, 11, 9, 0, or 1 when $D \equiv 5$, 12, 21, 29, 44, or 45 modulo 48.

Finally, assume that $f \nmid 12$ and that $\mathfrak{f} \mid \sqrt{D}$, which implies that $f \mid D$ when $D \equiv 1 \mod 4$ and $f \mid (D/2)$ when $D \equiv 0 \pmod 4$. In particular, f is squarefree except perhaps for a factor 4 when $D \equiv 8 \pmod{16}$. Since $f \nmid 12$ and is almost squarefree, f is divisible by some prime $\ell \geq 5$. By the Chinese remainder theorem we can find an integer x such that $x \equiv 1 \pmod \ell$ and $x \equiv -1 \pmod{2f/\ell}$. Since ℓ and $2f/\ell$ are coprime, we have $\mathcal{N}_{K/\mathbb{Q}}(x) = x^2 \equiv 1 \pmod{2f}$. Furthermore, $x \equiv -1 \pmod{\mathfrak{f}}$ is equivalent to $x \equiv -1 \pmod f$, which implies $x \equiv -1 \pmod \ell$, and this is impossible since $\ell > 2$. On the other hand, $x \equiv 1 \pmod{\mathfrak{f}}$ is equivalent to $x \equiv 1 \pmod f$,

hence to $-1 \equiv 1 \pmod{f/\ell}$, and this is possible if and only if $f \mid 2\ell$, in other words, since $f \nmid 12$, if and only if $f = \ell$ or $f = 2\ell$. Since $\ell \mid D$, ℓ is ramified, so $f = \ell$ and $\mathfrak{f} \mid \sqrt{D}$ imply that $\mathfrak{f} = \mathfrak{p}_\ell$, which is case e) of (3). If $f = 2\ell$, then since $\mathfrak{f} \mid \sqrt{D}$, D must be even, hence 2 is ramified, and $\mathfrak{f} = \mathfrak{p}_2^a \mathfrak{p}_\ell$ with $1 \le a \le 2$, and $a = 1$ is not possible because of Proposition 3.3.18, so $\mathfrak{f} = \mathfrak{p}_2^2 \mathfrak{p}_\ell = 2\mathfrak{p}_\ell$. Thus, if $\mathfrak{f}$ is not equal to $\mathfrak{p}_\ell$ or $2\mathfrak{p}_\ell$ with $\ell \mid D$, $\ell \ge 5$, we can take $\lambda = x$.

Finally, consider the case $\mathfrak{f} = 2\mathfrak{p}_\ell$, which implies $D \equiv 0 \pmod 4$ as we have just seen. The condition $D \equiv 8 \pmod{16}$ is case f) of (3), so assume that $D \equiv 12 \pmod{16}$. Choose x so that $x \equiv 1 \pmod \ell$ but $x \equiv 0 \pmod 4$. One easily checks that $\lambda = x + \sqrt{D}/2$ is suitable, and this finishes the proof of (4) and of Theorem 6.3.24. □

Remarks

(1) The compositum of $K(1)$ with $K(\zeta_m)$ should be computed using Algorithm 2.1.9 and not with Algorithm 2.1.8, since the coefficients obtained are much smaller.

(2) As seen in the proof, in the first four cases of (3), $K(1)$ and $K(\zeta_m)$ are linearly disjoint over K, so it is not necessary to perform a factoring step over K to obtain the compositum (step 6 of Algorithm 2.1.9). In the two remaining cases, the intersection of $K(1)$ and $K(\zeta_m)$ is equal to $K(\sqrt{\ell^*})$ (which is equal to K if and only if $D = -\ell$). As in the proof, we distinguish two different cases. If $D = -\ell$, the proof shows that the cyclotomic polynomial $\Phi_m(X)$ splits in $K[X]$ into a product of two conjugate irreducible factors of degree $\phi(m)/2$ (in other words, of degree $(\ell-1)/2$ if $m = \ell$ and of degree $\ell - 1$ if $m = 4\ell$), and so we compute the compositum using one of these factors as the defining polynomial for $K(\zeta_m)/K$. On the other hand, if $D \neq -\ell$, then $\Phi_m(X)$ is irreducible in $K[X]$. To compute the compositum, we may use two different methods. We can directly use Algorithm 2.1.9, and in step 6 we will find that the polynomial $R(X)$ (with the notation of that algorithm) splits in $K[X]$ into a product of two irreducible polynomials of degree $\phi(m)h/2$, and either one defines the desired defining polynomial. Since h may be large, this method involves factoring large-degree polynomials over K, however. An alternative and better method is to use our knowledge of the intersection $L = K(\sqrt{\ell^*})$ of $K(1)$ and $K(\zeta_m)$. We can compute a relative defining polynomial for $K(1)$ and for $K(\zeta_m)$ over L by factoring in $L[X]$ the corresponding polynomials. We then compute the compositum *over* L of the corresponding extensions, which will be linearly disjoint over L, and go back down to a relative defining polynomial over K by using Algorithm 2.1.11.

(3) As mentioned above, in [Sch4] and [Sch5] Schertz gives a more precise version of Theorem 6.3.23. Unfortunately, the exceptions treated in Theorem 6.3.24 remain exactly the same, hence the simpler version of his theorem is sufficient for algorithmic purposes. Note also that thanks to

Theorem 6.3.24, in the case where we will apply Theorem 6.3.23 we will take $e = 1$ so the coefficients will be as small as possible.

Thanks to this theorem, which complements Schertz's Theorem 6.3.23, we can now write an algorithm for computing ray class fields of imaginary quadratic fields. For simplicity, we will assume with Schertz that his theorem is still true if we remove the restriction that the class of $\lambda\mathbb{Z}_K$ is not of order 3 in the class group (we still assume that it is not of order 1, however; otherwise, the result is definitely not true in general). If by any chance this technical conjecture was false, the defining polynomial we would find at the end of Algorithm 6.3.27 below either would not have coefficients in $\mathbb{Z}_K$ or would not define $K(\mathfrak{f})/K(1)$, and this can easily be checked. In addition, if we do not want to depend on this assumption, it it is easily seen that it suffices to add a finite number of special cases in Theorem 6.3.24 (see Exercise 26).

Algorithm 6.3.27 (Ray Class Field Using $\sigma(z, L)$). Given an imaginary quadratic field $K = \mathbb{Q}(\sqrt{D})$ of discriminant D and a conductor $\mathfrak{f}$, this algorithm returns an irreducible polynomial $P \in K[X]$ such that the extension of K defined by a root of P is the full ray class field of K of conductor $\mathfrak{f}$. We let $(1, \omega)$ be an integral basis of K and f be the positive generator of $\mathfrak{f} \cap \mathbb{Z}$.

1. [Compute $K(1)$] Using Algorithm 6.3.8, compute a defining polynomial $P_1(X)$ for the Hilbert class field $K(1)$ of K. If $\mathfrak{f} = \mathbb{Z}_K$, output P_1 and terminate the algorithm.
2. [Special cases] If the pair $(D, \mathfrak{f})$ is in one of the special cases of Theorem 6.3.24 (1), (2), or (3), compute $P(X)$ either directly or as a compositum of $P_1(X)$ with a suitable cyclotomic polynomial using the remarks made after the theorem, and terminate the algorithm.
3. [Choose λ] For $a = 0, \ldots, 2f-1$ and $b = 0, \ldots, 2f-1$, set $\lambda \leftarrow a\omega + b$. If $\mathcal{N}_{K/\mathbb{Q}}(\lambda) = \lambda\overline{\lambda}$ is congruent to 1 modulo $2f$, test whether $\lambda - \varepsilon \in \mathfrak{f}$ for one of the 2 (if $D < -4$), 4 (if $D = -4$), or 6 (if $D = -3$) units ε of K. As soon as this is not the case, go to step 4 (by Theorem 6.3.24, such a λ will exist).
4. [Compute $Cl_{\mathfrak{f}}(K)$] Using Subalgorithm 6.3.28 below, compute a list $\mathcal{R}$ of primitive ideals coprime to f whose classes give the ray class group $Cl_{\mathfrak{f}}(K)$, and set $n \leftarrow |Cl_{\mathfrak{f}}(K)| = |\mathcal{R}|$.
5. [Initialize $P_2(X)$] Set $P_2(X) \leftarrow 1$ and $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{R}$).
6. [Loop in $Cl_{\mathfrak{f}}(K)$] Set $j \leftarrow j + 1$. If $j > n$, go to step 9. Otherwise, let $\mathfrak{c}$ be the jth element of $\mathcal{R}$.
7. [Compute lattice basis of $\mathfrak{f}\mathfrak{c}^{-1}$] Let (ω_1, ω_2) be an oriented $\mathbb{Z}$-basis of the ideal $\mathfrak{f}\mathfrak{c}^{-1}$ (for example, the HNF basis in reverse order to have the correct orientation, but any basis will do).
8. [Compute ϕ^* values] Using Algorithm 6.3.19 and the improvements suggested afterwards, compute $s \leftarrow \phi^*(\lambda, L)/\phi^*(1, L)$, where L is the complex lattice with basis (ω_1, ω_2), set $P_2(X) \leftarrow (X - s)P_2(X)$, and go to step 6.

9. [Round to algebraic] Write $P_2(X) = \sum_{0 \le i \le n} \beta_i X^i$, where the β_i are complex approximations to algebraic integers. For each i, use Subalgorithm 6.3.29 below to compute $\gamma_i \in \mathbb{Z}_K$, which closely approximate the β_i. If that algorithm fails for some i, the accuracy used was not sufficient. In that case, double the accuracy and go to step 5. Otherwise, set $P_2(X) \leftarrow \sum_{0 \le i \le n} \gamma_i X^i \in \mathbb{Z}_K[X]$.
10. [Terminate] Using [Coh0, Algorithm 3.6.4], check if $P_2(X)$ is irreducible in $K[X]$. If it is, output $P_2(X)$ as a defining polynomial for $K(\mathfrak{f})/K$. Otherwise, using Algorithm 2.1.8, compute the compositum of the extensions defined by $P_1(X)$ and $P_2(X)$, output the polynomial $P(X) \in K[X]$ defining this extension as a defining polynomial for $K(\mathfrak{f})/K$, and terminate the algorithm.

Remark. In step 3, we use a naive enumerative method to find a suitable value of λ. This is perfectly reasonable if f is rather small but may become slow if f is large, although the proof of Theorem 6.3.24 shows that many suitable λ will exist so the number of trials should be substantially lower than $4f^2$. An improvement on this naive method would be to use the proof of Theorem 6.3.24, which implies that in most cases we can take $\lambda = (t + \sqrt{D})/(t - \sqrt{D})$ for a suitable value of t and even gives a recipe for computing t. As in the proof, this involves looking at quite a number of special cases, so the details are left to the industrious reader (Exercise 29).

In step 4, we need to compute a list of primitive ideals coprime to $\mathcal{N}(\mathfrak{f})$ (or, equivalently, to f) whose classes enumerate the ray class group. There are at least two methods for doing this. The first one is to compute *some* integral ideals coprime to $\mathfrak{f}$ which enumerate the ray class group (this is easily done from the SNF), and then multiply these ideals by some $\alpha \equiv 1 \pmod{\mathfrak{f}}$ so as to make them primitive and coprime to f. Experiment shows that the ideals obtained in this manner are usually very large, so this method should not be used.

The second method is simply to compute the discrete logarithms in $Cl_{\mathfrak{f}}(K)$ of the primitive ideals of K coprime to f by increasing norm until all possible values have been obtained. Although this may take some time if $\mathfrak{f}$ and D are large, we must keep in mind that we perform this computation in order to compute a ray class field *defining polynomial*, hence that it is unreasonable to do this computation explicitly if $\mathfrak{f}$ and D are too large. We thus use the following algorithm.

Subalgorithm 6.3.28 (Primitive Representatives of Ray Class Group). Let K be an imaginary quadratic field, let $\mathfrak{f}$ be a conductor of K, let f be the positive generator of $\mathfrak{f} \cap \mathbb{Z}$, and let

$$Cl_{\mathfrak{f}}(K) = \bigoplus_{1 \le i \le r} (\mathbb{Z}/d_i\mathbb{Z})\overline{\mathfrak{a}_i}$$

be the SNF of the ray class group modulo $\mathfrak{f}$. This algorithm computes a list $\mathcal{R}$ of ideals $\mathfrak{c}$ that are primitive and coprime to f and whose classes form exactly the ray class group $Cl_{\mathfrak{f}}(K)$ (so that, in particular, $|\mathcal{R}| = |Cl_{\mathfrak{f}}(K)|$).

1. [Initialize] Set $\mathcal{R} \leftarrow \varnothing$, $n \leftarrow |Cl_{\mathfrak{f}}(K)|$, $B \leftarrow 10n$, $b \leftarrow 0$, $L \leftarrow (0,\ldots,0)$ (vector with n components initialized to 0), and $l \leftarrow 0$.
2. [Compute ideal list] Using Algorithm 2.3.23, compute the list $\mathcal{L}$ of ideals of K of norm less than or equal to B (where $\mathcal{L}_j$ contains the list of ideals of norm equal to j), and set $j \leftarrow b$.
3. [Loop on ideal norms] Set $j \leftarrow j+1$. If $j > B$, set $B \leftarrow 2B$, set $b \leftarrow B$, and go to step 2. Otherwise, if $\gcd(f,j) > 1$, go to step 3, and if not, set $S \leftarrow \mathcal{L}_j$, $s \leftarrow |S|$, and $k \leftarrow 0$ (k will be a pointer to the list S).
4. [Loop on elements of S] Set $k \leftarrow k+1$. If $k > s$, go to step 3. Otherwise, let $\mathfrak{c}$ be the kth element of S, given in HNF. If the bottom-right entry of the HNF of $\mathfrak{c}$ is not equal to 1 (in other words, if $\mathfrak{c}$ is not a primitive ideal), go to step 4.
5. [Compute discrete logarithm] (Here $\mathfrak{c}$ is a primitive ideal coprime to f.) Using Algorithm 4.3.2, compute the discrete logarithm $X = (x_1,\ldots,x_r)^t$ of the ideal $\mathfrak{c}$ in the ray class group $Cl_{\mathfrak{f}}(K)$, with $0 \le x_i < d_i$.
6. [Compute index and loop] Set

$$m \leftarrow 1 + x_r + d_r(x_{r-1} + d_{r-1}(\cdots x_2 + d_2 x_1)) \ .$$

   If the mth entry of the vector L is not equal to 0, go to step 4.
7. [Increase list $\mathcal{R}$] Set the mth entry of L equal to 1, set $\mathcal{R} \leftarrow \mathcal{R} \cup \{\mathfrak{c}\}$, and set $l \leftarrow l+1$. If $l < n$, go to step 4; otherwise, output the list $\mathcal{R}$ and terminate the subalgorithm.

Remark. There are still other methods for performing the above task. One is to take only prime ideals above prime numbers not dividing f. A better method is to compute a special list of ideals which directly computes primitive ideals coprime to f, and to compute discrete logarithms on the fly by computing only the discrete logarithms of the prime ideal factors. We leave the details to the reader (Exercise 27). Since the time for performing this computation is rather small compared to the time for computing the values of the function ϕ^* in Algorithm 6.3.27, it is probably not worthwhile to take the trouble of doing this.

Finally, we use the following simple algorithm for detecting whether a complex number is close to an element of $\mathbb{Z}_K$. Evidently this algorithm is much simpler than the corresponding algorithms in the real quadratic case given in Exercises 3 and 5.

Subalgorithm 6.3.29 (Is $\beta \in \mathbb{Z}_K$?). Let K be an imaginary quadratic field of discriminant D, and let β be given by a complex numerical approximation $\beta = x + iy$. This algorithm says whether or not it is plausible that $\beta \in \mathbb{Z}_K$; if it is, it outputs $(a,b) \in \mathbb{Z}^2$ such that $\beta = (a + b\sqrt{D})/2$.

1. [Compute a and b] Set $a_0 \leftarrow 2x$, $a \leftarrow \lfloor a_0 \rceil$, $b_0 \leftarrow 2y/\sqrt{|D|}$, and $b \leftarrow \lfloor b_0 \rceil$.

2. [Check and terminate] If $|a - a_0| > 10^{-5}$ or $|b - b_0| > 10^{-5}$ or $a - bD$ is odd, output a message saying that $\beta \notin \mathbb{Z}_K$. Otherwise, output (a, b) and terminate the subalgorithm.

6.4 Exercises for Chapter 6

1. In the situation of Section 6.1.2, show that we always have $L_S(0, \chi) = 0$ and that if χ is an even character we also have $L'_S(0, \chi) = 0$.
2. In the computation of Hilbert class fields of real quadratic fields, we need to compute $\phi(n) = E_1(cn)$ for a fixed constant c and $n = 1, \ldots, n = n_{\max}$, where $n_{\max}$ is a bound depending on the necessary accuracy.
   a) Let $u_k(n)$ be the value of the kth derivative of $\phi(x)$ at $x = n$. Show that for fixed n, $u_k(n)$ satisfies the following second-order linear recurrence relation:
   $$nu_k(n) = (cn + k)u_{k-1}(n) - cu_{k-2}(n) \ .$$
   b) Using Taylor's formula, deduce a method for computing $\phi(n - a)$ and $\phi'(n - a)$ from $\phi(n)$ and $\phi'(n)$ for any reasonably small a.
   c) Write an algorithm for simultaneously computing all the values $\phi(n)$ for $1 \leq n \leq n_{\max}$ by starting at $n_{\max}$ using the formula given in [Coh0, Proposition 5.6.12] and working backwards at variable speed. (When n is small, it will again be necessary to use [Coh0, Proposition 5.6.12].)
3. Let K be a real quadratic field, call τ_1 and τ_2 the two real embeddings of K, and let $\gamma \in \mathbb{Z}_K$. Assume given a good approximation β of $\tau_2(\gamma)$ (such that $|\tau_2(\gamma) - \beta| < \varepsilon$, say), and an upper bound B on $|\tau_1(\gamma)|$. As suggested in the text, write a naive algorithm that finds all possible γ (and that gives an error message if ε is not sufficiently small).
4. Complete the proof of Proposition 6.2.5 by showing that if $\varepsilon < 1/(3(B+1)(\sqrt{D}+1))$ we necessarily have $z = 0$. It will be important to use the following results on continued fractions.
   a) If $|\alpha - p/q| \leq 1/(2q^2)$, then p/q is a convergent to the continued fraction expansion of α.
   b) If p_n/q_n is the nth convergent to the continued fraction expansion of α and a_n is the nth partial quotient, then $|\alpha - p_n/q_n| \geq 1/((a_n + 2)q_n^2)$.
   c) The largest partial quotient occurring in the continued fraction expansion of $\omega = (\delta + \sqrt{D})/2$ with $\delta = D \bmod 2$ is equal to D.
5. Same exercise as Exercise 3, but using the Fincke–Pohst algorithm on the quadratic form
   $$q(x, y, z) = \left(\frac{B}{\varepsilon}\right)^2 (x + y\omega_2 - \beta z)^2 + (x + y\omega_1)^2 + B^2 z^2$$
   together with Proposition 6.2.5.
6. By taking for K suitable noncyclic complex cubic fields, show that when $K/\mathbb{Q}$ is not a cyclic extension, the conclusion of Theorem 6.2.8 can be true or false depending on K. In the specific case of cubic extensions, try to find a necessary and sufficient condition for the theorem's validity (I do not know the answer to this question, but apparently the conclusion of the theorem is, for example,

always false if the class number of K is equal to 2, and always true if it is equal to 3).

7. Modify the Polred algorithm ([Coh0, Algorithm 4.4.11]) so that it outputs as many polynomials as desired, and not exactly the degree of the polynomial to be reduced (you can, for example, use the Fincke–Pohst algorithm).
8. Assume the validity of the basic transformation formula $\eta(-1/\tau) = (\tau/i)^{1/2}\eta(\tau)$ with the principal part of the square root. Prove the complete transformation formula of η under $\mathrm{PSL}_2(\mathbb{Z})$ given in the text (this is tedious but not difficult).
9. Fill in the details of the proof of Algorithm 6.3.2.
10. Using the transformation formula for the η-function, prove the transformation formula for $g_{p,q}(\tau)$ given in Proposition 6.3.3.
11. Show that
$$g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a}^{-1}) = g_{\bar{\mathfrak{p}},\bar{\mathfrak{q}},e}(\mathfrak{a})$$
and that this is not always equal to $\overline{g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})}$.
12. Let K be an imaginary quadratic field of discriminant D, let p a prime such that $\left(\frac{D}{p}\right) \neq -1$, and let $\mathfrak{p}$ be a prime ideal of K above p. Denote as usual by $K(1)$ the Hilbert class field of K. Finally, let $\pi \in K(1)$ be such that $\mathfrak{p}\mathbb{Z}_{K(1)} = \pi\mathbb{Z}_{K(1)}$ (which must exist by Furtwängler's capitulation theorem). By using quotients of values of $\eta(z)$ at suitable points, show how one can compute π analytically without using the solution to the principal ideal problem in $K(1)$ (see [Sch1] for help).
13. Fill in the details of the proof of Proposition 6.3.11. In particular, prove the formula for the Weierstrass $\wp$ function used in the proof of (3) directly from the definition of $\wp$.
14. Using Proposition 6.3.11 to reduce to the fundamental domain, as well as Algorithm 6.3.13, write an algorithm for computing the Weierstrass ζ-function $\zeta(z, L)$.
15. Let $\eta(\tau) = q^{1/24}\prod_{n\geq 1}(1-q^n)$ be the Dedekind eta function.
    a) Show that $\eta'(\tau)/\eta(\tau) = 2i\pi E_2(\tau)/24$.
    b) Deduce from Corollary 6.3.12 that for any $\gamma = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{PSL}_2(\mathbb{Z})$, there exists a constant $v(\gamma) \in \mathbb{C}^*$ such that $\eta((a\tau+b)/(c\tau+d)) = v(\gamma)(c\tau+d)^{1/2}\eta(\tau)$.
    c) Show that $\eta(-1/\tau) = (\tau/i)^{1/2}\eta(\tau)$, where the square root is taken with positive real part.
    d) More generally, show that $v(\gamma)$ is a 24th root of unity and prove the formula for $v(\gamma)$ given before Algorithm 6.3.2 (this is long and tedious).
16. Fill in the details of the proof of Proposition 6.3.14, and using it, write an algorithm for computing $\sigma(z, L)$ by first reducing to the fundamental domain.
17. Give the details of the proof of Corollary 6.3.16.
18. By computing the product of two versions of the Jacobi triple-product identity (Proposition 6.3.15) and using the power series expansion of the eta function (Corollary 6.3.16), prove the *Jacobi quintuple-product identity*
$$\prod_{n\geq 1}(1-q^n u)\left(1-\frac{q^n}{u}\right)(1-q^n)(1-q^{2n-1}u^2)\left(1-\frac{q^{2n-1}}{u^2}\right)$$
$$= \sum_{k\geq 0}\left(\frac{1-k}{3}\right)q^{\frac{k(k+1)}{6}}\frac{u^{k+1/2}-u^{-k-1/2}}{u^{1/2}-u^{-1/2}},$$

where $\left(\frac{1-k}{3}\right)$ is the Legendre symbol equal to 1, 0, or -1 if k is congruent to 0, 1, or 2 modulo 3.

19. Show that
$$\prod_{n\geq 1,\ 4\nmid n}(1-q^n)^8+8q=q\left(\frac{\eta^8(\tau)}{\eta^8(4\tau)}+8\right)$$
is an even function of q (you may need some knowledge of modular forms to solve this exercise).

20. Give a good upper limit for the number of terms to be used in the series occurring in step 5 of Algorithm 6.3.19 for a given accuracy, depending only on $\mathrm{Im}(\tau)$ and on $|\mathrm{Im}(z_2)|$.

21. Modify Algorithm 6.3.19 so that it computes directly the ratio $\sigma(z,L)/\sigma(z',L)$ for two complex numbers z and z'.

22. Give an explicit upper bound for $|\phi^*(z,L)|$ on $\mathbb{C}$, which should even be independent of L.

23. Read [Sch1], [Sch2], and [Sch3], and write and implement the corresponding algorithm for computing ray class fields of imaginary quadratic fields using the Weierstrass $\wp$-function.

24. Modify Algorithm 6.3.27 so that it computes the ray class field corresponding to a congruence subgroup $(\mathfrak{f},C)$, in a way similar to Algorithm 6.3.8.

25. Give an example of an imaginary quadratic field K of discriminant $D<-4$, of a conductor $\mathfrak{f}$ (where as usual we set $\mathfrak{f}\cap\mathbb{Z}=f\mathbb{Z}$), and of a $\lambda\in\mathbb{Z}_K\setminus\mathbb{Z}$ such that $\mathcal{N}_{K/\mathbb{Q}}(\lambda)\equiv 1 \pmod{2f}$ but $K(\mathfrak{f})\neq K(1)(\theta_{\lambda,\mathbb{Z}_K})$.

26. Modify Theorem 6.3.24 so that its special cases also deal with $\lambda\mathbb{Z}_K$ of order 3 in the ray class group, and modify Algorithm 6.3.27 accordingly. (Hint: the special cases for $D<-4$ are $\mathfrak{f}=2\mathbb{Z}_K$ when $D\equiv 5 \pmod 8$ or $D\equiv 8 \pmod{16}$; $\mathfrak{f}=4\mathbb{Z}_K$ when $D\equiv 1 \pmod 4$; $\mathfrak{f}=\mathfrak{p}_3^a\mathfrak{p}_3'^b$ for $3\mathbb{Z}_K=\mathfrak{p}_3\mathfrak{p}_3'$ with $(a,b)=(2,0)$, $(1,1)$, $(0,2)$, $(2,1)$, $(1,2)$ and $(2,2)$ when $D\equiv 1 \pmod 3$; $\mathfrak{f}=\mathfrak{p}_3^2$ and $\mathfrak{f}=\mathfrak{p}_3^3$ when $D\equiv 0 \pmod 3$; $\mathfrak{f}=5\mathbb{Z}_K$ when $\left(\frac{D}{5}\right)=-1$; $\mathfrak{f}=\mathfrak{p}_7^a\mathfrak{p}_7'^b$ for $7\mathbb{Z}_K=\mathfrak{p}_7\mathfrak{p}_7'$ with $(a,b)=(1,0)$, $(0,1)$ and $(1,1)$ when $\left(\frac{D}{7}\right)=1$; $\mathfrak{f}=\mathfrak{p}_\ell$ if $\ell\mid D$ and $\ell>3$; $\mathfrak{f}=2\mathfrak{f}'$ if $\mathfrak{f}'$ is a conductor prime to 2 belonging to the above list and $D\equiv 5 \pmod 8$ or $D\equiv 8 \pmod{16}$; plus 18 conductors if $D=-3$ and 8 conductors if $D=-4$.)

27. Write an algorithm analogous to Algorithm 2.3.23, modified in two ways. First, it should directly compute only primitive ideals prime to f (in other words it should not compute all ideals first and then remove unsuitable ideals), and second it should compute the discrete logarithms of all the ideals by adding the discrete logarithms of the prime ideal factors along the way.

28. Complete the proof of Theorem 6.3.24 (4) for $D=-3$ and $D=-4$.

29. Write a detailed algorithm that uses the proof of Theorem 6.3.24 instead of a naive enumeration method to find a λ suitable for step 3 of Algorithm 6.3.27.

7. Variations on Class and Unit Groups

In Chapter 3 we studied variants of class and unit groups, the ray class groups $Cl_{\mathfrak{m}}(K)$, as well as the associated unit groups $U_{\mathfrak{m}}(K)$ of units multiplicatively congruent to 1 modulo $\mathfrak{m}$. The fundamental application of these notions through the deep theorems of class field theory is the construction of Abelian extensions.

In this chapter we will study other variations of class and unit groups. We will first study the generalization to the relative case of class groups, units, regulators, and related quantities. As in the preceding chapters dealing with relative information, the interest of this is twofold. On the one hand, a relative structure is almost always richer than an absolute one, and the absolute one can usually be deduced quite simply from the relative one. On the other hand, it generally gives much more efficient methods of computation. For example, the subexponential methods for class group and unit computations described in [Coh0, Chapter 6] become extremely costly for degrees larger than 20, say, even for moderate discriminants. If the field being considered is given as a relative extension, however (and for fields of large degree this is usually the case), we can use the methods described in this chapter to reach larger degrees.

The other variant of class and unit groups we will study are the S-class and unit groups. These have already been implicitly or explicitly used in the preceding chapters, for example, in the construction of Kummer extensions. As an application, we give results and algorithms for solving relative norm equations in many cases, mainly due to D. Simon.

Of course, all these variants may be combined, and we leave to the reader the study of these even more general groups.

7.1 Relative Class Groups

As usual, let K be a number field, which we take as base field, and let L be a relative extension of K. There are two maps that can be used to link objects attached to the field K with similar objects attached to the field L. One is the injection $i_{L/K}$ from K to L, and the second is the norm map $\mathcal{N}_{L/K}$ from L to K. Note that if $[K : L] = n$, we have $\mathcal{N}_{L/K} \circ i_{L/K} = [n]_K$, where for any integer m and Abelian group G, $[m]_G$ denotes the map that raises to the

mth power in G. We will thus have at least two different notions of relative objects, one obtained through the use of the map $i_{L/K}$, the other through the use of the map $\mathcal{N}_{L/K}$.

In this section, we want to define *relative class groups*. As mentioned above, there are (at least) two ways to do this.

7.1.1 Relative Class Group for $i_{L/K}$

We begin by the following definition.

Definition 7.1.1. (1) *A nonzero ideal I of L will be called* pseudo-principal *if there exist $\alpha \in L$ and an ideal $\mathfrak{a}$ of K such that $I = \alpha\mathfrak{a}\mathbb{Z}_L$. In a language that we have already used, I is pseudo-principal if it is generated by a single pseudo-element.*

(2) *Let $\mathcal{P}^*$ be the group of pseudo-principal ideals, and let $\mathcal{I}$ be the group of fractional ideals of L. The* relative pseudo-class group *is defined by*

$$Cl_i(L/K) = \mathcal{I}/\mathcal{P}^* \ ,$$

and we will call $h_i(L/K) = |Cl_i(L/K)|$ the relative pseudo-class number.

It is clear that $\mathcal{P}^*$ is a multiplicative subgroup of $\mathcal{I}$ containing the group $\mathcal{P}$ of fractional principal ideals of L, and hence $Cl_i(L/K)$ is a quotient of $Cl(L)$, and in particular is finite. Note also that if $h(K) = 1$, then $\mathcal{P}^* = \mathcal{P}$, and hence $Cl_i(L/K) = Cl(L)$. This can be made more precise as follows.

By abuse of notation, we will again denote by $i_{L/K}$ the natural map from $Cl(K)$ to $Cl(L)$ defined by $i_{L/K}(\overline{\mathfrak{a}}) = \overline{\mathfrak{a}\mathbb{Z}_L}$ for an ideal $\mathfrak{a}$ of K. It is clear that this is well-defined and is a group homomorphism from $Cl(K)$ to $Cl(L)$. Note that, contrary to the map $i_{L/K}$ defined on *elements*, it is *not* necessarily injective (see Exercise 1). It is then clear that

$$Cl_i(L/K) = Cl(L)/i_{L/K}(Cl(K)) = \mathrm{Coker}(i_{L/K}) \ .$$

Definition 7.1.2. *Let $i_{L/K}$ be the natural map from $Cl(K)$ to $Cl(L)$ as above. We will say that an ideal $\mathfrak{a}$ of K* capitulates *in L if $i_{L/K}(\mathfrak{a}) = \mathfrak{a}\mathbb{Z}_L$ is a principal ideal of $\mathbb{Z}_L$. The group $Cl_{i,L}(K) = \mathrm{Ker}(i_{L/K})$ of ideal classes of K which capitulate in L will be called the* capitulating subgroup *of $Cl(K)$ with respect to the extension L/K. Since the field L is usually understood, we will simply write $Cl_i(K)$ instead of $Cl_{i,L}(K)$.*

Note for future reference the following exact sequence, which is the exact translation of the definitions:

$$1 \longrightarrow Cl_i(K) \longrightarrow Cl(K) \overset{i_{L/K}}{\longrightarrow} Cl(L) \longrightarrow Cl_i(L/K) \longrightarrow 1 \ . \qquad (1)$$

Important example. Let L be the Hilbert class field of K (see Definition 3.5.4). Then a well-known theorem of Furtwängler tells us that *every* ideal

of K capitulates in L, so that $Cl_i(K) = Cl(K)$ in that case. Of course, this does *not* imply that $Cl(L)$ is trivial. In fact, a theorem of Golod–Shafarevitch says that there exist *infinite class field towers*, where each field in the tower is the Hilbert class field of the preceding one and, in particular, no field in the tower can have a trivial class group. By Theorem 2.5.1, in such a tower, the root discriminant stays constant, and in the totally complex case, the smallest known root discriminant is due to J. Martinet with the field $K = \mathbb{Q}(\cos(2\pi/11), \sqrt{-46})$ as base field, with root discriminant equal to $92.368\ldots$. For recent progress on this subject, where the root discriminant has been improved to $82.100\ldots$ by using towers of *tamely ramified* extensions instead of unramified extensions, see [Haj-Mai].

7.1.2 Relative Class Group for $\mathcal{N}_{L/K}$

We now come to the second notion of relative class group.

Let $\mathcal{N}_{L/K}$ be the norm map from fractional ideals of L to fractional ideals of K, which can also be considered as a group homomorphism from $Cl(L)$ to $Cl(K)$. Once again, this map is in general neither surjective nor injective (see Exercise 2).

Definition 7.1.3. *The* relative norm-class group *is the subgroup of* $Cl(L)$ *defined by*

$$Cl_N(L/K) = \mathrm{Ker}(\mathcal{N}_{L/K}) \ ,$$

and $h_N(L/K) = |Cl_N(L/K)|$ *will be called the* relative norm-class number.

Note that while the group $Cl_i(L/K)$ was a *quotient group* of $Cl(L)$, the group $Cl_N(L/K)$ is a *subgroup* of $Cl(L)$. Therefore, it is also finite and equal to $Cl(L)$ when $h(K) = 1$.

In a dual manner to the capitulating subgroup, we set the following definition.

Definition 7.1.4. *Let* $\mathcal{N}_{L/K}$ *be the norm map from* $Cl(L)$ *to* $Cl(K)$. *The group*

$$Cl_{N,L}(K) = \mathrm{Coker}(\mathcal{N}_{L/K}) = Cl(K)/\mathcal{N}_{L/K}(Cl(L))$$

will be called the norm-default quotient *of* $Cl(K)$ *in the extension* L/K *and will simply be denoted* $Cl_N(K)$ *when* L *is understood.*

As above, note for future reference the following exact sequence, which is the exact translation of the definitions:

$$1 \longrightarrow Cl_N(L/K) \longrightarrow Cl(L) \stackrel{\mathcal{N}_{L/K}}{\longrightarrow} Cl(K) \longrightarrow Cl_N(K) \longrightarrow 1 \ . \qquad (2)$$

A natural question to ask is whether the two notions of relative class group are related, apart from the fact that they are, respectively, a quotient

and subgroup of the absolute class group. This is essentially answered by Theorem 7.1.5 below and by the counterexamples following its proof.

Let $n = [L : K]$. As for elements, we have

$$\mathcal{N}_{L/K}\left(i_{L/K}\left(\overline{\mathfrak{a}}\right)\right) = \mathcal{N}_{L/K}\left(\overline{\mathfrak{a}\mathbb{Z}_L}\right) = \overline{\mathfrak{a}}^n \ ,$$

hence $\mathcal{N}_{L/K} \circ i_{L/K} = [n]_{Cl(K)}$. Furthermore, we have a natural map $\psi_{N,i}$ from $Cl_N(L/K)$ to $Cl_i(L/K)$ which sends an ideal class $\overline{I}$ of L belonging to $\mathrm{Ker}(\mathcal{N}_{L/K})$ to the class of $\overline{I}$ modulo $i_{L/K}(Cl(K))$. Finally, we have another map $\psi_{i,N}$ from $Cl_i(L/K)$ to $Cl_N(L/K)$ which sends an element $\overline{\overline{I}}$ of $Cl_i(L/K) = Cl(L)/i_{L/K}(Cl(K))$ to $\overline{I^{h(K)}}$ in $Cl_N(L/K)$, where $h(K)$ is the class number of K. Indeed, we have

$$\mathcal{N}_{L/K}(I^{h(K)}) = \mathcal{N}_{L/K}(I)^{h(K)} \ ,$$

which is a principal ideal of K, so $\overline{I^{h(K)}} \in \mathrm{Ker}(\mathcal{N}_{L/K})$. Furthermore, if we replace I by $\alpha\mathfrak{a}I$ for $\alpha \in L$ and $\mathfrak{a}$ an ideal of K, $\overline{I^{h(K)}}$ is multiplied by $\alpha^{h(K)}\mathfrak{a}^{h(K)}$, which is a principal ideal of L since $\mathfrak{a}^{h(K)}$ is a principal ideal of K; hence the map $\psi_{i,N}$ is well-defined and is evidently a group homomorphism.

By definition, we have $\psi_{i,N} \circ \psi_{N,i} = [h(K)]_{Cl_N(L/K)}$ and $\psi_{N,i} \circ \psi_{i,N} = [h(K)]_{Cl,(L/K)}$.

We have the following simple but important result.

Theorem 7.1.5. *Let L/K be a relative extension of degree n, and let $h(K)$ be the class number of K.*

(1) *If $(n, h(K)) = 1$ (for example, if K has class number 1), the natural map $\psi_{N,i}$ from $Cl_N(L/K)$ to $Cl_i(L/K)$ defined above is an isomorphism, and $Cl_N(K) = Cl_i(K) = \{1\}$.*

(2) *More generally, let p be a prime number, and for any group G, denote by G_p its p-Sylow subgroup. Then, if $p \nmid (n, h(K))$, the map $\psi_{N,i}$ induces an isomorphism between $Cl_N(L/K)_p$ and $Cl_i(L/K)_p$.*

(3) *Assume that $p \nmid h(K)$. Then*

$$Cl_N(L/K)_p \overset{\psi_{N,i}}{\simeq} Cl_i(L/K)_p \simeq Cl(L)_p \quad \textit{and}$$
$$Cl_N(K)_p = Cl_i(K)_p = \{1\} \ .$$

(4) *Assume that $p \nmid n$. Then*

$$Cl_N(L/K)_p \overset{\psi_{N,i}}{\simeq} Cl_i(L/K)_p \ ,$$
$$\mathcal{N}_{L/K}(Cl(L))_p \simeq \mathcal{N}_{L/K}(Cl(L)_p) \simeq (Cl(K)/Cl_i(K))_p \simeq Cl(K)_p \ , \ \textit{and}$$
$$Cl_N(K)_p = Cl_i(K)_p = \{1\} \ .$$

(5) *The exponents of the finite Abelian groups $Cl_i(K)$ and $Cl_N(K)$ both divide $(n, h(K))$.*

Another way of stating part (2) of this theorem is to say that the only primes for which the two class groups can differ are the primes dividing $(n, h(K))$.

Proof. It is clear that statements (3) and (4) imply statement (2), which implies statement (1). Furthermore, the formulas $\mathcal{N}_{L/K}\left(i_{L/K}\left(\overline{\mathfrak{a}}\right)\right) = \overline{\mathfrak{a}}^n$ and $\overline{\mathfrak{a}}^{h(K)} = \overline{1}$ in $Cl(K)$ immediately imply (5). Thus we need only to prove (3) and (4). Recall from Proposition 4.1.17 that in the category of finite Abelian groups, taking p-Sylow subgroups is an exact functor. Applying this to the two exact sequences (1) and (2), which define the two notions of class groups, we obtain the following two exact sequences:

$$1 \longrightarrow Cl_i(K)_p \longrightarrow Cl(K)_p \overset{i_{L/K,p}}{\longrightarrow} Cl(L)_p \longrightarrow Cl_i(L/K)_p \longrightarrow 1$$

and

$$1 \longrightarrow Cl_N(L/K)_p \longrightarrow Cl(L)_p \overset{\mathcal{N}_{L/K,p}}{\longrightarrow} Cl(K)_p \longrightarrow Cl_N(K)_p \longrightarrow 1 \ . \quad (3)$$

Assume first that p is a prime such that $p \nmid h(K)$, in other words that $Cl(K)_p = \{1\}$. The two exact sequences above imply immediately that $Cl_i(K)_p = Cl_N(K)_p = \{1\}$ and that $Cl_i(L/K)_p \simeq Cl(L)_p \simeq Cl_N(L/K)$, proving (3).

Assume now that p is a prime such that $p \nmid n$. As already mentioned, $\mathcal{N}_{L/K} \circ i_{L/K} = [n]_{Cl(K)}$; hence by restriction, this gives the map $[n]_{Cl(K)_p}$ from $Cl(K)_p$ to itself. Since $Cl(K)_p$ is a p-group, it follows that the map $[n]_{Cl(K)_p}$ is a bijection of $Cl(K)_p$ onto itself (its inverse is the map $[n']_{Cl(K)_p}$ for any integer n' such that $nn' \equiv 1 \pmod{|Cl(K)_p|}$). Therefore, with evident notation the map $i_{L/K,p}$ is injective, and the map $\mathcal{N}_{L/K,p}$ is surjective, so $Cl_i(K)_p = Cl_N(K)_p = \{1\}$, as claimed. By the exactness of $_p$ for finite Abelian groups, it also follows that

$$\mathcal{N}_{L/K}(Cl(L))_p \simeq \mathcal{N}_{L/K}(Cl(L)_p) \simeq (Cl(K)/Cl_i(K))_p \simeq Cl(K)_p \ .$$

The careful reader will note that the first two groups above are not the same, although they are isomorphic in this case.

We now show that $\psi_{N,i}$ is a bijection on the p-parts. Assume first that I is an ideal of L such that $\psi_{N,i}(\overline{I}) = 1_{Cl_i(L/K)}$ and that the class of I is in $Cl_N(L/K)_p$. This means that there exists an ideal $\mathfrak{a}$ of K and $\alpha \in L^*$ such that $I = \mathfrak{a}\alpha\mathbb{Z}_L$. Thus $\mathcal{N}_{L/K}(I) = \mathfrak{a}^n \mathcal{N}_{L/K}(\alpha)$. On the other hand, since the class of I is in $Cl_N(L/K)$, we have $\mathcal{N}_{L/K}(I) = \beta\mathbb{Z}_K$ for some $\beta \in K^*$. Thus $\mathfrak{a}^n$ is a principal ideal of K, and in particular $\mathfrak{a}^n\mathbb{Z}_L$ is a principal ideal of L. However, since $I \in Cl_N(L/K)_p$, I^p is a principal ideal of L so $\mathfrak{a}^p\mathbb{Z}_L$ is a principal ideal of L; since $(n,p) = 1$, it follows that $\mathfrak{a}\mathbb{Z}_L$, hence also $I = \mathfrak{a}\alpha\mathbb{Z}_L$, is a principal ideal of L, thus proving that $\psi_{N,i}$ is injective.

To prove that $\psi_{N,i}$ is surjective, it is now sufficient to prove that the groups $Cl_N(L/K)_p$ and $Cl_i(L/K)_p$ have the same cardinality. This follows immediately from the exact sequences (3), which give

$$|Cl_i(L/K)_p| = |Cl(L)_p| \,/\, |Cl(K)_p| = |Cl_N(L/K)_p| \ ,$$

thus finishing the proof of the theorem. □

Remarks

(1) It is easy to give examples for which $Cl_i(L/K)$ and $Cl_N(L/K)$ differ. The following example is due to D. Simon. Take $K = \mathbb{Q}(y)$ with $y^2 + 30 = 0$ and $L = K(\sqrt{y})$. Then one can show that $Cl(K) \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, $Cl(L) \simeq \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, $Cl_i(L/K) \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, and $Cl_N(L/K) \simeq \mathbb{Z}/2\mathbb{Z}$.

(2) The theory of the map $i_{L/K}$, and in particular of capitulation, is not well understood apart from the beautiful theorem of Furtwängler already mentioned and more recent generalizations (see [Suz]). To the contrary, the theory of the map $\mathcal{N}_{L/K}$ is quite well understood, and is essentially class field theory.

(3) In algorithmic practice, however, as we shall see in Section 7.3.3, the natural group that is computed is the group $Cl_i(L/K)$ and the associated unit group $U_i(L/K)$ that we will define below in Definition 7.2.2, and not the group $Cl_N(L/K)$, which can of course also be computed if desired, but less naturally.

7.2 Relative Units and Regulators

Since we know that units are intimately linked to class groups, we will now study the notion of relative units and relative regulators. As before, there are two such notions, one linked to the map $i_{L/K}$, the other to the map $\mathcal{N}_{L/K}$.

7.2.1 Relative Units and Regulators for $i_{L/K}$

Recall that the absolute unit group $U(K)$ could be defined by the following exact sequence:

$$1 \longrightarrow U(K) \longrightarrow K^* \xrightarrow{\phi} \mathcal{P} \longrightarrow 1 \ ,$$

where $\phi(\alpha) = \alpha\mathbb{Z}_K$. Hence, we will define the first notion of relative units in a similar way. Let $\mathcal{P}^*$ be the group of pseudo-principal ideals in the sense of Definition 7.1.1. We first must generalize the map ϕ. For this, we introduce the following definition.

Definition and Proposition 7.2.1. *Let L/K be a relative extension of number fields, and denote by $\mathcal{I}(K)$ the group of fractional ideals of K. The* extended multiplicative group L^{*K} *is the quotient group of $\mathcal{I}(K) \times L^*$ by the equivalence relation defined by*

$$(\mathfrak{a},\alpha)\ \mathcal{R}\ (\mathfrak{b},\beta) \iff \exists \gamma \in K^*,\ \mathfrak{b} = \mathfrak{a}\gamma \text{ and } \beta = \frac{\alpha}{\gamma}$$

or, equivalently, by

$$(\mathfrak{a},\alpha)\ \mathcal{R}\ (\mathfrak{b},\beta) \iff \alpha\mathfrak{a} = \beta\mathfrak{b}\ ,$$

where $\alpha\mathfrak{a}$ and $\beta\mathfrak{b}$ are considered as sub-$\mathbb{Z}_K$-modules of L.

Proof. We must show that the two definitions of the equivalence relation $\mathcal{R}$ are the same. Indeed, it is clear that $\mathfrak{b} = \mathfrak{a}\gamma$ and $\beta = \alpha/\gamma$ implies $\alpha\mathfrak{a} = \beta\mathfrak{b}$. Conversely, assume that this is true. Then $\mathfrak{a}\mathfrak{b}^{-1} = (\beta/\alpha)\mathbb{Z}_K$, hence $\beta/\alpha \in \mathfrak{a}\mathfrak{b}^{-1} \subset K$, so we can take $\gamma = \alpha/\beta$. □

Remark. In the case where $h(K) = 1$, and in particular in the absolute case $K = \mathbb{Q}$, it is clear that in every equivalence class, there is a representative of the form $(\mathbb{Z}_K, \alpha)$, where α is defined modulo units of K, hence $L^{*K} \simeq L^*/U(K)$. More generally, as a set (but of course not as a group), L^{*K} can be considered as the union of $h(K)$ copies of $L^*/U(K)$.

We can now define in a natural way the map ϕ from L^{*K} to $\mathcal{P}^*$ by setting

$$\phi\big(\overline{(\mathfrak{a},\alpha)}\big) = \alpha\mathfrak{a}\mathbb{Z}_L\ .$$

By definition of the equivalence relation $\mathcal{R}$, this map is well-defined and is a group homomorphism. Guided by the absolute case, we will define the group of *relative pseudo-units* as the kernel of ϕ.

Definition 7.2.2. *We say that an element $\overline{(\mathfrak{a},\alpha)}$ of L^{*K} is a* relative pseudo-unit *in L/K if $\alpha\mathfrak{a}\mathbb{Z}_L = \mathbb{Z}_L$. The set of relative pseudo-units in L/K is a multiplicative group denoted by $U_i(L/K)$.*

Proposition 7.2.3. (1) *We have the following exact sequence*

$$1 \longrightarrow U(K) \xrightarrow{i_{L/K}} U(L) \xrightarrow{j} U_i(L/K) \xrightarrow{\pi} Cl_i(K) \longrightarrow 1\ , \qquad (4)\ ,$$

where $j(\alpha) = \overline{(\mathbb{Z}_K,\alpha)}$ and $\pi\big(\overline{(\mathfrak{a},\alpha)}\big) = \bar{\mathfrak{a}}$ in $Cl_i(K)$.

(2) *The group $U_i(L/K)$ is a finitely generated Abelian group of rank $r(L) - r(K)$, where $r(N)$ denotes the unit rank of a number field N.*

Proof. (1). If $\alpha \in U(L)$, then $\alpha\mathbb{Z}_L = \mathbb{Z}_L$, so $\alpha\mathbb{Z}_K\mathbb{Z}_L = \mathbb{Z}_L$. Hence $j(\alpha) \in U_i(L/K)$ is well-defined and is a group homomorphism. If $\overline{(\mathfrak{a},\alpha)} = \overline{(\mathfrak{b},\beta)}$ in $U_i(L/K)$, there exists $\gamma \in K^*$ such that $\mathfrak{b} = \gamma\mathfrak{a}$ and $\beta = \alpha/\gamma$, and furthermore $\alpha\mathfrak{a}\mathbb{Z}_L = \mathbb{Z}_L$. It follows that $\mathfrak{a}$ and $\mathfrak{b}$ are in the same ideal class in $Cl(K)$ and that the class of $\mathfrak{a}$ is in the capitulating class group, so the map π is well-defined and is a group homomorphism.

Let us show exactness. Exactness at $U(K)$ follows from the fact that $i_{L/K}$ is injective *on elements* (recall that it was not injective on ideal classes). The

kernel of j is the set of $\alpha \in U(L)$ such that $\overline{(\mathbb{Z}_K, \alpha)}$ is the unit element of the group $U_i(L/K)$, which is the class of $\overline{(\mathbb{Z}_K, 1)}$. By definition of the equivalence relation, this means that there exists $u \in K^*$ such that $\mathbb{Z}_K = u\mathbb{Z}_K$ and $\alpha = u \cdot 1$, in other words that $\alpha = u$ is a unit of $\mathbb{Z}_K$, so the kernel of i is the group $U(K)$, thus showing exactness at $U(L)$.

The kernel of π is the set of $\overline{(\mathfrak{a}, \alpha)}$ such that $\mathfrak{a}$ is a principal ideal in $\mathbb{Z}_K$, hence such that there exists $\gamma \in K^*$ with $\mathfrak{a} = \gamma\mathbb{Z}_K$, and also $\alpha\mathfrak{a}\mathbb{Z}_L = \mathbb{Z}_L$. Hence, $\overline{(\mathfrak{a}, \alpha)} = \overline{(\mathbb{Z}_K, \alpha\gamma)}$, and $\mathbb{Z}_L = \alpha\mathfrak{a}\mathbb{Z}_L = \alpha\gamma\mathbb{Z}_L$, so $\alpha\gamma \in U(L)$ and the kernel of π is thus equal to $j(U(L))$, proving exactness at $U_i(L/K)$.

Finally, we must show that π is surjective. But if $\overline{\mathfrak{a}}$ is an ideal class in $Cl_i(K)$, there exists $\alpha \in L^*$ such that $\mathfrak{a}\mathbb{Z}_L = \alpha\mathbb{Z}_L$, and it is clear that $\overline{(\mathfrak{a}, 1/\alpha)} \in U_i(L/K)$ and satisfies $\pi\big(\overline{(\mathfrak{a}, 1/\alpha)}\big) = \overline{\mathfrak{a}}$, thus finishing the proof of (1).

(2). From (1), we deduce that the following short exact sequence is exact:

$$1 \longrightarrow U(L)/i_{L/K}(U(K)) \longrightarrow U_i(L/K) \xrightarrow{\pi} Cl_i(K) \longrightarrow 1 \ . \tag{5}$$

Since $Cl_i(K)$ is a finite Abelian group and $i_{L/K}$ is injective, it follows that the rank of $U_i(L/K)$ is equal to $r(L) - r(K)$, as claimed. □

Corollary 7.2.4. *Assume that* $(n, h(K)) = 1$. *Then the group* $U_i(L/K)$ *is isomorphic to the quotient group* $U(L)/i_{L/K}(U(K))$.

Proof. By Theorem 7.1.5 (1), when $(n, h(K)) = 1$ we have $Cl_i(K) = \{1\}$, and so the corollary follows immediately from Proposition 7.2.3. □

Corollary 7.2.5. *There exists a six-term exact sequence*

$$1 \longrightarrow U(K) \xrightarrow{i_{L/K}} U(L) \xrightarrow{j} U_i(L/K) \longrightarrow Cl(K) \xrightarrow{i_{L/K}} Cl(L) \longrightarrow Cl_i(L/K) \longrightarrow 1 \ . \tag{6}$$

Proof. This is a trivial consequence of exact sequences (1) and (4). □

This six-term exact sequence is reminiscent of exact sequences occurring in homology/cohomology theories. This is certainly not a coincidence, since one can interpret the class group $Cl(K)$ as the torsion part of the algebraic K-group $K_0(\mathbb{Z}_K)$ and the unit group $U(K)$ as the algebraic K-group $K_1(\mathbb{Z}_K)$ (see [Ros]).

Determining theoretically the torsion subgroup of $U_i(L/K)$ is much more difficult for several reasons, not the least of which being that in general the exact sequence (4) is not split. Before stating a result in this direction, we need the following generalization of [Coh0, Theorem 2.4.12].

Proposition 7.2.6. *Let A be an $m \times n$ integer matrix of rank n, hence with $m \geq n$. There exist two unimodular matrices U and V of size $n \times n$ and $m \times m$, respectively, such that*

$$B = VAU = \begin{pmatrix} d_1 & 0 & \dots & 0 \\ 0 & d_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & d_n \\ 0 & \dots & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & \dots & 0 \end{pmatrix}$$

with $d_{i-1} \mid d_i$ for $i > 1$.

Proof. As usual, denote by A^t the transpose of A, which is an $n \times m$ matrix of rank n. By the theorem on the HNF ([Coh0, Theorem 2.4.4]), there exists a unimodular $m \times m$ matrix V_1 and an $n \times n$ matrix H in HNF such that $A^t V_1 = (0|H)$. Calling V_2 the matrix obtained by putting the first $m - n$ columns of V_1 after the last n columns, the matrix V_2 is still unimodular and $A^t V_2 = (H|0)$, hence $V_2^t A = \binom{H^t}{0}$. Let V_3 and U be $n \times n$ unimodular matrices such that $V_3 H^t U = D = \mathrm{diag}(d_i)$ is in SNF. If we set

$$V = \begin{pmatrix} V_3 & 0 \\ 0 & I_{m-n} \end{pmatrix} V_2^t \ ,$$

then $VAU = \binom{D}{0} = B$ satisfies the conditions of the proposition. Note of course that the above proof gives an algorithm to compute U, V, and B. □

We can now give some indication on the structure of $U(L)/i_{L/K}(U(K))$. To simplify notation, we identify K with a subfield of L; in other words, we omit the map $i_{L/K}$.

Proposition 7.2.7. *There exists a generator ζ of the group of roots of unity of L of order $w(L)$, a system of fundamental units $(\eta_i)_{1 \leq i \leq r(L)}$ of L, an integer $w(L/K)$, and integers d_i and e_i for $1 \leq i \leq r(K)$ such that:*

(1) *$w(L/K) = w(L)/w(K)$ and $\zeta^{w(L/K)}$ is a generator of the group of roots of unity of K;*
(2) *$(\zeta^{e_i} \eta_i^{d_i})_{1 \leq i \leq r(K)}$ is a system of fundamental units of K;*
(3) *for each $i > 1$, we have $d_i \mid d_{i-1}$.*

Proof. Let ζ (resp., ζ') be a generator of the group of roots of unity of L (resp., K), let $(\eta_i')_{1 \leq i \leq r(L)}$ (resp., $(\varepsilon_i')_{1 \leq i \leq r(K)}$) be an arbitrary system of fundamental units of L (resp., K). Since we identify K with a subfield of L,

there exists an integral $(r(L)+1) \times r(K)$ matrix $A = (a_{i,j})$ such that for every j with $1 \le j \le r(K)$ we have

$$\varepsilon'_j = \prod_{0 \le i \le r(L)} \eta_i'^{a_{i,j}} ,$$

where we set $\eta'_0 = \zeta$ to simplify notation.

Write $A = \begin{pmatrix} A_0 \\ A_1 \end{pmatrix}$, where A_0 is a $1 \times r(K)$ matrix and A_1 is an $r(L) \times r(K)$ matrix. By Proposition 7.2.6 applied to the matrix A_1, we can find unimodular matrices U and V_1 such that $B_1 = V_1 A_1 U = \begin{pmatrix} D \\ 0 \end{pmatrix}$, where $D = \mathrm{diag}(d_1, \ldots, d_n)$ and $d_{i-1} \mid d_i$ for $i > 1$. If we set

$$V = \begin{pmatrix} 1 & 0 \\ 0 & V_1 \end{pmatrix} ,$$

then

$$B = VAU = \begin{pmatrix} A_0 U \\ D \\ 0 \end{pmatrix} .$$

Since

$$[\varepsilon'_1, \ldots, \varepsilon'_{r(K)}] = [\zeta, \ldots, \eta'_{r(L)}]A ,$$

we have

$$[\varepsilon_1, \ldots, \varepsilon_{r(K)}] = [\zeta, \eta_1, \ldots, \eta_{r(L)}]B ,$$

where

$$[\varepsilon_1, \ldots, \varepsilon_{r(K)}] = [\varepsilon'_1, \ldots, \varepsilon'_{r(K)}]U$$

and

$$[\zeta, \eta_1, \ldots, \eta_{r(L)}] = [\zeta, \eta'_1, \ldots, \eta'_{r(L)}]V^{-1} .$$

Since the matrix U is unimodular, $(\varepsilon_1, \ldots, \varepsilon_{r(K)})$ is a system of fundamental units of K, and similarly since V_1 is unimodular, $(\eta_1, \ldots, \eta_{r(L)})$ is a system of fundamental units of L.

Note finally that $\zeta' = \zeta^m$ for some integer m, and it is clear that $\zeta^{w(L/K)}$ with $w(L/K) = (m, w(L))$ is also a generator of the group of roots of unity of K and that $w(L/K) \mid w(L)$ (see Exercise 4). Since this group is of order $w(K)$, we have $w(L/K) = w(L)/w(K)$, thus finishing the proof of the proposition if we write $A_0 U = (e_1, \ldots, e_{r(K)})$. □

Remark. It could be hoped that one can take $e_i = 0$ for all i, thus giving the cleaner formulas $\varepsilon_i = \eta_i^{d_i}$ for $1 \le i \le r(K)$. This is not possible in general (see Exercise 5).

Corollary 7.2.8. *Keep the notation of the above proposition, denote as usual by $\mu(L)$ the group of roots of unity of L, and set $w(L/K) = w(L)/w(K)$. Then:*

(1)
$$\frac{U(L)}{\mu(L)\cdot i_{L/K}(U(K))} = \bigoplus_{1\le i\le r(K)} (\mathbb{Z}/d_i\mathbb{Z})\overline{\eta_i} \oplus \bigoplus_{r(K)<i\le r(L)} \mathbb{Z}\overline{\eta_i}\ ,$$

where $\overline{\eta}$ denotes the class of η modulo $\mu(L)\cdot i_{L/K}(U(K))$.

(2) *The integers d_i are unique, in other words they do not depend on the particular choice made for the η_i satisfying Proposition 7.2.7.*

(3) *For all i we have $d_i \mid n = [L:K]$.*

(4) *We have $n \equiv w(K) \pmod{w(L/K)}$ and in particular $(w(L/K), w(K)) \mid n$.*

(5) *We can choose the e_i in any fixed residue set modulo $(w(L/K), d_i)$, for example, such that $-(w(L/K), d_i)/2 < e_i \le (w(L/K), d_i)/2$.*

Proof. Statement (1) follows immediately from Proposition 7.2.7. Note that because of the presence of the integers e_i, we cannot give such a clean formula for $U(L)/i_{L/K}(U(K))$.

By (1), $\mathrm{diag}(d_i)$ is the Smith normal form of the torsion submodule of $U(L)/(\mu(L)\cdot i_{L/K}(U(K)))$. Hence the uniqueness of the d_i follows from the uniqueness of the SNF, proving (2).

By definition of d_i and e_i, we have $\varepsilon_i = \zeta^{e_i}\eta_i^{d_i}$, where the (ε_i) form a system of fundamental units of K. Taking the norm from L to K, we obtain

$$\varepsilon_i^n = \zeta^{ke_i}\,\mathcal{N}_{L/K}(\eta_i)^{d_i}$$

for some integer k such that $\mathcal{N}_{L/K}(\zeta) = \zeta^k \in K$. Since the ε_i form a system of fundamental units and $\zeta^{w(L/K)}$ generates $\mu(K)$, there exist integers x_j such that

$$\mathcal{N}_{L/K}(\eta_i) = \zeta^{x_0 w(L/K)} \prod_j \varepsilon_j^{x_j}\ .$$

Since the ε_j are independent, if we replace in the formula for ε_i^n, we obtain $x_j = 0$ for $j \ne i$ and $j \ne 0$, $d_i x_i = n$, and $ke_i + d_i x_0 w(L/K) \equiv 0 \pmod{w(L)}$. In particular, we see that $d_i \mid n$, proving (3).

Since $\mathcal{N}_{L/K}(\zeta)$ is a root of unity of K, there exists an integer m such that $\mathcal{N}_{L/K}(\zeta) = \zeta^{w(L/K)m}$. On the other hand, since $\zeta^{w(L/K)} \in K$, we have $\mathcal{N}_{L/K}(\zeta^{w(L/K)}) = \zeta^{w(L/K)n}$, hence

$$\zeta^{w(L/K)(w(L/K)m-n)} = 1$$

so $w(L) \mid w(L/K)(w(L/K)m - n)$ or, equivalently, $w(K) \mid w(L/K)m - n$. It follows in particular that $(w(K), w(L/K)) \mid n$, proving (4).

For any integers u_i and v_i we can write

$$\varepsilon_i\zeta^{u_i w(L/K)} = \zeta^{e_i + u_i w(L/K) - v_i d_i}(\eta_i\zeta^{v_i})^{d_i}\ .$$

Since the $(\varepsilon_i \zeta^{u_i w(L/K)})$ (resp., the $(\eta_i \zeta^{v_i})$) still form a fundamental system of units of K (resp., of L) with the same d_i, it follows that we can freely replace e_i by $e_i + u_i w(L/K) - v_i d_i$ or, equivalently, by $e_i + k(w(L/K), d_i)$ for any integer k, and in particular we can choose e_i to be in any fixed residue set modulo $(w(L/K), d_i)$, proving (5). □

Remark. We still have much additional freedom for modifying the e_i. It is not clear, however, if there is some canonical choice of the e_i analogous to a Hermite normal form. Since we do not know of such a choice, please note that contrary to the d_i, the e_i are *not* invariants of the extension L/K since they depend on the specific choices of generating units in K and L.

We can now define the notion of relative regulator for the map $i_{L/K}$.

Definition 7.2.9. *Let L/K be a relative extension of number fields.*

(1) *We define $U_{i,L}(K)$ (abbreviated to $U_i(K)$) to be the kernel of the map $i_{L/K}$ from $U(K)$ to $U(L)$.*
(2) *We define the* relative regulator $R_i(L/K)$ *associated to the map $i_{L/K}$ by the formula*

$$R_i(L/K) = \frac{1}{|U_i(K)|} \frac{R(L)}{R(K)} \ .$$

Evidently we have $U_i(K) = \{1\}$ hence $R_i(L/K) = R(L)/R(K)$, but introducing this makes the definition similar to the one that we give for $\mathcal{N}_{L/K}$ (Definition 7.2.12).

7.2.2 Relative Units and Regulators for $\mathcal{N}_{L/K}$

The second notion of relative unit is of course linked with the norm.

Definition 7.2.10. (1) *We say that a pair $(\overline{\mathfrak{a}}, \alpha)$ is a* relative norm-unit *if $\overline{\mathfrak{a}} \in Cl_N(K)$ and if $\alpha \in U(L)$ is such that $\mathcal{N}_{L/K}(\alpha) \in \mu(K)$ is a root of unity of K. The group of relative norm-units is denoted $U_N(L/K)$.*
(2) *We define $U_{N,L}(K)$ (usually abbreviated as $U_N(K)$) to be the cokernel of the map $\mathcal{N}_{L/K}$ from $U(L)$ to $U(K)/\mu(K)$; in other words,*

$$U_N(K) = \frac{U(K)}{\mathcal{N}_{L/K}(U(L)) \cdot \mu(K)} \ .$$

We will identify the group $U_{N,0}(L/K)$ of units α whose relative norm is a root of unity of K with pairs $(\overline{\mathbb{Z}_K}, \alpha)$. By abuse of notation, if $\mathcal{N}_{L/K}(\alpha)$ is a root of unity, we will also call α a relative norm-unit, and $U_N(L/K) = Cl_N(K) \times U_{N,0}(L/K)$.

Proposition 7.2.11. (1) *We have the following exact sequences:*

$$1 \longrightarrow Cl_N(K) \longrightarrow U_N(L/K) \longrightarrow U(L) \overset{\mathcal{N}_{L/K}}{\longrightarrow} \frac{U(K)}{\mu(K)} \longrightarrow U_N(K) \longrightarrow 1 \tag{7}$$

and

$$1 \longrightarrow Cl_N(L/K) \longrightarrow Cl(L) \overset{\mathcal{N}_{L/K}}{\longrightarrow} Cl(K) \longrightarrow U_N(L/K) \longrightarrow U(L) \overset{\mathcal{N}_{L/K}}{\longrightarrow} \frac{U(K)}{\mu(K)} \longrightarrow U_N(K) \longrightarrow 1 \ .$$

(2) *The group $U_N(K)$ is finite and its exponent divides $n = [L : K]$.*
(3) *The rank of $U_N(L/K)$ is equal to $r(L) - r(K)$.*

Proof. Note that the map from $Cl_N(K)$ to $U_N(L/K)$ is the map sending $\overline{\mathfrak{a}}$ to $(\overline{\mathfrak{a}}, 1)$, and the map from $U_N(L/K)$ to $U(L)$ is the map sending $(\overline{\mathfrak{a}}, \alpha)$ to α, so the exact sequences of (1) immediately follow from the definitions. For (2), we note that as a quotient of $U(K)$, $U_N(K)$ is a finitely generated Abelian group, and if $u \in U(K)$ we have $\mathcal{N}_{L/K}(i_{L/K}(u)) = u^n$, hence $u^n \in \mathcal{N}_{L/K}(U(L))$, so the exponent of $U_N(K)$ divides n, and in particular since it is finitely generated, it is finite. (3) follows from the finiteness of $Cl_N(K)$, $\mu(K)$, $U_N(K)$, and the exact sequence (7). □

Remarks

(1) It would have been more natural to give a direct definition of $U_N(L/K)$ in a manner similar to what we have done for $U_i(L/K)$, and obtain this proposition as a consequence and not simply as the definition. I do not see how to do this.
(2) Note that the group $U_N(K)$ is not necessarily trivial; in other words, there may exist units of K that are not equal to norms of units of L even up to roots of unity (see Exercise 6).
(3) Similarly to the relative pseudo-units, this can also be considered as an exact sequence of K-groups, in a certain sense dual to the preceding one. The nontriviality of $U_N(K)$ means that it is reasonable to continue the sequence with something like

$$U(L) \overset{\mathcal{N}_{L/K}}{\longrightarrow} \frac{U(K)}{\mu(K)} \longrightarrow K_2(\mathbb{Z}_L/\mathbb{Z}_K) \longrightarrow K_2(\mathbb{Z}_L) \longrightarrow K_2(\mathbb{Z}_K)$$

and define the relative K-group $K_2(\mathbb{Z}_L/\mathbb{Z}_K)$ accordingly. We will not pursue this further; see [Ros] for details.

Definition 7.2.12. *The* relative regulator $R_N(L/K)$ *associated to the map $\mathcal{N}_{L/K}$ is defined by the formula*

$$R_N(L/K) = \frac{1}{|U_N(K)|} \frac{R(L)}{R(K)} \ .$$

Since the group $U_N(K)$ is in general nontrivial, $R_N(L/K)$ is not always equal to $R_i(L/K) = R(L)/R(K)$.

Example. I thank D. Simon for this example. Let $K = \mathbb{Q}(y)$, where y is a root of $y^2 - y - 26 = 0$, and let L be the extension of K generated by a root of $x^2 - x - 184 = 0$. Then one can show that

• $Cl(K) \simeq C_2$, $Cl(L) \simeq C_6 \times C_2$, $Cl_i(L/K) \simeq C_6 \times C_2$, $Cl_N(L/K) \simeq C_6$, $Cl_i(K) \simeq C_2$, $Cl_N(K) = \{1\}$.

• We have $\mu(K) = \mu(L) = \{\pm 1\}$, and if u_1 is a fundamental unit of K (in fact, $u_1 = 8y + 37$), one can prove that there exist units u_2 and u_3 of L such that (u_1, u_2, u_3) is a system of fundamental units of L satisfying $\mathcal{N}_{L/K}(u_2) = \mathcal{N}_{L/K}(u_3) = 1$. It follows in particular that $U(L)/i_{L/K}(U(K))$ is a free Abelian group generated by the classes of u_2 and u_3.

• One can also show that $U_N(K)$ is a group of order 2 generated by the class of u_1 (hence is nontrivial), that $U_i(L/K)$ is a free Abelian group of rank 2 generated by the classes of $\mathfrak{p}_2\alpha$ and u_3, where $\mathfrak{p}_2$ is one of the ideals of K above 2 (which is not principal) and α is such that $\mathfrak{p}_2\mathbb{Z}_L = (1/\alpha)\mathbb{Z}_L$ (so $\mathfrak{p}_2$ capitulates in L/K), and that $U_N(L/K) \simeq U(L)/U(K)$ is a free Abelian group of rank 2 generated by the classes of u_2 and of u_3, hence is not equal to $U_i(L/K)$.

7.3 Algorithms for Computing Relative Class and Unit Groups

Let L/K be an extension of number fields. Using the methods of [Coh0, Chapter 6], we can assume that we can solve all the usual problems in K, and in particular that we know its class and unit group and a principal ideal algorithm.

In this section, we give algorithms for computing the relative class groups $Cl_i(L/K)$, $Cl_i(K)$, $Cl_N(L/K)$, $Cl_N(K)$ and the relative unit groups $U_i(L/K)$, $U_N(L/K)$, and $U_N(K)$ (recall that $U_i(K) = \{1\}$). There are two ways to do this. One is to compute directly the absolute groups $Cl(L)$ and $U(L)$ using the methods of [Coh0, Chapter 6]. This may be expensive, since the absolute degree $[L : \mathbb{Q}] = [L : K][K : \mathbb{Q}]$ can be large. The other method is to use only relative techniques, and this is of course in the spirit of this book and in general much more efficient. We present both methods.

7.3.1 Using Absolute Algorithms

If the absolute degree of L is not too large, we can use the methods of [Coh0, Chapter 6] to compute $Cl(L)$ and $U(L)$ and to generate data so that we can have a principal ideal algorithm in L. Using the exact sequence techniques studied in Section 4.1, it is not difficult to compute the relative class and unit groups. Let us see in detail how to proceed. For the sake of completeness,

we will write the formal algorithms. These algorithms are quite technical, however, and otherwise quite easy to reconstruct, so the reader is advised to skip the algorithms (but not the comments) of this section at first.

• Thanks to the exact sequence (1), computing $Cl_i(L/K)$ and $Cl_i(K)$ amounts to computing the cokernel and the kernel of the map $i_{L/K}$ considered as a map from $Cl(K)$ to $Cl(L)$. This is done immediately by applying the algorithms of Section 4.1. Note that the map $i_{L/K}$ is computed explicitly by using the method explained in Section 2.5.3. This gives the following algorithm.

Algorithm 7.3.1 (Computation of $Cl_i(L/K)$ and of $Cl_i(K)$). Let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis, let $Cl(K) = (B, D_B)$ and $Cl(L) = (C, D_C)$ be the SNF of the absolute class groups $Cl(K)$ and $Cl(L)$, where $B = (\mathfrak{b}_i)$ and $C = (I_i)$. This algorithm computes the SNF of $Cl_i(L/K)$ and the HNF left divisor of D_B defining $Cl_i(K)$.

1. [Compute absolute bases] For each j and k, compute a $\mathbb{Z}$-basis $(\beta_{i,j,k})_i$ of the ideal $\mathfrak{a}_j\mathfrak{b}_k$ and set $\alpha_{i,j,k} \leftarrow \omega_j\beta_{i,j,k}$ (the $(\alpha_{i,j,k})_{i,j}$ now form a $\mathbb{Z}$-basis of $\mathfrak{b}_k\mathbb{Z}_L$).
2. [Use principal ideal algorithm in $Cl(L)$] Using the principal ideal algorithm in $Cl(L)$ ([Coh0, Algorithm 6.5.10]), compute a matrix $P = (p_{i,k})$ such that $\overline{\mathfrak{b}_k\mathbb{Z}_L} = \prod_i \overline{I_i}^{p_{i,k}}$ in $Cl(L)$.
3. [Compute $Cl_i(L/K)$] Apply Algorithm 4.1.3 to the system of generators and relations $\big((\overline{I_i}), (P|D_C)\big)$, output the corresponding SNF as the SNF of $Cl_i(L/K)$ and if desired the auxiliary matrix U_a so as to be able to solve the discrete logarithm problem in $Cl_i(L/K)$.
4. [Compute U_1] Apply an HNF algorithm to the matrix $(P|D_C)$, and let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(P|D_C)U = (0|H)$. We can discard the matrices U_2, U_3, U_4, and H (note that this computation can be done during the SNF computation in step 3).
5. [Compute $Cl_i(K)$] Let H_B be the HNF matrix of the matrix $(U_1|D_B)$. Output H_B as the HNF left divisor of D_B representing $Cl_i(K)$ and terminate the algorithm.

• Similarly, thanks to the exact sequence (2), computing $Cl_N(L/K)$ and $Cl_N(K)$ amounts to computing the kernel and cokernel of the map $\mathcal{N}_{L/K}$ considered as a map from $Cl(L)$ to $Cl(K)$, which is done once again by using the algorithms of Section 4.1. Note that the map $\mathcal{N}_{L/K}$ is computed explicitly by using Algorithm 2.5.2. The corresponding algorithm, which is almost identical to Algorithm 7.3.1, is as follows.

Algorithm 7.3.2 (Computation of $Cl_N(L/K)$ and of $Cl_N(K)$). Let $(\omega_i, \mathfrak{a}_i)$ be an integral pseudo-basis, let $Cl(K) = (B, D_B)$ and $Cl(L) = (C, D_C)$ be the SNF of the absolute class groups $Cl(K)$ and $Cl(L)$, where $B = (\mathfrak{b}_i)$ and $C =$

(I_i). This algorithm computes the HNF left divisor of D_B defining $Cl_N(L/K)$ and the SNF of $Cl_N(K)$.

1. [Compute relative norms] For each j, use Algorithm 2.5.2 to compute $\mathfrak{n}_j \leftarrow \mathcal{N}_{L/K}(I_j)$.
2. [Use principal ideal algorithm in $Cl(K)$] Using the principal ideal algorithm in $Cl(K)$ ([Coh0, Algorithm 6.5.10]), compute a matrix $P = (p_{i,k})$ such that $\overline{\mathfrak{n}_j} = \prod_i \overline{\mathfrak{b}_i}^{p_{i,j}}$ in $Cl(K)$.
3. [Compute $Cl_N(K)$] Apply Algorithm 4.1.3 to the system of generators and relations $\left(\left(\overline{\mathfrak{n}_j}\right), (P|D_B)\right)$, output the corresponding SNF as the SNF of $Cl_N(K)$ and if desired the auxiliary matrix U_a so as to be able to solve the discrete logarithm problem in $Cl_N(K)$.
4. [Compute U_1] Apply an HNF algorithm to the matrix $(P|D_B)$, and let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(P|D_B)U = (0|H)$. We can discard the matrices U_2, U_3, U_4, and H (note that this computation can be done during the SNF computation in step 3).
5. [Compute $Cl_N(L/K)$] Let H_C be the HNF matrix of the matrix $(U_1|D_C)$. Output H_C as the HNF left divisor of D_C representing $Cl_N(L/K)$ and terminate the algorithm.

• Computing $U(L)/i_{L/K}(U(K))$ is done using Algorithm 4.1.7 modified to take into account finitely generated but infinite Abelian groups. Note that to compute explicitly the map $i_{L/K}$ on elements, one simply needs to express a primitive element of K in terms of a primitive element of L, which is done using Algorithm 2.1.11, and discrete logarithms in $U(L)$ are computed using Algorithm 5.3.10. In the following algorithm on units, we will denote by r_K (resp., r_L) the rank of the group of units of K (resp., L) and by $r_{L/K}$ the rank of the relative unit groups, so that $r_{L/K} = r_L - r_K$. The formal algorithm for $U(L)/i_{L/K}(U(K))$ is as follows.

Algorithm 7.3.3 (Computation of $U(L)/i_{L/K}(U(K))$). Let $K = \mathbb{Q}(\alpha)$, let $L = K(\theta)$, and let $A \in K[X]$ such that $\alpha = A(\theta)$, found by using Algorithm 2.1.11. On the other hand, let ζ_K (resp., ζ_L) be a generator of $\boldsymbol{\mu}(K)$ (resp., $\boldsymbol{\mu}(L)$), let $(\varepsilon_i)_{1\le i\le r_K}$ (resp., $(\eta_i)_{1\le i\le r_L}$) be a system of fundamental units of K (resp., L), where ζ_K and the ε_i are given as polynomials in α, and ζ_L and the η_i as polynomials in θ. This algorithm computes the group $U(L)/i_{L/K}(U(K))$ in the form

$$U(L)/i_{L/K}(U(K)) = \bigoplus_{0\le i\le r_K} (\mathbb{Z}/d_i\mathbb{Z})\overline{\eta_i'} \oplus \bigoplus_{r_K<i\le r_L} \mathbb{Z}\overline{\eta_i'} \ .$$

1. [Compute matrix P] Using the polynomial A, compute the images in L by $i_{L/K}$ of ζ_K and the ε_i, and using Algorithm 5.3.10, compute the $(r_L+1)\times(r_K+1)$ matrix P such that

$$i_{L/K}([\zeta_K, \varepsilon_1, \ldots, \varepsilon_{r_K}]) = [\zeta_L, \eta_1, \ldots, \eta_{r_L}]P \ .$$

2. [Compute SNF] Using the algorithm described in the proof of Proposition 7.2.6, compute unimodular matrices U and V such that $VPU = \left(\frac{D}{0}\right)$.
3. [Terminate] Let $D = \mathrm{diag}(d_0, d_1, \ldots, d_{r_K})$, and set

$$[\eta'_0, \ldots, \eta'_{r_L+1}] \leftarrow [\zeta_L, \eta_1, \ldots, \eta_{r_L}]V^{-1}$$

in the usual multiplicative sense. Output the d_i and $\overline{\eta'_i}$ (classes modulo $i_{L/K}(U(K))$), and terminate the algorithm.

The determination of $U(L)/(\mu(L) \cdot i_{L/K}(U(K)))$ is done in a very similar manner.

Algorithm 7.3.4 (Computation of $U(L)/(\mu(L) \cdot i_{L/K}(U(K)))$). Let $K = \mathbb{Q}(\alpha)$, let $L = K(\theta)$, and let $A \in K[X]$ such that $\alpha = A(\theta)$, found by using Algorithm 2.1.11. On the other hand, let ζ_K (resp., ζ_L) be a generator of $\mu(K)$ (resp., $\mu(L)$), let $(\varepsilon_i)_{1 \le i \le r_K}$ (resp., $(\eta_i)_{1 \le i \le r_L}$) be a system of fundamental units of K (resp., L), where ζ_K and the ε_i are given as polynomials in α, and ζ_L and the η_i as polynomials in θ. This algorithm computes the group $U(L)/(\mu(L) \cdot i_{L/K}(U(K)))$ in the form

$$U(L)/(\mu(L) \cdot i_{L/K}(U(K))) = \bigoplus_{1 \le i \le r_K} (\mathbb{Z}/d_i\mathbb{Z})\overline{\eta'_i} \oplus \bigoplus_{r_K < i \le r_L} \mathbb{Z}\overline{\eta'_i} \ .$$

1. [Compute matrix P] Using the polynomial A, compute the images in L by $i_{L/K}$ of the ε_i, and using Algorithm 5.3.10, compute the $(r_L+1) \times r_K$ matrix P such that

$$i_{L/K}([\varepsilon_1, \ldots, \varepsilon_{r_K}]) = [\zeta_L, \eta_1, \ldots, \eta_{r_L}]P \ .$$

2. [Compute SNF] Let M be the matrix obtained from P by discarding the first row. Using the algorithm described in the proof of Proposition 7.2.6, compute unimodular matrices U and V such that $VMU = \left(\frac{D}{0}\right)$.
3. [Terminate] Let $D = \mathrm{diag}(d_1, \ldots, d_{r_K})$, and set

$$[\eta'_1, \ldots, \eta'_{r_L}] \leftarrow [\eta_1, \ldots, \eta_{r_L}]V^{-1}$$

in the usual multiplicative sense. Output the d_i and $\overline{\eta'_i}$ (classes modulo $\mu(L) \cdot i_{L/K}(U(K))$), and terminate the algorithm.

This algorithm can easily be extended so that it also computes the integers e_i occurring in Proposition 7.2.7 (Exercise 7).

• Since we now know how to compute $U(L)/i_{L/K}(U(K))$ and $Cl_i(K)$, we can use Algorithm 4.1.8 applied to the exact sequence (5) to compute $U_i(L/K)$. Formally, this gives the following.

Algorithm 7.3.5 (Computation of $U_i(L/K)$). Given the groups $Cl(K) = (B, D_B)$, $Cl(L) = ((I_i), D_C)$, $U(K)$, and $U(L)$, and an integral pseudo-basis $(\omega_i, \mathfrak{a}_i)$ of L, this algorithm computes the group $U_i(L/K)$ in the form

$$U_i(L/K) = \bigoplus_i (\mathbb{Z}/u_i\mathbb{Z})\overline{\gamma_i \mathfrak{c}_i} \oplus \bigoplus_{1 \le i \le r_{L/K}} \mathbb{Z}\,\overline{\delta_i \mathbb{Z}_K} \ .$$

1. [Compute $Cl_i(K)$] Using Algorithm 7.3.1, compute $Cl_i(K)$ as a left HNF divisor H_B of D_B, and applying Algorithm 4.1.3 to the system of generators and relations $(BH_B, H_B^{-1}D_B)$, compute the SNF of $Cl_i(K)$ as $Cl_i(K) = \bigoplus_i (\mathbb{Z}/c_i\mathbb{Z})\overline{c_i'}$. Let D_{Cl} be the diagonal matrix of the c_i.
2. [Compute $U(L)/i_{L/K}(U(K))$] Using Algorithm 7.3.3, compute

$$U(L)/i_{L/K}(U(K)) = \bigoplus_{0 \le i \le r_K} (\mathbb{Z}/d_i\mathbb{Z})\overline{\eta_i'} \oplus \bigoplus_{r_K < i \le r_L} \mathbb{Z}\overline{\eta_i'} \ ,$$

   and keep the unimodular matrix V obtained during the algorithm. Let D_U be the diagonal matrix of the d_i for $0 \le i \le r_K$.
3. [Compute $\mathfrak{c}_k'\mathbb{Z}_L$] For each j and k, compute a $\mathbb{Z}$-basis $(\beta_{i,j,k})_i$ of the ideal $\mathfrak{a}_j\mathfrak{c}_k'$ and set $\alpha_{i,j,k} \leftarrow \omega_j\beta_{i,j,k}$ (the $(\alpha_{i,j,k})_{i,j}$ now form a $\mathbb{Z}$-basis of $\mathfrak{c}_k'\mathbb{Z}_L$).
4. [Use principal ideal algorithm in $Cl(L)$] Using the principal ideal algorithm in $Cl(L)$ ([Coh0, Algorithm 6.5.10]), compute elements $\gamma_k' \in L$ such that $\mathfrak{c}_k'\mathbb{Z}_L = (1/\gamma_k')\mathbb{Z}_L$.
5. [Use principal ideal algorithm in $Cl(K)$] Using the same algorithm, compute elements $\beta_k \in K$ such that $\mathfrak{c}_k'^{\,c_k} = \beta_k\mathbb{Z}_K$.
6. [Compute discrete logarithms in $U(L)/i_{L/K}(U(K))$] (Here, we have $\gamma_k'^{\,c_k}\beta_k \in U(L)$ for all k). Using Algorithm 5.3.10, compute a matrix P such that the columns of P express the $\gamma_k'^{\,c_k}$ on ζ_L and the η_i.
7. [Terminate] Let M be the matrix formed by the first $r_K + 1$ rows of the matrix VP. Apply Algorithm 4.1.3 to the system of generators $(\eta_0', \ldots, \eta_{r_K}', \gamma_1'\mathfrak{c}_1', \ldots)$ and relations $\left(\begin{smallmatrix} D_U & -M \\ 0 & D_{Cl} \end{smallmatrix}\right)$, let $(\gamma_i\mathfrak{c}_i)$ be the system of generators, and $\mathrm{diag}(u_i)$ the SNF thus obtained. For $1 \le i \le r_{L/K}$, set $\delta_i \leftarrow \eta_{i-r_K}'$. Output the u_i, the classes of $\gamma_i\mathfrak{c}_i$ and of the $\delta_i\mathbb{Z}_K$, and terminate the algorithm.

• Computing $U_N(L/K)$ is done by applying Algorithm 4.1.13 to the exact sequence (7), or more simply by applying Algorithm 4.1.11 to compute $U_{N,0}(L/K)$, and using that $U_N(L/K) = Cl_N(K) \times U_{N,0}(L/K)$. Formally, this gives the following.

Algorithm 7.3.6 (Computation of $U_N(L/K)$). Given the groups $Cl(K)$, $Cl(L)$, $U(K)$, and $U(L)$, this algorithm computes the groups $U_{N,0}(L/K)$ and $U_N(L/K)$ in the form

$$U_{N,0}(L/K) = (\mathbb{Z}/w(L)\mathbb{Z})\zeta \oplus \bigoplus_{1 \le i \le r_{L/K}} \mathbb{Z}\eta_i'$$

and

$$U_N(L/K) = \bigoplus_i (\mathbb{Z}/c_i'\mathbb{Z})\left(\overline{\mathfrak{b}_i}, \zeta^{b_i}\right) \oplus \bigoplus_{1 \le i \le r_{L/K}} \mathbb{Z}\left(\overline{\mathbb{Z}_K}, \eta_i'\right) \ .$$

1. [Compute matrix P] If $(\eta_i)_{1\le i\le r_L}$ is a system of fundamental units of $U(L)$ (not including a generator of $\mu(L)$), using Algorithm 5.3.10, for each i compute the discrete logarithms of $\mathcal{N}_{L/K}(\eta_i)$ on a system of fundamental units of $U(K)$ together with a generator of $\mu(K)$, and let P be the $r_K \times r_L$ matrix whose columns are these discrete logarithms where the component on the generator of $\mu(K)$ is omitted.
2. [Compute integral kernel] Apply an HNF algorithm to the matrix P, and let $U = (U_1|U_2)$ be a unimodular matrix and H an HNF matrix such that $P(U_1|U_2) = (0|H)$. We can discard the matrices U_2 and H.
3. [Compute $U_{N,0}(L/K)$] Let

$$[\eta'_1, \ldots, \eta'_{r_{L/K}}] \leftarrow [\eta_1, \ldots, \eta_{r_L}]U_1$$

(of course multiplicatively), let ζ be a generator of $\mu(L)$, and output

$$U_{N,0}(L/K) \leftarrow (\mathbb{Z}/w(L)\mathbb{Z})\zeta \oplus \bigoplus_{1\le i\le r_{L/K}} \mathbb{Z}\eta'_i \ .$$

4. [Compute $Cl_N(K)$] Using Algorithm 7.3.2, compute an SNF for the group $Cl_N(K) = Cl(K)/\mathcal{N}_{L/K}(Cl(L))$ as $Cl_N(K) = \bigoplus_{1\le j\le k}(\mathbb{Z}/c_j\mathbb{Z})\mathfrak{c}_j$.
5. [Terminate] Set $G \leftarrow \left((\overline{\mathfrak{c}_j}, 1)_j, (\overline{\mathbb{Z}_K}, \zeta)\right)$ and $M \leftarrow \mathrm{diag}(c_1, \ldots, c_k, w(L))$. Apply Algorithm 4.1.3 to the system of generators and relations (G, M), thus obtaining $\bigoplus_{1\le i\le k+1}(\mathbb{Z}/c'_i\mathbb{Z})(\overline{\mathfrak{b}_i}, \zeta^{b_i})$, output

$$U_N(L/K) = \bigoplus_i (\mathbb{Z}/c'_i\mathbb{Z})(\overline{\mathfrak{b}_i}, \zeta^{b_i}) \oplus \bigoplus_{1\le i\le r_{L/K}} \mathbb{Z}(\overline{\mathbb{Z}_K}, \eta'_i) \ ,$$

and terminate the algorithm.

Finally, to compute $U_N(K)$, we compute $\mathcal{N}_{L/K}(U(L))$ as a subgroup of $U(K)$ using Algorithm 4.1.10, then $\mu(K)\cdot\mathcal{N}_{L/K}(U(L))$ by using Algorithm 4.1.14, and finally $U_N(K) = U(K)/(\mu(K)\cdot\mathcal{N}_{L/K}(U(L)))$ by using Algorithm 4.1.7. We leave to the reader the write-up of the corresponding formal algorithm (Exercise 8).

The proofs of the validity of all the above algorithms are easy but technical and are left to the reader (see Exercises 9 and 10).

7.3.2 Relative Ideal Reduction

The main ingredient necessary for class and unit group algorithms is the possibility to *reduce* an ideal. Whatever this means, starting with an ideal I, we must be able to compute an ideal J that is in the same ideal class as I (hence of the form I/α in the absolute case or of the form $I/(\alpha\mathfrak{a})$ with $\mathfrak{a}$ an ideal of the base field K in the relative case) and that in some sense is "small", which for us means that it must have a reasonably good chance of

having nonzero valuation only at small prime ideals. We have already given such a definition (although not very pleasing) in the relative quadratic case (Definition 2.6.11).

Instead of giving formal definitions, which probably will be superseded by further research, I prefer to give a generic driver algorithm and an instance that achieves the above-mentioned result.

Algorithm 7.3.7 (Reduction of an Ideal (Driver Algorithm)). Given an ideal I of L by a pseudo-matrix $(H, \mathfrak{c}_j)$ on some relative integral basis $(\omega_i, \mathfrak{q}_i^{-1})$, this algorithm outputs a pseudo-element $\mathfrak{a}\alpha$ (as usual, with $\mathfrak{a}$ ideal of K and $\alpha \in L^*$) and the ideal $J = I/(\mathfrak{a}\alpha)$ such that J is a primitive ideal (necessarily equivalent to I in $Cl_i(L/K)$) that should be reasonably "small" for our purposes.

1. [Initialize] Set $\mathfrak{a} \leftarrow \mathbb{Z}_K$, $J \leftarrow I$, $\alpha \leftarrow 1$, $n \leftarrow \prod_{1 \le i \le n} \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{q}_i)$, and $m \leftarrow 0$.
2. [Reduce to primitive] If $J = ((h_{i,j}), \mathfrak{c}_j)$, set $\mathfrak{b} \leftarrow \oplus_{1 \le i \le j} h_{i,j}\mathfrak{c}_j\mathfrak{q}_i$, $J \leftarrow \mathfrak{b}^{-1}J$ (in other words for all j, set $\mathfrak{c}_j \leftarrow \mathfrak{b}^{-1}\mathfrak{c}_j$), and $\mathfrak{a} \leftarrow \mathfrak{a}\mathfrak{b}$.
3. [Compute norm of J] Compute $d \leftarrow n\big(\prod_j \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{c}_j)\big)$.
4. [Finished?] If $m > 0$ and $d \ge m$, output $\mathfrak{a}\alpha$ and J, and terminate the algorithm. Otherwise, set $m \leftarrow d$.
5. [Reduce] Using a subalgorithm such as the one given below, choose some element $\beta \in L^*$, set $\alpha \leftarrow \alpha\beta$ and $J \leftarrow J/\beta$ (see remark below), and go to step 2.

Proof. By Proposition 2.3.5, the ideal $\mathfrak{b}$ computed in step 2 is the content of the ideal J, hence $\mathfrak{b}^{-1}J$ is a primitive ideal. By Proposition 2.3.1, the number d computed in step 3 is equal to the absolute norm of the ideal J. The number m initially contains 0 to indicate that no ideal norms have yet been computed, but afterwards is a strictly decreasing sequence of positive integers (the successive norms of the ideals J), hence the algorithm terminates. Furthermore, we clearly have throughout the algorithm the equality $J = I/(\mathfrak{a}\alpha)$, proving the algorithm's validity. □

Remark. The computation of $J \leftarrow J/\beta$ in step 5 is performed as follows. Let $(\gamma_j, \mathfrak{c}_j)$ be the pseudo-basis of J corresponding to the pseudo-matrix $((h_{i,j}), \mathfrak{c}_j)$. Then $(\gamma_j/\beta, \mathfrak{c}_j)$ is a pseudo-basis of J/β which is in general not in relative HNF. To apply step 2 it is necessary to have an HNF pseudo-matrix on the relative pseudo-basis; hence we compute the matrix H' of the (γ_j/β) on the ω_i, and we apply a relative HNF algorithm (for example, Algorithm 1.6.2) to the pseudo-matrix $(H', \mathfrak{c}_j)$.

Evidently, the most important and difficult thing that remains to be done is to explain the choice of $\beta \in L^*$ in step 5. A naive method is as follows.

Subalgorithm 7.3.8 (Naive Relative Ideal Reduction). Given a primitive ideal J by a relative HNF $((h_{i,j}, \mathfrak{c}_j))$ on an integral pseudo-basis $(\omega_i, \mathfrak{q}_i^{-1})$, this algorithm computes a $\beta \in L^*$ suitable for step 5 of Algorithm 7.3.7.

1. [Find small HNF matrix for J] Using Algorithm 2.3.3, compute a small HNF pseudo-matrix $((h'_{i,j}), \mathfrak{c}'_j)$ for the ideal J.
2. [Reduce] Let $\beta \leftarrow \omega_2 + h'_{1,2}\omega_1$ be the element of L^* corresponding to the second column of the matrix $(h'_{i,j})$, output β, and terminate the subalgorithm.

Remarks

(1) This choice of β is quite arbitrary, but it is the simplest. We could also modify the subalgorithm by taking for β each of the successive columns of the matrix $(h'_{i,j})$ and seeing which gives an ideal of smallest norm once it is reduced to a primitive ideal in step 2 of the main algorithm. This would be much slower, and it is not clear that it would bring much improvement. This idea could, however, be usefully applied upon termination of the main algorithm to see if some further reduction can be achieved by using the other columns.

(2) This algorithm is an exact generalization of the algorithm indicated in the quadratic case after Definition 2.6.11. Indeed, the LLL-reduction step (step 1 of the subalgorithm) is the same, and it is easy to check that the stopping condition $d \geq m$ of step 4 of the main algorithm is the same as the condition $\mathcal{N}(\mathfrak{a}) \leq \mathcal{N}(\mathfrak{c})$ given in Definition 2.6.11.

The above subalgorithm is based on a relative HNF representation of the ideal J. It would clearly be preferable to use a relative LLL representation so that we could generalize [Coh0, Section 6.5.1]. Attempts in this direction have been made (see, for example, [Fie-Poh]) but are not sufficient in practice.

7.3.3 Using Relative Algorithms

The main inefficiency of the absolute methods described in Section 7.3.1 is the necessity to compute the absolute invariants $Cl(L)$ and $U(L)$ directly.

In the present section, we show how it is possible to compute the relative invariants directly. This is, of course, one of the main motivations for introducing these relative invariants in the first place.

We will proceed as follows. We will first compute the invariants $Cl_i(L/K)$, $Cl_i(K)$, $U_i(L/K)$ attached to the map $i_{L/K}$. Indeed, as we will see, this is the natural setting for relative algorithms, while this is not the case for the invariants attached to the map $\mathcal{N}_{L/K}$.

Once we know the invariants attached to $i_{L/K}$, we use Algorithm 4.1.9 on the exact sequence (1) to compute $Cl(L)$ and Algorithm 4.1.13 on the exact sequence (4) to compute $U(L)$. We can then compute the groups attached to $\mathcal{N}_{L/K}$ as explained in the preceding section.

We briefly sketch the complete algorithm, without writing it formally. Most details are very close to the absolute case, and we refer to [Coh0, Chapter 6] for this. The main work is to write an implementation of relative prime

ideal decomposition (Algorithm 2.4.13), of the algorithms to compute valuations (Algorithms 2.3.13 and 2.3.14), of a pseudo–two-element representation of an ideal (Algorithm 2.3.8), and of ideal multiplication and powering (as explained in Section 2.3.4). The rest of the implementation can be copied almost verbatim from the absolute case. A preliminary version of this algorithm for the relative quadratic case can be found in [Co-Di-Ol5].

Let L/K be a given extension.

A) As a first initialization step, compute everything that will be needed about the base field K, including its class and unit groups as well as data for using the principal ideal algorithm in K.

B) Compute basic data about the *relative* extension L/K, in particular a pseudo-integral basis and the data allowing to go back and forth from ideals of K to L.

C) After choosing a suitable constant A, compute the prime ideals $\mathfrak{P}$ of L of absolute norm less than A which are not above inert primes of K, represented with five elements as explained after Algorithm 2.4.13. This will be the factor base, and the five-element representation will be used to compute $\mathfrak{P}$-adic valuations of ideals of L using Algorithm 2.3.13. At the same time, store the corresponding trivial relations including the pseudo-elements.

D) Choose small values s, l_1, and l_2 (for example, $s = 3$, $l_1 = -8$, $l_2 = 8$). Extract from the factor base the s unramified prime ideals $\mathfrak{P}_j$ of smallest norm, and compute in relative HNF the ideals $\mathfrak{P}_j^m$ for $1 \le j \le s$ and $l_1 \le m < l_2$.

E) For $1 \le j \le s$, choose random exponents m_j such that $l_1 \le m_j < l_2$, and compute a reduced ideal F equivalent to $\prod_{1 \le j \le s} \mathfrak{P}_j^{m_j}$. Using $\mathfrak{P}$-adic valuations, try to factor F on our factor base. If it does factor, store the resulting relation in the format explained above: a column vector of integer exponents, together with a pseudo-element generating a principal ideal.

F) If one believes that one has enough relations, simultaneously compute the Hermite normal form of the relation matrix and the corresponding pseudo-elements. As in the absolute case, the pseudo-elements that correspond to zero columns will be relative units for $i_{L/K}$, in other words, pseudo-elements αa such that $\alpha a \mathbb{Z}_L = \mathbb{Z}_L$.

G) Compute a tentative relative class group (the Smith normal form of our relation matrix) and class number (its determinant). From this and knowledge of the class group of the base field K, one easily deduces a tentative absolute class number. Similarly, using the principal ideal algorithm in the base field, from the relative units that we have obtained we can obtain a set of absolute units of L and compute a tentative absolute regulator.

H) As in the absolute case, since we have assumed GRH, we check that a suitable partial Euler product coming from the absolute Dedekind zeta function of L is sufficiently close (up to a factor of 2) to the tentative product of the class number by the regulator. If it is not, compute more relations and

go back to step F (equivalently, we can compute a partial Euler product coming from the quotient $\zeta_L(s)/\zeta_K(s)$).

I) Otherwise, we have computed the relative class group and regulator under some reasonable hypotheses. By definition, the ideals occurring in the pseudo-elements that have been kept for computing the relative units will generate the capitulation subgroup $Cl_i(K)$, which we thus compute at the same time. As in the absolute case, we can also compute a fundamental system of units if desired.

A word about the correctness of the result. As in the absolute case, we need to assume GRH in two essential places: first in the numerical verification of the product $h(L)R(L)$, to ensure fast convergence of the Euler product; second, we also need our factor base to generate the class group, by taking as constant A the value $12\log(|d(L)|)^2$ and using a theorem of E. Bach. As in the absolute case, however, we choose a much lower constant A, and then we must "be honest", that is, we must check that all the prime ideals of norm between A and Bach's bound are generated in the relative class group by the prime ideals of norm less than A. This can easily be done by generating more random relations involving the specific prime that is considered.

Finally, note that, as in the absolute case, if we keep the full HNF of the reduction matrix and the corresponding pseudo-elements (and not only the class group and the relative units), it is easy to obtain a principal ideal algorithm in L (more precisely, a pseudo-principal ideal algorithm).

7.3.4 An Example

The following example was given to us by C. Fieker. It shows some of the limitations of the absolute method, hence the usefulness of relative methods. Let $L = \mathbb{Q}(\zeta_9, \sqrt{-4201})$, where ζ_9 is a primitive ninth root of unity. Compute its class group, regulator, units, and so forth. For completeness, note that $\mathbb{Q}(\zeta_9)$ has class number equal to 1 while $\mathbb{Q}(\sqrt{-4201})$ has class group isomorphic to C_{36}. The field L enters naturally if you want to apply Kummer theory to the construction of the Hilbert class field of $\mathbb{Q}(\sqrt{-4201})$ (which can in this specific case of an imaginary quadratic field be constructed very simply by using complex multiplication; see Section 6.3).

The field L is a totally complex number field of degree 12 over $\mathbb{Q}$, with root discriminant approximately 673.6, so neither the degree nor the discriminant is too large compared to what can be presently attacked. However, if you feed it to the best existing programs (such as Kant and Pari), even a week of CPU time on a good workstation does not seem to produce enough relations in the class group. This is due mainly to the fact that L has many subfields. Thus, when we search, for example, for elements of small norm, they tend to be in the smaller subfields, and so the relations they generate are highly dependent.

We choose $K = \mathbb{Q}(\zeta_9)$. Since K has class number 1, the relative and absolute class groups coincide, and the pseudo-elements are simply elements.

As parameters in the algorithm described above, we chose (almost arbitrarily) $A = 600$, $s = 5$, $l_1 = -8$, $l_2 = 8$. To be completely honest, at least modulo GRH, we must check that the ideals up to Bach's bound (here 73291) are generated by the small ones.

In approximately one hour of CPU time on a workstation, we find that the class group is isomorphic to $C_{377244} \times C_6$ (note that $377244 = 2^2 \cdot 3^3 \cdot 7 \cdot 499$) and the absolute regulator is approximately equal to 3338795.5921522.... We also explicitly find the generators of the class group and the fundamental units themselves, as well as the information necessary to use a principal ideal algorithm in L.

As already mentioned, this example shows once again that the basic theoretical notions useful in a relative computational context are the notions linked to the map $i_{L/K}$: in other words, the relative class group $Cl_i(L/K)$, the unit group $U_i(L/K)$, the capitulation subgroup $Cl_i(K)$, and the notion of pseudo-element.

The main weakness of the algorithm, which is completely independent of the rest, is that we have not been able to develop a reasonable theory of relative reduction of ideals. The example that we have just given shows, however, that even a very naive definition of reduction such as the one used here suffices to give highly nontrivial results.

Philosophical Remark. To conclude these sections on relative class and unit groups, I would like to make a remark concerning the definitions that have been introduced. Most of these definitions do not occur in the existing literature, except the definitions (with different notation) of the capitulating group $Cl_i(K)$ and of the relative regulator $R_N(L/K)$ with respect to the norm (see [Ber-Mar]). The algorithms sketched above show that the relative groups $Cl_i(L/K)$, $U_i(L/K)$, and $Cl_i(K)$ arise very naturally from the relative algorithms, hence they are aesthetically and mathematically pleasing. The definitions of $Cl_N(L/K)$ and $Cl_N(K)$ also seem quite natural, although they do not occur naturally in the relative algorithms. This is confirmed by Theorem 7.1.5, which shows that the class group notions relative to $i_{L/K}$ and $\mathcal{N}_{L/K}$ are closely related, more precisely are identical outside primes dividing $(n, h(K))$.

As stated above, the group $U_i(L/K)$ is natural, and this is confirmed by the exact sequence (4) or the longer six-term exact sequence (6). On the other hand, I must admit that the definitions of $U_N(L/K)$ and of $U_N(K)$ are not satisfactory: they are artificially defined in such a way that the exact sequence (7) exists, and the artificialness of this is confirmed by the fact that natural maps between $U_i(L/K)$ and $U_N(L/K)$ do not seem to exist. Perhaps this is in the nature of things, but it is also possible that there is a better definition of these groups.

7.4 Inverting Prime Ideals

In this section, we would like to consider what happens in the very common situation where we "invert" certain prime ideals, leading to the notions of S-integers, S-units, S-class groups, and so forth, and to give the corresponding algorithms.

7.4.1 Definitions and Results

Recall from Definition 1.2.6 that a *place* of K is an equivalence class of nontrivial field norms and can be represented either by a prime ideal $\mathfrak{p}$ of $\mathbb{Z}_K$ or by one of the $r_1 + r_2$ embeddings σ_i of K into $\mathbb{C}$ (since $\sigma_{r_1+r_2+i} = \overline{\sigma_{r_1+i}}$, we do not need to consider the σ_i for $i > r_1 + r_2$).

In the rest of this chapter, the letter S will stand for a *finite* set of places of K containing the Archimedean places. The elements of S will be identified with prime ideals $\mathfrak{p}$ and embeddings σ as above.

Definition 7.4.1. *Let S be a finite set of places of a number field K.*

(1) *We say that an element $x \in K$ is an* S-integer *if $v_\mathfrak{p}(x) \geq 0$ (or, equivalently, $|x|_\mathfrak{p} \leq 1$) for all $v_\mathfrak{p} \notin S$. The ring of S-integers in K is denoted $\mathbb{Z}_{K,S}$.*

(2) *We say that an element $x \in K$ is an* S-unit *if $v_\mathfrak{p}(x) = 0$ (or, equivalently, $|x|_\mathfrak{p} = 1$) for all $v_\mathfrak{p} \notin S$. The group of S-units of K is denoted $U_S(K)$.*

It is easily checked that $\mathbb{Z}_{K,S}$ is a ring such that $\mathbb{Z}_K \subset \mathbb{Z}_{K,S} \subset K$, and $U_S(K)$ is a group such that $U(K) \subset U_S(K) \subset K^*$.

Proposition 7.4.2. *Let K and S be defined as above.*

(1) *The maps $I \mapsto I \cap \mathbb{Z}_K$ and $\mathfrak{a} \mapsto \mathfrak{a}\mathbb{Z}_{K,S}$ are inverse bijections from the set of integral ideals of $\mathbb{Z}_{K,S}$ to the set of integral ideals of $\mathbb{Z}_K$ coprime to all the prime ideals belonging to S. These maps preserve strict inclusion and prime and maximal ideals.*

(2) *The ring $\mathbb{Z}_{K,S}$ is a Dedekind domain.*

Proof. (1). Let $\mathfrak{a} = I \cap \mathbb{Z}_K$. It is clear that $\mathfrak{a}$ is an Abelian group, stable under multiplication by $\mathbb{Z}_K$, hence is an ideal of $\mathbb{Z}_K$. Assume by contradiction that $\mathfrak{a}$ is not coprime to the prime ideals of S, and let $\mathfrak{p} \in S$ be a prime ideal dividing $\mathfrak{a}$, so that $\mathfrak{a} = \mathfrak{p}\mathfrak{b}$ with $\mathfrak{b} \subset \mathbb{Z}_K$. Since $\mathfrak{p} \in S$, we have $\mathfrak{p}\mathbb{Z}_{K,S} \subset \mathbb{Z}_{K,S}$ and $\mathfrak{p}^{-1}\mathbb{Z}_{K,S} \subset \mathbb{Z}_{K,S}$, so that $\mathbb{Z}_{K,S} = \mathfrak{p}\mathbb{Z}_{K,S}$. On the other hand, we have $\mathfrak{p}\mathfrak{b} \subset I$, hence $\mathfrak{p}\mathfrak{b}\mathbb{Z}_{K,S} \subset I$, so $\mathfrak{b}\mathbb{Z}_{K,S} \subset I$, and hence $\mathfrak{b} \subset I$. Since $\mathfrak{b} \subset \mathbb{Z}_K$, we deduce that $\mathfrak{b} \subset \mathfrak{a} = \mathfrak{p}\mathfrak{b}$, which is absurd since $\mathfrak{p}$ is an invertible ideal not equal to $\mathbb{Z}_K$. Therefore, $\mathfrak{a}$ is an ideal of $\mathbb{Z}_K$ coprime to all the prime ideals belonging to S.

We clearly have $\mathfrak{a}\mathbb{Z}_{K,S} \subset I\mathbb{Z}_{K,S} \subset I$. To show the reverse inclusion, let $x \in I$. We have $0 \in \mathfrak{a}\mathbb{Z}_{K,S}$, so we assume $x \neq 0$. By definition of $\mathbb{Z}_{K,S}$, we

have $x\prod_{\mathfrak{p}\in S}\mathfrak{p}^{-v_{\mathfrak{p}}(x)}\subset\mathbb{Z}_K$. Since $x\in I$ and I is an ideal of $\mathbb{Z}_{K,S}$, it follows that $x\prod_{\mathfrak{p}\in S}\mathfrak{p}^{-v_{\mathfrak{p}}(x)}\subset I$, so that $x\prod_{\mathfrak{p}\in S}\mathfrak{p}^{-v_{\mathfrak{p}}(x)}\subset\mathfrak{a}$, or in other words

$$x\in\Big(\prod_{\mathfrak{p}\in S}\mathfrak{p}^{v_{\mathfrak{p}}(x)}\Big)\mathfrak{a}\ .$$

By definition of $\mathbb{Z}_{K,S}$, if $\mathfrak{p}\in S$ we have $\mathfrak{p}^k\subset\mathbb{Z}_{K,S}$ for any $k\in\mathbb{Z}$. It follows that $x\in\mathfrak{a}\mathbb{Z}_{K,S}$, so $I\subset\mathfrak{a}\mathbb{Z}_{K,S}$, and so $(I\cap\mathbb{Z}_K)\mathbb{Z}_{K,S}=I$.

Conversely, let $\mathfrak{a}$ be an ideal of $\mathbb{Z}_K$ coprime to all the prime ideals belonging to S. Set $\mathfrak{b}=(\mathfrak{a}\mathbb{Z}_{K,S})\cap\mathbb{Z}_K$. Applying what we have just proved to $I=\mathfrak{a}\mathbb{Z}_{K,S}$, we deduce that $\mathfrak{b}\mathbb{Z}_{K,S}=\mathfrak{a}\mathbb{Z}_{K,S}$ and that $\mathfrak{b}$ is an ideal of $\mathbb{Z}_K$ coprime to all the prime elements of S. Assume by contradiction that $\mathfrak{a}\neq\mathfrak{b}$. Then, there exists a prime ideal $\mathfrak{p}$, necessarily not in S, such that $v_{\mathfrak{p}}(\mathfrak{a})\neq v_{\mathfrak{p}}(\mathfrak{b})$. Assume, for example, that $v_{\mathfrak{p}}(\mathfrak{a})>v_{\mathfrak{p}}(\mathfrak{b})=v$. Write $\mathfrak{b}=\mathfrak{p}^v\mathfrak{b}'$, with $\mathfrak{p}\nmid\mathfrak{b}'$, and $\mathfrak{a}=\mathfrak{p}^{v+1}\mathfrak{a}'$ for some integral ideal $\mathfrak{a}'$. Hence we obtain $\mathfrak{b}'\mathbb{Z}_{K,S}=\mathfrak{p}\mathfrak{a}'\mathbb{Z}_{K,S}\subset\mathfrak{p}\mathbb{Z}_{K,S}$, and intersecting with $\mathbb{Z}_K$, we obtain

$$\mathfrak{b}'\subset(\mathfrak{b}'\mathbb{Z}_{K,S})\cap\mathbb{Z}_K\subset(\mathfrak{p}\mathbb{Z}_{K,S})\cap\mathbb{Z}_K\ . \tag{8}$$

Note that $(\mathfrak{p}\mathbb{Z}_{K,S})\cap\mathbb{Z}_K=\mathbb{Z}_K$ implies $\mathfrak{p}^{-1}\subset\mathbb{Z}_{K,S}$, hence $\mathfrak{p}\in S$, contrary to our hypothesis. Since $\mathfrak{p}$ is a maximal ideal and $(\mathfrak{p}\mathbb{Z}_{K,S})\cap\mathbb{Z}_K$ contains $\mathfrak{p}$, we must therefore have $(\mathfrak{p}\mathbb{Z}_{K,S})=\mathfrak{p}$. Equation (8) thus gives $\mathfrak{b}'\subset\mathfrak{p}$, in other words $\mathfrak{p}\mid\mathfrak{b}'$, which is a contradiction. Therefore, $\mathfrak{b}=\mathfrak{a}$, and our maps are indeed inverse maps, as claimed. Since they are bijective, they preserve strict inclusion and maximal ideals. Finally, this also implies that if $\mathfrak{P}$ is a nonzero prime ideal of $\mathbb{Z}_{K,S}$, then $\mathfrak{P}\cap\mathbb{Z}_K$ is a nonzero ideal of $\mathbb{Z}_K$ different from $\mathbb{Z}_K$, which clearly satisfies the definition of a prime ideal. Conversely, if $\mathfrak{p}$ is a nonzero prime ideal of $\mathbb{Z}_K$ not in S, then $\mathfrak{p}$ is a maximal ideal, hence $\mathfrak{P}=\mathfrak{p}\mathbb{Z}_{K,S}$ is also maximal, finishing the proof of (1).

For (2), we note that $\mathbb{Z}_{K,S}\subset K$ is an integral domain. By (1), any strictly increasing sequence of ideals of $\mathbb{Z}_{K,S}$ gives rise to a strictly increasing sequence of ideals of $\mathbb{Z}_K$, hence is finite, so $\mathbb{Z}_{K,S}$ is a Noetherian ring. If $\mathfrak{P}$ is a nonzero prime ideal of $\mathbb{Z}_{K,S}$, $\mathfrak{P}\cap\mathbb{Z}_K$ is a nonzero prime ideal of $\mathbb{Z}_K$, hence is maximal. Thus, by (1), $\mathfrak{P}$ is a maximal ideal of $\mathbb{Z}_{K,S}$, so every nonzero prime ideal of $\mathbb{Z}_{K,S}$ is maximal.

Finally, let $x\in K$ be a root of a monic equation

$$x^n+a_{n-1}x^{n-1}+\cdots+a_0=0$$

with $a_i\in\mathbb{Z}_{K,S}$. We may assume $x\neq0$. Let $\mathfrak{p}\notin S$. If $v=v_{\mathfrak{p}}(x)<0$, then $v_{\mathfrak{p}}(a_{n-1}x^{n-1}+\cdots+a_0)\geq(n-1)v$ while $v_{\mathfrak{p}}(x^n)=vn$, which is absurd since $v<0$. Hence for every $\mathfrak{p}\notin S$, $v_{\mathfrak{p}}(x)\geq0$; in other words, $x\in\mathbb{Z}_{K,S}$, so $\mathbb{Z}_{K,S}$ is integrally closed and hence is a Dedekind domain, as claimed. □

Since $\mathbb{Z}_{K,S}$ is a Dedekind domain, in addition to its ideals and prime ideals completely described in terms of those of $\mathbb{Z}_K$, we can define its class group $Cl_S(K)$ in the usual way.

Definition 7.4.3. *Let S be a finite set of prime ideals of K.*

(1) *If $\mathfrak{a}$ is an ideal of K, we say that $\mathfrak{a}$ is* S-integral *if $v_{\mathfrak{p}}(\mathfrak{a}) \geq 0$ for all prime ideals $\mathfrak{p} \notin S$.*
(2) *We define $Cl_S(K)$ to be the class group of the Dedekind domain $\mathbb{Z}_{K,S}$; in other words, by Proposition 7.4.2, the quotient of the group of S-integral ideals by the subgroup of S-integral principal ideals of K.*

Proposition 7.4.4. *We have a canonical isomorphism*

$$Cl_S(K) \simeq Cl(K)/ < \overline{\mathfrak{p}_i} >_{\mathfrak{p}_i \in S} ,$$

where $< \overline{\mathfrak{p}_i} >$ denotes the subgroup of $Cl(K)$ generated by ideal classes of the prime ideals in S.

Proof. Let $\overline{I}$ be an ideal class in $Cl_S(K)$, and define $f(\overline{I}) = \overline{I \cap \mathbb{Z}_K}$ in $Cl(K)/ < \overline{\mathfrak{p}_i} >$. The map f is well-defined and is a group homomorphism. Assume that $f(\overline{I}) = 1$. This means that $I \cap \mathbb{Z}_K = \alpha \prod_i \mathfrak{p}_i^{x_i}$ for some $\alpha \in K^*$. Multiplying by $\mathbb{Z}_{K,S}$ and using Proposition 7.4.2 and $\mathfrak{p}\mathbb{Z}_{K,S} = \mathbb{Z}_{K,S}$ for $\mathfrak{p} \in S$, we obtain $I = \alpha\mathbb{Z}_{K,S}$, so $\overline{I}$ is trivial, hence f is injective. Finally, if $\overline{\mathfrak{a}}$ is some ideal class in $Cl(K)/ < \overline{\mathfrak{p}_i} >$, by Corollary 1.2.11 we can choose as representative an integral ideal $\mathfrak{a}$ coprime to the product of all prime ideals of S, and then by Proposition 7.4.2, we have $f(\overline{\mathfrak{a}\mathbb{Z}_{K,S}}) = \overline{\mathfrak{a}}$, so f is surjective. □

Corollary 7.4.5. *There exists S_1 such that for any $S \supset S_1$, the ring $\mathbb{Z}_{K,S}$ is a principal ideal domain.*

Proof. Let $(\overline{\mathfrak{a}_i})$ be generators of $Cl(K)$, and let S_1 be a set containing all Archimedean places and all prime ideals $\mathfrak{p}$ such that $v_{\mathfrak{p}}(\mathfrak{a}_i) \neq 0$ for some i. This set is finite, and the proposition implies that if $S \supset S_1$ then $Cl_S(K)$ is trivial, so that $\mathbb{Z}_{K,S}$ is a principal ideal domain. □

7.4.2 Algorithms for the S-Class Group and S-Unit Group

Using Section 4.1.3, Proposition 7.4.4 allows us to give an algorithm to compute $Cl_S(K)$.

Algorithm 7.4.6 (S-Class Group). Let $Cl(K) = (B, D_B)$ be the SNF of the class group of K, where $B = (\overline{\mathfrak{b}_i})$ and the $\mathfrak{b}_i$ are ideals of K. This algorithm computes the SNF (C, D_C) of $Cl_S(K)$, where $C = (\overline{\mathfrak{c}_i})$ and the $\mathfrak{c}_i$ are ideals of $\mathbb{Z}_{K,S}$.

1. [Compute discrete logarithms] Using the principal ideal algorithm ([Coh0, Algorithm 6.5.10]), compute the matrix P whose columns are the discrete logarithms of $\overline{\mathfrak{p}}$ with respect to B, for each $\mathfrak{p} \in S$.

2. [Terminate] Apply Algorithm 4.1.3 to the system of generators and relations $(B,(P|D_B))$, thus obtaining an SNF (C',D_C). Let $C'=(\overline{\mathfrak{c}_i'})$ for ideals $\mathfrak{c}_i'$ of K. Using Algorithm 1.3.14, compute ideals $\mathfrak{c}_i''$ in the same ideal class as $\mathfrak{c}_i'$ which are coprime to all the prime ideals belonging to S. Set $\mathfrak{c}_i \leftarrow \mathfrak{c}_i''\mathbb{Z}_{K,S}$ and $C \leftarrow (\overline{\mathfrak{c}_i})$, output (C,D_C), and terminate the algorithm.

We leave to the reader the proof that this is the algorithmic translation of Proposition 7.4.4.

Similarly, we can deal with the group $U_S(K)$ of S-units. The result is as follows. We let s be the number of prime ideals in S.

Proposition 7.4.7. *Let $Cl(K)=(B,D_B)$ with $B=(\overline{\mathfrak{b}_i})$ be the SNF of $Cl(K)$, where $\mathfrak{b}_i$ are ideals and $D_B=\mathrm{diag}(b_i)$. Let*

$$U(K)=(\mathbb{Z}/w(K)\mathbb{Z})\varepsilon_0\oplus\bigoplus_{1\le i\le r}\mathbb{Z}\varepsilon_i$$

be the unit group of K in SNF. Let $\beta_i\in K$ be such that $\mathfrak{b}_i^{b_i}=(1/\beta_i)\mathbb{Z}_K$. For each prime ideal $\mathfrak{p}_j\in S$, write

$$\mathfrak{p}_j=\alpha_j\prod_i\mathfrak{b}_i^{p_{i,j}}$$

with $\alpha_j\in K^*$ and $p_{i,j}\in\mathbb{Z}$, and let $P=(p_{i,j})$. Finally, let $U=\left(\begin{smallmatrix}U_1 & U_2\\ U_3 & U_4\end{smallmatrix}\right)$ be the unimodular matrix such that $(P|D_B)U=(0|H)$, where H is the HNF of $(P|D_B)$. If

$$[\gamma_1,\dots,\gamma_s]=[\alpha_1,\dots,\alpha_s,\beta_1,\dots,\beta_m]\begin{pmatrix}U_1\\ U_3\end{pmatrix}$$

in the usual multiplicative sense used in Chapter 4, then

$$U_S(K)=(\mathbb{Z}/w(K)\mathbb{Z})\varepsilon_0\oplus\bigoplus_{1\le i\le r}\mathbb{Z}\varepsilon_i\oplus\bigoplus_{1\le i\le s}\mathbb{Z}\gamma_i\ .$$

In particular, the torsion subgroup of $U_S(K)$ is equal to that of $U(K)$, and the rank of $U_S(K)$ is equal to $r+s=r_1+r_2-1+|S_0|$, where S_0 denotes the set of prime ideals belonging to S (this can also be written $|S|-1$, since there are r_1+r_2 Archimedean places in S).

Proof. Let S_0 be the row vector of the prime ideals belonging to S and let $B'=[\mathfrak{b}_1,\dots,\mathfrak{b}_m]$ be the row vector of the ideals $\mathfrak{b}_i$. We use again matrix notation as we did in Chapter 4. I first claim that

$$S_0U_1=[\gamma_1\mathbb{Z}_K,\dots,\gamma_m\mathbb{Z}_K]\ .$$

Indeed, by definition, we have

$$B'(P|D_B)=((\alpha_j^{-1}\mathfrak{p}_j)|(\beta_j^{-1}\mathbb{Z}_K))\ .$$

Hence, multiplying on the right by $\binom{U_1}{U_3}$, we obtain

$$[\mathbf{Z}_K,\dots,\mathbf{Z}_K] = ((\alpha_j^{-1})|(\beta_j^{-1}))\binom{U_1}{U_3}\cdot S_0U_1 = (\gamma_j^{-1})S_0U_1 \ ,$$

proving my claim.

Now,

$$\alpha \in U_S(K) \iff \forall \mathfrak{p} \notin S,\ v_{\mathfrak{p}}(\alpha) = 0 \iff \exists X \in \mathbf{Z}^{S_0},\ \alpha\mathbf{Z}_K = S_0X \ .$$

Taking ideal classes, we see that $\overline{S_0}X = \mathbf{1}$ in $Cl(K)$. But, by definition,

$$\begin{aligned}\overline{S_0}X = \mathbf{1} &\iff BPX = \mathbf{1} \iff \exists Y \in \mathbf{Z}^m,\ PX = D_BY \\ &\iff \binom{X}{-Y} \in \mathrm{Ker}(P|D_B) \ ,\end{aligned}$$

where of course Ker denotes the integer kernel. If $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ is the unimodular matrix such that $(P|D_B)U = (0|H)$, where H is the HNF of $(P|D_B)$, we have seen several times (see, for example, Section 4.1.6) that the integer kernel is generated by the columns of $\binom{U_1}{U_3}$, hence $\overline{S_0}X = \mathbf{1} \iff X = U_1Z$ for some vector $Z \in \mathbf{Z}^{S_0}$.

Coming back to α, using $S_0U_1 = (\gamma_i\mathbf{Z}_K)$, we see that

$$\alpha \in U_S(K) \iff \exists Z \in \mathbf{Z}^{S_0},\ \alpha\mathbf{Z}_K = S_0U_1Z \iff \alpha\mathbf{Z}_K = \prod_i \gamma_i^{z_i}$$

for some integers z_i. Thus $\alpha/\prod_i \gamma_i^{z_i}$ is a unit; hence $U_S(K)$ is generated by the γ_i and the ε_i.

The relations between these generators are as follows. Assume that $\prod_i \gamma_i^{z_i}$ is a unit, with $Z = (z_i)$, and set $X = U_1Z$. Then

$$S_0X = S_0U_1Z = \prod_i \gamma_i^{z_i}\mathbf{Z}_K = \mathbf{Z}_K \ .$$

This means that $\prod_{\mathfrak{p}_i \in S} \mathfrak{p}_i^{x_i} = \mathbf{Z}_K$, and since the $\mathfrak{p}_i$ are distinct prime ideals, this implies that $x_i = 0$ for all i, in other words, that $X = 0$. However, by Lemma 4.1.12, U_1 has a nonzero determinant, hence $Z = 0$, and so there are no extra relations among the γ_i. □

Although this completely and explicitly answers the problem of computing $U_S(K)$, it is still not completely satisfactory from a practical point of view since the new generators γ_i can have very large coefficients and not even be algebraic integers. As usual, the problem of reducing the size of generators is not completely straightforward and can be done using variants of the LLL algorithm. However, it is easy to see that one can choose the γ_i to be algebraic integers as follows.

Let V be a unimodular matrix such that U_1V is in Hermite normal form. In particular, the entries of U_1V are nonnegative. Define

$$[\gamma_1', \ldots, \gamma_s'] = [\gamma_1, \ldots, \gamma_s]V \ .$$

Since V is unimodular, the γ_i' generate (multiplicatively) the same lattice as the γ_i, so we can use them instead as generators of $U_S(K)$. Furthermore, they are algebraic integers. In fact, since U_1V has nonnegative entries, $S_0U_1V = [\mathfrak{a}_1, \ldots, \mathfrak{a}_s]$ is a vector of integral ideals, but

$$S_0U_1V = [\gamma_1\mathbf{Z}_K, \ldots, \gamma_s\mathbf{Z}_K]V = [\gamma_1'\mathbf{Z}_K, \ldots, \gamma_s'\mathbf{Z}_K] \ ,$$

so $\gamma_i'\mathbf{Z}_K = \mathfrak{a}_i$; hence γ_i' is indeed an algebraic integer, as claimed.

We are thus led to the following algorithm for computing $U_S(K)$.

Algorithm 7.4.8 (S-Unit Group). Let $Cl(K) = (B, D_B)$ be the SNF of the class group of K, where $B = (\overline{\mathfrak{b}_i})$ and the $\mathfrak{b}_i$ are ideals of K. This algorithm computes algebraic integers γ_i for $1 \le i \le s$ such that

$$U_S(K) = U(K) \oplus \bigoplus_{1 \le i \le s} \mathbf{Z}\gamma_i \ .$$

We let $\mathfrak{p}_j$ be the prime ideals of S.

1. [Compute discrete logarithms] Using the principal ideal algorithm ([Coh0, Algorithm 6.5.10]), compute the matrix P whose columns are the discrete logarithms of $\overline{\mathfrak{p}}$ with respect to B, for each $\mathfrak{p} \in S$ (this step is, of course, identical to step 1 of Algorithm 7.4.6).
2. [Compute big HNF] Using one of the algorithms for HNF computations, compute the unimodular matrix $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ such that $(P|D_B)U = (0|H)$ with H in HNF.
3. [Compute $\gamma_i\mathbf{Z}_K$] Compute the HNF W of the matrix U_1, and set $[\mathfrak{a}_1, \ldots, \mathfrak{a}_s] \leftarrow [\mathfrak{p}_1, \ldots, \mathfrak{p}_s]W$.
4. [Find generators] (Here the $\mathfrak{a}_j$ are principal ideals.) Using the principal ideal algorithm again, for each j, find γ_j such that $\mathfrak{a}_j = \gamma_j\mathbf{Z}_K$. Output the γ_j and terminate the algorithm.

Remark. Note that, although not really necessary, in this algorithm we use the principal ideal algorithm twice. Indeed, in step 1 we could keep the extra information given by the principal ideal algorithm, in other words, the α_j of Proposition 7.4.7. However, since the $\mathfrak{p}_j$ are not principal ideals in general, these α_j are usually large, and the subsequent operations may make them even larger. Hence, it is usually advisable to find the generators at the very end, when the computations are finished and we know that we have principal ideals, as we have done in step 4 of the algorithm.

We can reduce even more the size of the γ_i by replacing γ_i by γ_i/ε for a suitable unit ε, which still gives generators of $U_S(K)$. To do this, we multiply γ_i recursively by very small powers of a generating set of the unit group as long as the size of γ_i (measured in any reasonable way) decreases.

7.5 Solving Norm Equations

7.5.1 Introduction

In this section, we explain how the notions of S-units and S-integers can be used to solve absolute or relative norm equations. We closely follow a paper of D. Simon [Sim1], whom we heartily thank for the present section.

Let L/K be an extension of number fields, and let $a \in K^*$. We would like to know if there exists $x \in L$ such that $a = \mathcal{N}_{L/K}(x)$; additionally, we want to give an algorithm for finding such an x if it exists. In addition, if $a \in \mathbf{Z}_K$, we may want to additionally require that $x \in \mathbf{Z}_L$. In this section we give solutions to all these problems.

We begin by noticing that some nontrivial phenomena may occur. Consider the following example. Let $K = \mathbb{Q}$ and $L = \mathbb{Q}(\sqrt{34})$. Since the fundamental unit of L is $35 + 6\sqrt{34}$, which is of norm $+1$, the norm of a unit of L is always equal to $+1$, so the equation $\mathcal{N}_{L/K}(x) = -1$ is not soluble with $x \in U(L)$, hence also with $x \in \mathbb{Z}_L$. On the other hand, we check that, for example,

$$\mathcal{N}_{L/K}\left(\frac{5+\sqrt{34}}{3}\right) = \mathcal{N}_{L/K}\left(\frac{3+\sqrt{34}}{5}\right) = \mathcal{N}_{L/K}\left(\frac{27+5\sqrt{34}}{11}\right) = -1 \ ,$$

hence the equation $\mathcal{N}_{L/K}(x) = -1$ does have solutions (in fact, an infinite number of nonassociate solutions) if we do not restrict to $x \in U(L)$ or to $x \in \mathbb{Z}_L$. The primes occurring in the denominator of the solutions happen to be split primes in L such that the prime ideals above them are not principal, in other words generate the class group, which is here of order 2. As we will see, this is indeed the general behavior, at least when L/K is a Galois extension.

We make the following abuse of notation. In the rest of this chapter, since Archimedean places are not used, if S is a finite set of prime ideals of a number field we will write $\mathbb{Z}_{K,S}$, $U_S(K), \ldots$ instead of $\mathbb{Z}_{K,S'}$, $U_{S'}(K), \ldots$, where S' is the union of S with all Archimedean places. In addition, we make the following very useful convention. If S is a finite set of prime ideals of the base field K, and if T is the set of prime ideals of L above the prime ideals of S, we will still write $\mathbb{Z}_{L,S}$, $U_S(L)$, and $Cl_S(L)$ instead of $\mathbf{Z}_{L,T}$, $U_T(L)$, and $Cl_T(L)$ to avoid explicitly introducing the set T.

Recall that the *exponent* of an Abelian group $\mathcal{A}$ is the least positive integer m such that g^m is the unit element of $\mathcal{A}$ for all $g \in \mathcal{A}$. Note the following trivial result.

Proposition 7.5.1. *The exponent of the quotient group*

$$(\mathcal{N}_{L/K}(L^*) \cap U_S(K))/\mathcal{N}_{L/K}(U_S(L))$$

divides $[L : K]$.

Proof. Indeed, if $a \in U_S(K) \subset K^*$, then

$$a^{[L:K]} = \mathcal{N}_{L/K}(a) \in \mathcal{N}_{L/K}(U_S(L)) \ .$$

□

Definition 7.5.2. *We will say that a finite set S_0 of prime ideals is* suitable *for the extension L/K if for all finite $S \supset S_0$ we have*

$$\mathcal{N}_{L/K}(L^*) \cap U_S(K) = \mathcal{N}_{L/K}(U_S(L)) \ ;$$

in other words, if every S-unit of K that is the norm of an element of L is in fact the norm of an S-unit of L.

The following proposition is immediate.

Proposition 7.5.3. *Assume that S_0 is suitable for the extension L/K, let $a \in K^*$, and call S_a the set of prime ideals $\mathfrak{p}$ of K such that $v_{\mathfrak{p}}(a) \neq 0$. Then the equation $\mathcal{N}_{L/K}(x) = a$ is soluble with $x \in L$ if and only if it is soluble with $x \in U_{L,S_0 \cup S_a}$.*

Proof. Indeed, if $S = S_0 \cup S_a$, then $a \in U_S(K)$ by definition, and since S_0 is suitable and $S \supset S_0$, if the equation $\mathcal{N}_{L/K}(x) = a$ is soluble, we have $a \in \mathcal{N}_{L/K}(U_S(L))$. □

Thus, once a suitable S_0 has been found, this proposition allows us to solve norm equations by looking only in the group $U_{L,S}$ for a certain S, which is much easier to control.

7.5.2 The Galois Case

We start with the Galois case, which is much simpler than the general case.

Theorem 7.5.4. *Let L/K be a Galois extension, and let S_0 be a set of prime ideals of K such that $Cl_i(L/K)$ can be generated by the classes of ideals divisible only by prime ideals of L above the ideals of S_0. Then S_0 is suitable for the extension L/K in the sense of Definition 7.5.2: in other words, for all $S \supset S_0$, we have $\mathcal{N}_{L/K}(U_S(L)) = \mathcal{N}_{L/K}(L^*) \cap U_S(K)$. In addition, we also have $\mathcal{N}_{L/K}(\mathbb{Z}_{L,S}) = \mathcal{N}_{L/K}(L^*) \cap \mathbb{Z}_{K,S}$.*

Thanks to Proposition 7.5.3, this theorem will allow us to solve norm equations in the Galois case.

We first need a lemma.

Lemma 7.5.5. *Let L/K be a Galois extension of number fields, let S be a finite set of prime ideals of K, and let I and J be S-integral ideals of L such that $\mathcal{N}_{L/K}(I) = \mathfrak{a}\mathcal{N}_{L/K}(J)$ for some S-integral ideal $\mathfrak{a}$ of K. Assume that for $1 \leq i \leq k$ there exist prime ideals $\mathfrak{P}_i$ of L such that $\prod_{1 \leq i \leq k} \mathfrak{P}_i \mid J$ as ideals of $\mathbb{Z}_{L,S}$. Then for each $i \leq k$ there exists $\sigma_i \in \mathrm{Gal}(L/K)$ such that $\prod_{1 \leq i \leq k} \sigma_i(\mathfrak{P}_i) \mid I$ as ideals of $\mathbb{Z}_{L,S}$.*

Proof. We prove the lemma by induction on k, the case $k = 0$ being trivial. Assume first $k = 1$, and let $\mathfrak{P}$ be a prime ideal dividing J. Since $\mathcal{N}_{L/K}(\mathfrak{P})$ divides $\mathcal{N}_{L/K}(J)$, it also divides $\mathcal{N}_{L/K}(I)$ since $\mathfrak{a}$ is S-integral. Since the extension is Galois, by the remark made at the end of Section 2.2.5, we have for any ideal I, $\mathcal{N}_{L/K}(I) = \prod_{1\le i\le n}\sigma_i(I)$; hence, in particular, there exists $\sigma \in \mathrm{Gal}(L/K)$ such that $\mathfrak{P} \mid \sigma(I)$, and hence $\sigma^{-1}(\mathfrak{P}) \mid I$, as claimed.

Assume now that the lemma is true for some $k \ge 1$, and assume that $\prod_{1\le i\le k+1}\mathfrak{P}_i$ divides J. In particular, the product of the first k primes divides J; hence by our induction hypothesis there exist $\sigma_i \in \mathrm{Gal}(L/K)$ such that $\prod_{1\le i\le k}\sigma_i(\mathfrak{P}_i)$ divides I. It follows that

$$\mathcal{N}_{L/K}\left(I/\prod_{1\le i\le k}\sigma_i(\mathfrak{P}_i)\right) = \mathfrak{a}\mathcal{N}_{L/K}\left(J/\prod_{1\le i\le k}\mathfrak{P}_i\right) .$$

We conclude by applying the case $k = 1$ proved above. □

Proof of Theorem 7.5.4. The inclusions $\mathcal{N}_{L/K}(U_S(L)) \subset \mathcal{N}_{L/K}(L^*) \cap U_S(K)$ and $\mathcal{N}_{L/K}(\mathbb{Z}_{L,S}) \subset \mathcal{N}_{L/K}(L^*) \cap \mathbb{Z}_{K,S}$ are trivial. Conversely, let $a \in \mathcal{N}_{L/K}(L^*) \cap \mathbb{Z}_{K,S}$, and let x, y in L be such that $\mathcal{N}_{L/K}(x/y) = a$. We may, of course, assume that x and y are in $\mathbb{Z}_{K,S}$ (in fact, in $\mathbb{Z}_K$ if desired). By definition, we have $\mathcal{N}_{L/K}(x) = a\mathcal{N}_{L/K}(y)$. Let $y\mathbb{Z}_L = \prod\mathfrak{P}_i$ be the prime ideal factorization of the principal ideal $y\mathbb{Z}_L$, with repeated prime ideals $\mathfrak{P}_i$ if necessary. By Lemma 7.5.5 there exist conjugates $\sigma_i(\mathfrak{P}_i)$ of $\mathfrak{P}_i$ and an S-integral ideal J such that

$$x\mathbb{Z}_L = J\prod_i \sigma_i(\mathfrak{P}_i) .$$

By the hypothesis of the theorem, the prime ideals of L above those of S_0 generate the relative class group $Cl_i(L/K)$. It follows that each of the ideals $\mathfrak{P}_i$ can be written in the form

$$\mathfrak{P}_i = \alpha_i\mathfrak{a}_i I_i$$

with $\alpha_i \in L^*$, $\mathfrak{a}_i$ ideal of K, and I_i a product of prime ideals of L above prime ideals in S. Therefore,

$$y\mathbb{Z}_L = \prod_i \alpha_i \prod_i (\mathfrak{a}_i\mathbb{Z}_L)\prod_i I_i .$$

Since the $\mathfrak{a}_i$ are fixed by $\mathrm{Gal}(L/K)$, and since a product of prime ideals above those of S is transformed into another such product by $\mathrm{Gal}(L/K)$, it also follows that

$$x\mathbb{Z}_L = J\prod_i \sigma_i(\alpha_i)\prod_i (\mathfrak{a}_i\mathbb{Z}_L)\prod_i I_i' .$$

Set

$$u = \frac{x}{y} \frac{\prod_i \alpha_i}{\prod_i \sigma_i(\alpha_i)} \ .$$

It is clear that $\mathcal{N}_{L/K}(u) = \mathcal{N}_{L/K}(x/y) = a$. On the other hand, the above formulas for $x\mathbb{Z}_L$ and $y\mathbb{Z}_L$ show that

$$u\mathbb{Z}_L = J \prod_i (I_i'/I_i) \ .$$

Since I_i and I_i' are products of prime ideals above those of S, it follows that I_i'/I_i is an S-integral ideal. Hence, since J is also an S-integral ideal, $u \in \mathbb{Z}_{L,S}$, proving the second equality of the theorem.

If, in addition, we have $a \in U_S(K)$, then necessarily $u \in U_S(L)$. Indeed, if $\mathfrak{P}$ is a prime ideal of L above a prime ideal $\mathfrak{p}$ of S such that $v_{\mathfrak{P}}(u) > 0$, then $\mathfrak{p} \mid \mathcal{N}_{L/K}(\mathfrak{P}) \mid \mathcal{N}_{L/K}(u) = a$, which is absurd, thus proving the first equality of the theorem. □

It is not difficult to prove more precise statements than Theorem 7.5.4. In particular, we have the following result.

Proposition 7.5.6. *Let L/K be a Galois extension, and let r be an integer such that* $\mathrm{Gal}(L/K)$ *can be generated by r elements. In addition, for any finite set S of prime ideals of K, denote by $Cl_{i,S}(L/K)$ the quotient of $Cl_i(L/K)$ by the group generated by the classes of ideals divisible only by prime ideals of L above the ideals of S. Then for all S the quotient group $(\mathcal{N}_{L/K}(L^*) \cap U_S(K))/\mathcal{N}_{L/K}(U_S(L))$ is a subquotient (in other words, a quotient of a subgroup) of $Cl_{i,S}(L/K)^r$.*

I refer to [Sim1] for the proof.

Corollary 7.5.7. *Let L/K be a Galois extension, and let S_0 be a set of prime ideals such that $|Cl_{i,S_0}(L/K)|$ is coprime to $[L : K]$. Then S_0 is suitable for the extension L/K.*

Proof. This follows immediately from the above proposition and Proposition 7.5.1. This corollary is a slight strengthening of part of Theorem 7.5.4, which asserts only that S_0 is suitable if $|Cl_{i,S_0}(L/K)| = 1$. □

Remark. If L/K is not only Galois but also cyclic, then C. Chevalley's "ambiguous class number formula" gives explicitly for all S the quotient $(\mathcal{N}_{L/K}(L^*) \cap U_S(K))/\mathcal{N}_{L/K}(U_S(L))$ in terms of ambiguous class groups, see [Che] and [Sim1].

7.5.3 The Non-Galois Case

In the non-Galois case, the situation is much more complicated because we will need to look at many different class groups of large degree fields. The

method is useful when the degree of the Galois closure is not too large, but it becomes impractical if it is too large.

Let L/K be a Galois extension, and let N/K be its Galois closure in some algebraic closure $\overline{K}$ of K (recall that if $L = K(\alpha)$, then N can be taken as the field generated over K by all the conjugates of α). Write $G = \mathrm{Gal}(N/K)$, $H = \mathrm{Gal}(N/L) \subset G$, and $n = [L : K]$.

We begin with the following.

Proposition 7.5.8. *Let S be a finite set of primes of K. The exponent of $(\mathcal{N}_{L/K}(L^*) \cap U_S(K))/\mathcal{N}_{L/K}(U_S(L))$ divides the GCD $(n, |H| \cdot |Cl_{i,S}(N/K)|)$.*

Proof. By Proposition 7.5.1, we already know that this exponent divides n. Set $h = |Cl_{i,S}(N/K)|$, and let $a \in \mathcal{N}_{L/K}(L^*) \cap U_S(K)$, so that we may write $a = \mathcal{N}_{L/K}(x)$ for some $x \in L^*$. We have

$$a^{|H|} = \mathcal{N}_{L/K}\left(x^{|H|}\right) = \mathcal{N}_{L/K}(\mathcal{N}_{N/L}(x)) = \mathcal{N}_{N/K}(x) \ ;$$

hence $a^{|H|} \in \mathcal{N}_{N/K}(N^*) \cap U_{K,S}$. Applying Proposition 7.5.6 to the Galois extension N/K, we deduce in particular that the exponent of the quotient group $(\mathcal{N}_{N/K}(N^*) \cap U_{K,S})/\mathcal{N}_{N/K}(U_{N,S})$ divides h; hence there exists $s \in U_{N,S}$ such that

$$a^{h|H|} = \mathcal{N}_{N/K}(s) = \mathcal{N}_{L/K}(\mathcal{N}_{N/L}(s)) \ ,$$

so $a^{h|H|} \in \mathcal{N}_{L/K}(U_S(L))$, as claimed. □

Corollary 7.5.9. *Keep the above notation, and let S_0 be a finite set of primes of K. Assume that $n = [L : K]$ is coprime to $|H| = |\mathrm{Gal}(N/L)|$ and to $|Cl_{i,S_0}(N/K)|$. Then S_0 is suitable for L/K.*

Although not general, this corollary is already sufficiently powerful in many cases. Note first that it covers the Galois case in the more precise form of Corollary 7.5.7 (here $|H| = 1$ and $N = L$). But if $n = [L : K]$ is prime, then $[N : K]$ divides $n!$, hence $|H|$ divides $(n-1)!$ and thus is coprime to n, so the condition of the corollary in this case is simply $n \nmid |Cl_{i,S_0}(N/K)|$. In fact, in relative degree $n \leq 5$, the only cases where this corollary cannot be used are the cases where $G \simeq D_4$ ($n = 4$, $|H| = 2$) and $G \simeq S_4$ ($n = 4$, $|H| = 6$).

Note that (unfortunately) it is easy to give examples where it is essential to use the relative class group of N and not only of L.

The general result that we need in the non-Galois case, due to D. Simon, is the following.

Theorem 7.5.10. *Keep the above notation, and let S_0 be a finite set of primes of K* containing all the prime ideals of K ramified in L/K. *Assume that $|Cl_{i,S_0}(N/K)|$ is coprime to n and that for all cyclic subgroups C of $G = \mathrm{Gal}(N/K)$ of prime power order p^a with $p \mid (n, |H|)$, $\left|Cl_{i,S_0}(N^C/K)\right|$ is coprime to $(n, |H|)$. Then S_0 is suitable for L/K.*

We refer to [Sim1] for the (quite technical) proof. Please note the condition that S_0 must contain all the ramified prime ideals, which did not occur in the previous results. It is not difficult to give examples showing that this condition (or a similar one) is necessary (see [Sim1]). Note also that, at least in the known proof, it is necessary to consider the cyclic subgroups of G, and not only of H.

In the special case of D_4 and S_4 extensions, which are not covered by Corollary 7.5.9, the above theorem can be refined to give the following results.

Proposition 7.5.11. *Assume that* $n = [L : K] = 4$ *and that* $G = \mathrm{Gal}(N/K) \simeq D_4$. *Assume that* S_0 *contains all the ramified primes in* L/K *and that* $Cl_{i,S_0}(N/K)$ *and* $Cl_{i,S_0}(L/K)$ *have odd order. Then* S_0 *is suitable for* L/K.

Proposition 7.5.12. *Assume that* $n = [L : K] = 4$ *and that* $G = \mathrm{Gal}(N/K) \simeq S_4$. *Call* C *any one of the three subgroups of* $H = \mathrm{Gal}(N/L)$ *of order* 2. *Assume that* S_0 *contains all the ramified primes in* L/K *and that* $Cl_{i,S_0}(N/K)$ *and* $Cl_{i,S_0}(N^C/K)$ *have odd order. Then* S_0 *is suitable for* L/K.

The interested reader can find many more results of this type as well as examples and counterexamples in [Sim1].

7.5.4 Algorithmic Solution of Relative Norm Equations

We now have all the theoretical and practical tools necessary to give algorithms for solving relative norm equations. We first write the two simple algorithms for determining the necessary set S_0.

Subalgorithm 7.5.13 (Compute S_0 for Galois Extensions). Given a Galois extension L/K, this auxiliary algorithm computes the set S_0 necessary for solving norm equations.

1. [Compute $Cl_i(L/K)$] Using Algorithm 7.3.1 or the methods of Section 7.3.3, compute the relative class group $Cl_i(L/K)$, and let $\overline{I_i}$ be generators of the $[L : K]$-part of $Cl_i(L/K)$ (in other words, the part of $Cl_i(L/K)$ involving only primes dividing $[L : K]$).
2. [Modify and terminate] If desired, multiply the ideals I_i by pseudo-elements of L so that they become prime ideals. Output the set S_0 of prime ideals of K below a prime ideal of L dividing one of the I_i and terminate the subalgorithm.

Subalgorithm 7.5.14 (Compute S_0 for Non-Galois Extensions). Given a non-Galois extension L/K of degree n, this auxiliary algorithm computes the set S_0 necessary for solving norm equations.

1. [Compute Galois closure] If $L = K(\alpha)$, let α_i be the conjugates of α in $\overline{K}$, and set $N \leftarrow K(\alpha_1, \ldots, \alpha_n)$. Using an algorithm for Galois group computation (see below), compute the Galois group $G = \mathrm{Gal}(N/K)$, as well as the subgroup $H = \mathrm{Gal}(N/L)$ fixing L.
2. [Compute subgroup list] Compute the list $\mathcal{C}$ of all cyclic subgroups of G of prime power order p^a for $p \mid (n, |H|)$ up to conjugacy.
3. [Compute relative class groups] Using Algorithm 7.3.1 or the methods of Section 7.3.3, compute the n-part of the relative class group $Cl_i(N/K)$ and the $(n, |H|)$-part of the relative class groups $Cl_i(N^C/K)$ for all $C \in \mathcal{C}$.
4. [Compute initial S_0] If desired, multiply the generators of the relative class groups obtained in step 3 by pseudo-elements of L so that they become prime ideals. Let S_0 be the set of prime ideals of K below a prime ideal of L dividing one of the generators of all the class groups found.
5. [Add ramified primes and terminate] For each prime ideal $\mathfrak{p}$ of K ramified in L/K, set $S_0 \leftarrow S_0 \cup \{\mathfrak{p}\}$, output S_0, and terminate the subalgorithm.

Remarks

(1) Note that this subalgorithm can be used only in very small cases (say, $|G| \leq 24$) because of the difficulty of computing $Cl_i(N/K)$ when $|G| = [N : K]$ is large. This is in marked contrast to the Galois case.

(2) To compute the Galois group G, it is easy to adapt the methods given in [Coh0, Section 6.3] to the relative case (Exercise 12). Since [Coh0] treats only degrees up to 7, we refer to [Eic-Oli] and [Gei] for degrees up to 12. Note that here the degree refers to the degree of L/K and not, of course, to the degree of N/K.

We can now write the main algorithm for solving norm equations, whose proof is immediate from the results of the preceding sections.

Algorithm 7.5.15 (Solving Relative Norm Equations). Let L/K be an extension of number fields and $a \in K^*$. This algorithm finds an $x \in L^*$ such that $a = \mathcal{N}_{L/K}(x)$, or outputs a message saying that x does not exist.

1. [Compute S_0] Using either Algorithm 7.5.13 if L/K is Galois or Algorithm 7.5.14 if L/K is non-Galois, compute a suitable set S_0 of prime ideals of K.
2. [Compute S] Set $S \leftarrow S_0$. Using Algorithm 2.3.22, factor the ideal $a\mathbf{Z}_K$ into a power product of prime ideals, and for every $\mathfrak{p} \mid a$, set $S \leftarrow S \cup \{\mathfrak{p}\}$.
3. [Compute S-units] Using Algorithm 7.4.8, compute S-units $\varepsilon_0, \ldots, \varepsilon_s$ of K and S-units $\eta_0, \ldots, \eta_t$ of L such that

$$U_S(K) = (\mathbf{Z}/w(K)\mathbf{Z})\varepsilon_0 \oplus \bigoplus_{1 \leq i \leq s} \mathbf{Z}\varepsilon_i$$

and

$$U_S(L) = (\mathbf{Z}/w(L)\mathbf{Z})\eta_0 \oplus \bigoplus_{1 \leq j \leq t} \mathbf{Z}\eta_j \ .$$

4. [Compute discrete logarithms] Using a discrete logarithm algorithm in $U_S(K)$ (Exercise 13), compute exponents y_i and $p_{i,j}$ such that

$$a = \prod_{0 \le i \le s} \varepsilon_i^{y_i} \quad \text{and} \quad \mathcal{N}_{L/K}(\eta_j) = \prod_{0 \le i \le s} \varepsilon_i^{p_{i,j}} \ .$$

5. [Solve system] Using Algorithm 4.1.23, look if there exists a solution $X = (x_j)$ to the mixed system of linear equations and congruences $\sum_{0 \le j \le t} p_{i,j} x_j = y_i$ for $1 \le i \le s$ and $\sum_{0 \le j \le t} p_{0,j} x_j \equiv y_0 \pmod{w(K)}$. If such a solution does not exist, our norm equation has no solution, so terminate the algorithm.
6. [Terminate] Output $x \leftarrow \prod_{0 \le j \le t} \eta_j^{x_j}$ as a solution of our norm equation and terminate the algorithm.

Remarks

(1) Since the norm of a root of unity is again a root of unity, we will clearly have $p_{i,0} = 0$ for $i > 0$.
(2) It is easy to modify the algorithm so that it gives the complete solution of the norm equation (Exercise 14).

Consider finally the case of norm equations where we look specifically for integral solutions, in other words we assume $a \in \mathbf{Z}_K$ and we look for $x \in \mathbf{Z}_L$ such that $\mathcal{N}_{L/K}(x) = a$. We could in fact just as easily treat the case $a \in \mathbf{Z}_{K,S}$ and $x \in \mathbf{Z}_{L,S}$, but since the initial problem is by far the most common, we leave the general case to the reader (Exercise 15).

Although it looks very similar to the preceding problem, the solution is much simpler and we do not need the results of the preceding sections. Indeed, let $a\mathbf{Z}_K = \prod_i \mathfrak{p}_i^{v_i}$ be the prime ideal decomposition of $a\mathbf{Z}_K$, and let $\mathfrak{P}$ be a prime ideal of L dividing the solution $x \in \mathbf{Z}_L$ that we are looking for. Then $\mathcal{N}_{L/K}(\mathfrak{P}) \mid a\mathbf{Z}_K$, so if $\mathfrak{p}$ is the ideal of K below $\mathfrak{P}$ and $f = f(\mathfrak{P}/p)$ is the residual degree of $\mathfrak{P}$, we have $\mathfrak{p}^f \mid a$; hence since $\mathfrak{p}$ is a prime ideal, $\mathfrak{p}$ is one of the $\mathfrak{p}_i$. Thus, we may write

$$x\mathbf{Z}_L = \prod_i \prod_{\mathfrak{P}_{i,j} \mid \mathfrak{p}_i} \mathfrak{P}_{i,j}^{x_{i,j}}$$

for nonnegative integers $x_{i,j}$.

There are two necessary conditions that must be satisfied by the $x_{i,j}$. First we must have $\mathcal{N}_{L/K}(x\mathbf{Z}_L) = a\mathbf{Z}_K$; hence for all i we must have

$$\sum_j f_{i,j} x_{i,j} = v_i \ ,$$

where we have set $f_{i,j} = f(\mathfrak{P}_{i,j}/\mathfrak{p}_i)$.

The second condition is that $\prod_i \prod_{\mathfrak{P}_{i,j} \mid \mathfrak{p}_i} \mathfrak{P}_{i,j}^{x_{i,j}}$ must be a principal ideal. As usual this can be transformed into a linear system by introducing the SNF

of the class group $Cl(L) = \bigoplus_k (\mathbf{Z}/d_k\mathbf{Z})\overline{I_k}$ for some ideals I_k of L. We can write for all (i,j) $\mathfrak{P}_{i,j} = \alpha_{i,j}\prod_k I_k^{a_{i,j,k}}$ for some $\alpha_{i,j} \in L$; hence we obtain the additional linear congruences for all k

$$\sum_{i,j} a_{i,j,k}x_{i,j} \equiv 0 \pmod{d_k} \ .$$

This mixed system of linear equations and congruences can be solved using Algorithm 4.1.23. Note that we must only keep solutions of our system such that the $x_{i,j}$ are all nonnegative. In particular, there are only a finite number of such solutions, since it is clear that $0 \le x_{i,j} \le v_i/f_{i,j}$.

Conversely, if we have a solution $(x_{i,j})$ to this system with $x_{i,j} \ge 0$ for all (i,j), we know that $\prod_i \prod_{\mathfrak{P}_{i,j}|\mathfrak{p}_i} \mathfrak{P}_{i,j}^{x_{i,j}}$ must be a principal ideal $y\mathbf{Z}_L$, say, where $y \in \mathbf{Z}_L$ can be found using the principal ideal algorithm in L, and we also know that $\mathcal{N}_{L/K}(y\mathbf{Z}_L) = \mathcal{N}_{L/K}(y)\mathbf{Z}_K = a\mathbf{Z}_K$, hence $a = \varepsilon\mathcal{N}_{L/K}(y)$ for some unit $\varepsilon \in \mathbf{Z}_K$. For a given solution $(x_{i,j})$, we may only modify y by multiplying it by an ordinary unit $\eta \in U(L)$, hence we must solve the norm equation $\mathcal{N}_{L/K}(\eta) = \varepsilon$. But once again this can be transformed into a linear system: let $U(K) = (\mathbf{Z}/w(K)\mathbf{Z})\varepsilon_0 \oplus \bigoplus_{1\le i\le r}\mathbf{Z}\varepsilon_i$ and $U(L) = (\mathbf{Z}/w(L)\mathbf{Z})\eta_0 \oplus \bigoplus_{1\le i\le R}\mathbf{Z}\eta_i$. Using a discrete logarithm algorithm in $U(K)$ (Algorithm 5.3.10), we compute exponents $u_{i,j}$ such that $\mathcal{N}_{L/K}(\eta_j) = \prod_i \varepsilon_i^{u_{i,j}}$ and exponents b_i such that $\varepsilon = \prod_i \varepsilon_i^{b_i}$. Then if $\eta = \prod_j \eta_j^{x_j}$, we must have $\sum_j u_{i,j}x_j = b_i$ for $1 \le i \le r$ and $\sum_j u_{0,j}x_j \equiv b_0 \pmod{w(K)}$.

To summarize, we can find all solutions to $\mathcal{N}_{L/K}(x) = a$ with $a \in \mathbf{Z}_K$ and $x \in \mathbf{Z}_L$ as follows. We first find the finite number of nonnegative solutions $x_{i,j}$ to the mixed linear system given above. For each such solution, we compute ε and b_i, and we find the (possibly infinite) solutions to $\mathcal{N}_{L/K}(\eta) = \varepsilon$ by solving the mixed linear system that we have just described. Writing all this formally gives the following algorithm, in which we output only one solution.

Algorithm 7.5.16 (Solving Relative Integral Norm Equations). Let L/K be an extension of number fields and $a \in \mathbf{Z}_K$, $a \ne 0$. This algorithm either finds an $x \in \mathbf{Z}_L$ such that $a = \mathcal{N}_{L/K}(x)$ or outputs a message saying that x does not exist. We assume computed the SNF (B, D_B) of $Cl(L)$ with $B = (\overline{I_k})$ and $D_B = \mathrm{diag}(d_k)$, the unit groups $U(K) = (\varepsilon_i)$ and $U(L) = (\eta_j)$ as above, and the necessary information to solve all the corresponding discrete logarithm problems.

1. [Factor a] Using Algorithm 2.3.22, factor the ideal $a\mathbf{Z}_K$ into a power product of prime ideals as $a\mathbf{Z}_K = \prod_i \mathfrak{p}_i^{v_i}$.
2. [Compute $\mathfrak{P}_{i,j}$ and $f_{i,j}$] Using Algorithm 2.4.13, for each $\mathfrak{p}_i \mid a$ compute the prime ideals $\mathfrak{P}_{i,j}$ of L above $\mathfrak{p}_i$, and let $f_{i,j} \leftarrow f(\mathfrak{P}_{i,j}/\mathfrak{p}_i)$ be their residual degrees.
3. [Use principal ideal algorithm] Using the principal ideal algorithm in L, for each pair (i,j) as above compute integers $a_{i,j,k}$ such that $\mathfrak{P}_{i,j} = \alpha_{i,j}\prod_k I_k^{a_{i,j,k}}$ for some $\alpha_{i,j} \in L$, that may be discarded.

4. [Compute discrete logarithms in $U(K)$] Using Algorithm 5.3.10, compute integers $u_{i,j}$ such that $\mathcal{N}_{L/K}(\eta_j) = \prod_i \varepsilon_i^{u_{i,j}}$.
5. [Solve mixed linear system] Using Algorithm 4.1.23, solve the mixed linear system in the unknowns $x_{i,j}$: $\sum_j f_{i,j}x_{i,j} = v_i$ for all i and $\sum_{i,j} a_{i,j,k}x_{i,j} \equiv 0 \pmod{d_k}$ for all k. The solution will be of the form $X_0 + HZ$ for a (not necessarily square) HNF matrix H and Z any integral column vector.
6. [Find nonnegative solutions] By looking at the rows from bottom up and using the fact that H is in HNF, find necessary and sufficient inequalities on the entries of Z so that $X_0 + HZ$ has only nonnegative entries, let $\mathcal{X}$ be the finite list of such vectors $X = X_0 + HZ$, and let $s \leftarrow 0$ (s will be a pointer on the list $\mathcal{X}$).
7. [Compute ideal product] Set $s \leftarrow s + 1$. If $s > |\mathcal{X}|$, output a message saying that our norm equation has no solution, so terminate the algorithm. Otherwise, let $X = (x_{i,j})$ be the sth element of $\mathcal{X}$, and let $I \leftarrow \prod_{i,j} \mathfrak{P}_{i,j}^{x_{i,j}}$.
8. [Find generator] Using the principal ideal algorithm in L, compute $y \in \mathbf{Z}_L$ (which must exist) such that $I = y\mathbf{Z}_L$, and set $\varepsilon \leftarrow a/\mathcal{N}_{L/K}(y)$.
9. [Solve unit system] Using Algorithm 5.3.10, compute integers b_i such that $\varepsilon = \prod_i \varepsilon_i^{b_i}$. Then using Algorithm 4.1.23, compute a solution to the mixed linear system $\sum_j u_{i,j}x_j = b_i$ for $i \geq 1$ and $\sum_j u_{0,j}x_j \equiv b_0 \pmod{w(K)}$. If this system has no solution, go to step 7.
10. [Terminate] Output $x \leftarrow y \prod_j \eta_j^{x_j}$ as a solution to our norm equation and terminate the algorithm.

It is of course easy to modify this algorithm so that it gives the complete solution to the norm equation. We must simply modify step 10 so that if a solution is found we go back to step 7, and modify step 9 so that all solutions of the mixed system are found, and not only one. We leave the details to the reader (Exercise 16).

7.6 Exercises for Chapter 7

1. Let L be a relative extension of K. Show that the map $i_{L/K}$ from $Cl(K)$ to $Cl(L)$ is not always surjective (take $K = \mathbf{Q}$ and $L = \mathbf{Q}(\sqrt{-23})$) and not always injective (take $K = \mathbf{Q}(\sqrt{-23})$ and $L = K(\theta)$, where θ is a root of the cubic polynomial $X^3 - X - 1 = 0$, which is the Hilbert class field of K).
2. Let L be a relative extension of K. By considering the same examples as in Exercise 1, show that the map $\mathcal{N}_{L/K}$ from $Cl(L)$ to $Cl(K)$ is not always injective nor surjective.
3. (E. Friedman) Denote by H_K and H_L the Hilbert class fields of K and L, respectively. Continuing the previous exercise, show the following assertions.
   a) The group $Cl_N(L/K)$ is isomorphic to $\mathrm{Gal}(H_L/LH_K)$ via the Artin reciprocity map Art. In particular, the map $\mathcal{N}_{L/K}$ from $Cl(L)$ to $Cl(K)$ is injective if and only if $LH_K = H_L$.

b) The map $\mathcal{N}_{L/K}$ from $Cl(L)$ to $Cl(K)$ is injective if and only if $[L \cap H_K : K] = |Cl(L)| / |Cl(K)|$.
c) The group $Cl_{N,L}(K)$ is isomorphic to $\mathrm{Gal}(H_K/K)/\mathrm{Gal}(H_K/L \cap H_K)$ via the Artin reciprocity map Art. In particular, the map $\mathcal{N}_{L/K}$ from $Cl(L)$ to $Cl(K)$ is surjective if and only if $L \cap H_K = K$.

4. With the notation of the proof of Proposition 7.2.7, show that $\zeta^{(m,w(L))}$ is a generator of the group of roots of unity of K.
5. Give an example of an extension L/K of number fields for which the integers e_i defined in Proposition 7.2.7 cannot be taken equal to zero (take, for example, $K = \mathbb{Q}(\sqrt{21})$ and $L = K(\theta)$, where θ is a root of the polynomial $X^4 + u^2 = 0$ and u is the fundamental unit of K).
6. Let L be a relative extension of K. Show that the norm map from $U(L)$ to $U(K)/\mu(K)$ is not always surjective (take, for example, $K = \mathbb{Q}(\sqrt{2})$ and $L = K(\sqrt{-1})$).
7. Extend Algorithm 7.3.4 so that it also computes the integers e_i occurring in Proposition 7.2.7.
8. Write a formal algorithm for the computation of $U_N(K)$ using the method explained in the text at the end of Section 7.3.1.
9. Prove the validity of all the algorithms of Section 7.3.1.
10. At the end of Algorithm 7.3.6, show that if we set by convention $c_0 = 0$ and $c_{k+1} = 1$, then we have $c'_i = c_i(w(L), c_{i-1})/(w(L), c_i)$ for $1 \leq i \leq k+1$. Find a similar formula for the b_i.
11. By giving an explicit example, show that Theorem 7.5.4 is false if we replace $Cl_i(L/K)$ by $Cl_N(L/K)$ (this is again another example showing the superiority of $i_{L/K}$ over $\mathcal{N}_{L/K}$ in relative class and unit group definitions).
12. Generalize the algorithms for Galois group computation given in [Coh0, Section 6.3] so that they are also valid in the relative case. (This can constitute a small research project.)
13. Write an algorithm that solves the principal ideal problem in S-class groups (given by Algorithm 7.4.6) and the discrete logarithm problem in the S-unit group (given by Algorithm 7.4.8).
14. Modify Algorithm 7.5.15 so that it outputs in a reasonable manner the complete set of solutions in L^* of the norm equation $\mathcal{N}_{L/K}(x) = a$.
15. Generalize Algorithm 7.5.16 to the case where $a \in \mathbb{Z}_{K,S}$ and $x \in \mathbb{Z}_{L,S}$ for some finite set S of prime ideals of K.
16. Modify Algorithm 7.5.16 so that it outputs in a reasonable manner the complete set of solutions in $\mathbb{Z}_L$ of the norm equation $\mathcal{N}_{L/K}(x) = a$.

8. Cubic Number Fields

In [Coh0, Chapter 5], we studied quadratic fields in great detail. The goal of the present chapter is to do the same for cubic fields. We have already studied them in some detail in [Coh0, Chapter 6], but in the present chapter we will deal with deeper subjects and also show how to generate tables of cubic fields almost as efficiently as tables of quadratic fields. The spirit of this chapter is slightly different from that of the preceding chapters, which essentially deal with relative extensions, but the results are sufficiently important to be included in a textbook. Had they been known when [Coh0] was first published, they would, of course, have been included there.

The initial results of this chapter are due to H. Hasse, but the main results are due to H. Davenport and H. Heilbronn ([Dav-Hei1], [Dav-Hei2]). These have been completed and transformed into efficient algorithms by K. Belabas (see [Bel1], [Bel3]), and I thank him for useful conversations on this subject.

The reader is warned that many of the proofs given in this chapter are essentially elementary but rather tedious, in that they consist in a study of a sometimes large number of special cases. Thus, we strongly advise the reader to skip the proofs and read only the results and algorithms, at least at first.

8.1 General Binary Forms

Before specializing to the cubic case, we consider the general case of binary forms of degree n. Let K be a field (usually $\mathbb{Q}$, $\mathbb{R}$, or $\mathbb{C}$).

A *binary form* of degree n with coefficients in K is a homogeneous polynomial in two variables of degree n with coefficients in K, in other words an expression of the form

$$F(x,y) = \sum_{i=0}^{n} a_i x^{n-i} y^i$$

with $a_i \in K$. We will write $F = (a_0, a_1, \ldots, a_n)$ as an abbreviation for the above notation. In particular, $F = (a,b,c)$ is the binary quadratic form $F(x,y) = ax^2 + bxy + cy^2$, and $F = (a,b,c,d)$ is the binary cubic form $F(x,y) = ax^3 + bx^2y + cxy^2 + dy^3$.

We will write $\Phi_n(K)$ (or simply Φ_n if the field K is understood) for the K-vector space of binary forms of degree n with coefficients in K.

The *roots* of F in the algebraic closure $\overline{K}$ of K are the solutions $(x : y) \in \mathbb{P}_1(\overline{K})$ of $F(x,y) = 0$. If $a_j = 0$ for $0 \leq j < m$ and $a_m \neq 0$, the point at infinity $(1 : 0)$ is a root of order exactly m (if $m = 0$ or, equivalently, if $a_0 \neq 0$, it is of course not a root), and the other roots are the roots in $\overline{K}$ of the polynomial $F(x,1)$ of degree $n - m$. In particular, if K is algebraically closed, F always has exactly n roots in K, counted with multiplicity. Note that the point at infinity is rational over any base field, algebraically closed or not.

Denote by $(\alpha_i : \beta_i) \in \mathbb{P}_1(\overline{K})$ (with $1 \leq i \leq n$) the roots of F in $\overline{K}$. It is easily seen that we can choose representatives in $\mathbb{P}_1(\overline{K})$ so that we have

$$F(x,y) = \prod_{1 \leq i \leq n} (\beta_i x - \alpha_i y) \ .$$

Of course, the choice of representative is not unique: for each i we can change (α_i, β_i) into $(\lambda_i \alpha_i, \lambda_i \beta_i)$ as long as $\prod_{1 \leq i \leq n} \lambda_i = 1$. We will always assume that the representatives of the roots are chosen in this manner.

We define the *discriminant* of the form F by the following formula:

$$\mathrm{disc}(F) = \prod_{1 \leq i < j \leq n} (\alpha_i \beta_j - \alpha_j \beta_i)^2 \ .$$

This makes sense since if we change (α_i, β_i) into $(\lambda_i \alpha_i, \lambda_i \beta_i)$ with $\prod_{1 \leq i \leq n} \lambda_i = 1$, the product is multiplied by

$$\prod_{1 \leq i < j \leq n} (\lambda_i \lambda_j)^2 = \Big(\prod_{1 \leq i \leq n} \lambda_i \Big)^{2n-2} = 1 \ .$$

By Galois theory, it is easy to see that $\mathrm{disc}(F) \in K$. In fact, if $F(x,1)$ is a polynomial of degree exactly equal to n (that is, if $a_0 \neq 0$), then we immediately check that $\mathrm{disc}(F) = \mathrm{disc}(F(x,1))$ with the usual meaning of discriminant.

In degrees up to 3 we have the following formulas:

$$\begin{aligned} \mathrm{disc}(ax + by) &= 1 \ ; \\ \mathrm{disc}(ax^2 + bxy + cy^2) &= b^2 - 4ac \ ; \\ \mathrm{disc}(ax^3 + bx^2 y + cxy^2 + dy^3) &= b^2c^2 - 27a^2d^2 + 18abcd - 4ac^3 - 4b^3 d \ . \end{aligned}$$

If F is a form of degree n and $\gamma = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$ is a 2×2 matrix with entries in K, we define the action of γ on F by

$$F \circ \gamma(x,y) = F(Ax + By, Cx + Dy) \ .$$

Proposition 8.1.1. *Let* $\gamma = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$. *Then*

$$\mathrm{disc}(F \circ \gamma) = (AD - BC)^{n(n-1)} \, \mathrm{disc}(F) \ .$$

Proof. Let $(\alpha_i : \beta_i)$ be the roots of F chosen as above so that

$$F(X,Y) = \prod_{1\le i\le n} (\beta_i x - \alpha_i y) \ .$$

Then

$$F\circ\gamma(X,Y) = \prod_{1\le i\le n} (\beta_i(Ax+By) - \alpha_i(Cx+Dy)) = \prod_{1\le i\le n} (\beta_i' x - \alpha_i' y) \ ,$$

with

$$\begin{pmatrix}\alpha_i'\\ \beta_i'\end{pmatrix} = \begin{pmatrix} D & -B\\ -C & A\end{pmatrix}\begin{pmatrix}\alpha_i\\ \beta_i\end{pmatrix} \ .$$

Hence,

$$\begin{aligned}\operatorname{disc}(F\circ\gamma) &= \prod_{1\le i<j\le n} (\alpha_i'\beta_j' - \alpha_j'\beta_i')^2\\ &= \prod_{1\le i<j\le n} ((AD-BC)(\alpha_i\beta_j - \alpha_j\beta_i))^2\\ &= (AD-BC)^{n(n-1)}\operatorname{disc}(F) \ .\end{aligned}$$

□

This proposition implies that the discriminant is *invariant* under the action of $\mathrm{GL}_2(\mathbb{Z})$. More precisely, since it is a polynomial of degree $2n-2$ in the variables of the form, and since the exponent of $(AD-BC)$ in the transformation formula is $n(n-1)$, we say that it is an invariant of degree $2n-2$ and weight $n(n-1)$.

More generally, we can give the following definition.

Definition 8.1.2. (1) *A map f from Φ_n to K is called an* invariant *of degree d and weight w if it is a homogeneous polynomial map of degree d in the coefficients of the forms such that for all $\gamma = \left(\begin{smallmatrix} A & B\\ C & D\end{smallmatrix}\right) \in \mathrm{GL}_2(K)$ and for all $F \in \Phi_n$ we have*

$$f(F\circ\gamma) = (AD-BC)^w f(F) \ .$$

(2) *More generally, a map f from Φ_n to Φ_m is called a* covariant *of degree d and weight w if it is a homogeneous polynomial map of degree d in the coefficients of the forms such that for all γ as above and for all $F \in \Phi_n$ we have*

$$f(F\circ\gamma) = (AD-BC)^w f(F)\circ\gamma \ .$$

Thus, an invariant is the special case $m = 0$ of a covariant. A trivial but important covariant for $m = n$, degree 1, and weight 0 is the identity map, which we will denote by I. By Proposition 8.1.1, the discriminant is an

invariant of weight $n(n-1)$, and it is not difficult to see that its degree is equal to $2n-2$ (see Exercise 1).

Apparently, there are three numbers associated to a covariant on Φ_n: its degree d, its weight w, and the degree m of the image forms. In fact, it is not difficult to show that these numbers are linked by the simple relation

$$w = \frac{nd-m}{2} ,$$

where the 2 in the denominator comes from the fact that we deal with *binary* forms (see Exercise 2). Since a product of covariants is clearly again a covariant, and since the degree, weight, and m are additive, we will consider the algebra of covariants of Φ_n as a *bigraded* algebra, the bidegree being the pair (d,m). We can recover the weight from the above relation.

More generally, the following proposition allows us to construct new covariants from old (the case $h=0$ corresponds to the product of covariants). I thank J. Cremona for having pointed out to me the existence of such a result.

Proposition 8.1.3. *Let f_1 and f_2 be two covariants on Φ_n of degree d_1, d_2, weight w_1, w_2, and with values in Φ_{m_1} and Φ_{m_2}, respectively. For any nonnegative integer h and $F \in \Phi_n$, set*

$$\phi_h(f_1,f_2)(F) = \sum_{j=0}^{h}(-1)^j\binom{h}{j}\frac{\partial^h}{\partial X^{h-j}\partial Y^j}f_1(F)\frac{\partial^h}{\partial X^j\partial Y^{h-j}}f_2(F) .$$

Then $\phi_h(f_1,f_2)$ is a covariant on Φ_n of degree d_1+d_2, weight w_1+w_2+h, with values in $\Phi_{m_1+m_2-2h}$.

Proof. To simplify notation, write ∂_X for $\partial/\partial X$ and ∂_Y for $\partial/\partial Y$. The only operators that occur are ∂_X and ∂_Y, which commute, and the multiplication operator, which does not commute with ∂_X or ∂_Y. Let $\gamma = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) \in \mathrm{GL}_2(K)$ and set $G = \phi_h(f_1,f_2)(F\circ\gamma)$. Then

$$\begin{aligned}
G &= \sum_{j=0}^{h}(-1)^j\binom{h}{j}(A\partial_X+C\partial_Y)^{h-j}(B\partial_X+D\partial_Y)^j f_1(F\circ\gamma) \\
&\quad\cdot(A\partial_X+C\partial_Y)^j(B\partial_X+D\partial_Y)^{h-j}f_2(F\circ\gamma) \\
&= (AD-BC)^{w_1+w_2}\big(-(B\partial_X+D\partial_Y)f_1(F)\circ\gamma(A\partial_X+C\partial_Y)f_2(F) \\
&\quad+(A\partial_X+C\partial_Y)f_1(F)(B\partial_X+D\partial_Y)f_2(F)\circ\gamma\big)^h \\
&= (AD-BC)^{w_1+w_2+h}\left(\partial_X f_1(F)\partial_Y f_2(F)-\partial_Y f_1(F)\partial_X f_2(F)\right)^h\circ\gamma \\
&= (AD-BC)^{w_1+w_2+h}\phi_h(f_1,f_2)(F)\circ\gamma .
\end{aligned}$$

In addition, it is clear that each term of $\phi_h(f_1,f_2)(F)$ is in $\Phi_{m_1+m_2-2h}$ and is of degree d_1+d_2 in the coefficients of the form, proving the proposition. □

It can be proved (see [Gor]) that this proposition allows us to construct *all* the covariants. For our purposes, we simply isolate two special cases.

Corollary 8.1.4. (1) *The function defined by*

$$H(F) = \phi_2(I,I)(F)/2 = \frac{\partial^2 F}{\partial X^2}\frac{\partial^2 F}{\partial Y^2} - \left(\frac{\partial^2 F}{\partial X \partial Y}\right)^2$$

(*called the* Hessian) *is a covariant of degree* 2, *weight* 2, *with values in* Φ_{2n-4}.

(2) *If* f_1 *and* f_2 *are two covariants on* Φ_n *of degree* d_1, d_2, *weight* w_1, w_2, *and with values in* Φ_{m_1} *and* Φ_{m_2}, *respectively, the function defined by*

$$\phi_1(f_1,f_2)(F) = \frac{\partial f_1(F)}{\partial X}\frac{\partial f_2(F)}{\partial Y} - \frac{\partial f_1(F)}{\partial Y}\frac{\partial f_2(F)}{\partial X}$$

(*called the* Jacobian *of the covariants* f_1 *and* f_2) *is a covariant of degree* d_1+d_2, *weight* w_1+w_2+1, *with values in* $\Phi_{m_1+m_2-2}$.

(3) *In particular,*

$$J(F) = \phi_1(I,H) = \frac{\partial F}{\partial X}\frac{\partial H(F)}{\partial Y} - \frac{\partial F}{\partial Y}\frac{\partial H(F)}{\partial X}$$

is a covariant of degree 3, *weight* 3, *with values in* Φ_{3n-6}, *which we can call the* Jacobian covariant *of* F.

Let us specialize to forms of degree $n \leq 3$.

In degree 1, we have already seen that the discriminant is equal to 1, and it is trivial to show that the covariants are all constant multiples of I^k, which is of degree k, weight 0, and with values in Φ_k. Thus, the bigraded algebra of covariants is equal to $K[I]$, where I is of bidegree $(1,1)$.

In degree 2, we already have the covariants I and the discriminant disc. For example, the Hessian $H(F)$ is equal to $-\operatorname{disc}(F)$. It is also easily proved that I and disc generate all covariants. More precisely, the space of covariants with values in Φ_{2k} and of weight 2ℓ is one-dimensional and generated by $I^k \operatorname{disc}^\ell$, and the degree is necessarily equal to $k+2\ell$. There are no nonzero covariants with values with Φ_m for m odd or of odd weight. In other words, the bigraded algebra of covariants is equal to $K[I,\operatorname{disc}]$, where I is of bidegree $(1,2)$ and disc of bidegree $(2,0)$.

In degree 3, as in degree 2, we already have the covariants I and the discriminant disc. However, Corollary 8.1.4 allows us to construct some new covariants. First we have the Hessian $H(F)$, given explicitly by the formula

$$H(ax^3+bx^2y+cxy^2+dy^3) = -4(Px^2+Qxy+Ry^2)$$

with

$$P = b^2-3ac, \quad Q = bc-9ad, \quad R = c^2-3bd\ .$$

(To remove this factor 4, we will in fact set $H_F = -H(F)/4$ and call it by abuse of language the Hessian of F, but for the moment we keep the above normalization.)

This is a covariant of degree 2, weight 2, with values in Φ_2.

The function $J(F) = \phi_1(I, H)$ is a covariant of degree 3, weight 3, with values in Φ_3, so that $J(F)$ is another cubic form. We will simply call it *the* cubic covariant of F. It is given by

$$J(F) = -4(a'x^3 + b'x^2y + c'xy^2 + d'y^3)$$

with

$$a' = -27a^2d + 9abc - 2b^3 \ , \quad b' = -27abd + 18ac^2 - 3b^2c \ ,$$
$$c' = 27acd - 18b^2d + 3bc^2 \ , \quad d' = 27ad^2 - 9bcd + 2c^3 \ .$$

Direct computation shows that $\mathrm{disc}(J(F)) = 2^8 3^6 \mathrm{disc}(F)^3$, $H(J(F)) = 2^4 3^3 \mathrm{disc}(F) H(F)$, and $J(J(F)) = -2^8 3^6 \mathrm{disc}(F)^2 F$.

Thus, we do not obtain any new covariants, and it can indeed be shown that all the covariants are generated by I, disc, H, and J. Slightly more subtle is the fact that there exists a *syzygy*, which is by definition a relation between these covariants. This relation is given by

$$J^2 + H^3 + 2^4 3^3 I^2 \mathrm{disc} = 0 \ .$$

One can show that this is the *only* relation. Thus, the bigraded algebra of covariants is equal to

$$K[I, H, J, \mathrm{disc}]/(J^2 + H^3 + 2^4 3^3 I^2 \mathrm{disc}) \ ,$$

where I is of bidegree $(1,3)$, H is of bidegree $(2,2)$, J is of bidegree $(3,3)$, and disc is of bidegree $(4,0)$. Thus, it is no longer a free polynomial algebra over K.

We now restrict to the case of *integral* binary forms, in other words to binary forms $F(x,y) = \sum_{i=0}^{n} a_i x^{n-i} y^i$ with $a_i \in \mathbb{Z}$ for all i. Since $n(n-1)$ is even, Proposition 8.1.1 tells us that the action of $\mathrm{GL}_2(\mathbb{Z})$ preserves the discriminant of F. This would also be the case for any covariant or invariant of even weight.

We will say that the form F is *irreducible* if $F(x,y)$ is irreducible as a polynomial in $\mathbb{Q}[x,y]$. Equivalently, F is irreducible if $a_0 \neq 0$ and the polynomial $F(x,1)$ is irreducible in $\mathbb{Q}[x]$ (or $\mathbb{Z}[x]$).

We will say that an integral form F is *primitive* if the GCD of all its coefficients is equal to 1.

Proposition 8.1.5. *Let F be an integral form and $\gamma \in \mathrm{GL}_2(\mathbb{Z})$. Then $F \circ \gamma$ is irreducible if and only if F is irreducible, and $F \circ \gamma$ is primitive if and only F is primitive.*

Proof. This immediately follows from the fact that the action of $\mathrm{GL}_2(\mathbb{Z})$ is reversible, that is, $F = (F \circ \gamma) \circ \gamma^{-1}$. □

8.2 Binary Cubic Forms and Cubic Number Fields

The aim of this section is to generalize to the cubic case the well-known correspondence between binary quadratic forms and quadratic number fields. These results are due to Davenport and Heilbronn (see [Dav-Hei1] and [Dav-Hei2]). Before stating and proving the main theorem, we need a few preliminary results. We let Φ be the set of classes under $\mathrm{GL}_2(\mathbb{Z})$ of primitive irreducible binary cubic forms. This makes sense, thanks to Proposition 8.1.5. Note that we consider classes under $\mathrm{GL}_2(\mathbb{Z})$, and not under $\mathrm{SL}_2(\mathbb{Z})$ as in the quadratic case.

Let K be a cubic number field and $(1,\alpha,\beta)$ an integral basis of $\mathbb{Z}_K$ with first element equal to 1. Denote by $d(K)$ the discriminant of the number field K, and let K^g be a normal closure of K, which is equal to K itself if K is a cyclic cubic field and otherwise is a number field of degree 6 over $\mathbb{Q}$. If $x \in K$, denote by $x_1 = x$, x_2, and x_3 the conjugates of x in K^g, and let $\mathrm{disc}(x)$ be the discriminant of the (monic) characteristic polynomial of x, so that $\mathrm{disc}(x) = ((x_1-x_2)(x_1-x_3)(x_2-x_3))^2$.

Note that if $x \in K$ and $x = x'/d$ with $x' \in \mathbb{Z}_K$ and $d \in \mathbb{Z}$, then $\mathrm{disc}(x) = \mathrm{disc}(x')/d^6 = d(K)f^2/d^6$, where $f = [\mathbb{Z}_K : \mathbb{Z}[x']]$, so that for every $x \in K$ we have $\sqrt{\mathrm{disc}(x)/d(K)} \in \mathbb{Q}$.

Proposition 8.2.1. *Let $\mathcal{B} = (1,\alpha,\beta)$ be an integral basis of a cubic number field K as above. For x and y elements of $\mathbb{Q}$, set*

$$F_{\mathcal{B}}(x,y) = \frac{\prod_{1\le i<j\le 3}((\beta_i-\beta_j)x-(\alpha_i-\alpha_j)y)}{\sum_{1\le i<j\le 3}(-1)^{i-j}(\alpha_i\beta_j-\alpha_j\beta_i)} .$$

(1) *For x and y in $\mathbb{Q}$, we have*

$$F_{\mathcal{B}}(x,y) = \pm\sqrt{\frac{\mathrm{disc}(\beta x-\alpha y)}{d(K)}} = \pm\frac{\mathcal{N}_{K/\mathbb{Q}}((\beta-\beta')x-(\alpha-\alpha')y)}{\sqrt{d(K)}} .$$

(2) *The function $F_{\mathcal{B}}$ is the restriction to $\mathbb{Q}\times\mathbb{Q}$ of a binary cubic form (again denoted by $F_{\mathcal{B}}$) with rational coefficients.*

(3) $\mathrm{disc}(F_{\mathcal{B}}) = d(K)$.

(4) *The form $F_{\mathcal{B}}$ is an integral, primitive, irreducible cubic form.*

(5) *The class of $F_{\mathcal{B}}$ in Φ is independent of the integral basis $(1,\alpha,\beta)$ that we have chosen, so we will denote this class by F_K.*

(6) *Let the number field K be defined by a root θ of the polynomial x^3+px^2+qx+r with p, q, r in $\mathbb{Z}$, such that there exists an integral basis of the form $(1,\theta,(\theta^2+t\theta+u)/f)$ with t, u, f in $\mathbb{Z}$ and $f = [\mathbb{Z}_K : \mathbb{Z}[\theta]]$ (this is always possible). If we choose $\alpha = \theta$ and $\beta = (\theta^2+t\theta+u)/f$, we have explicitly*

$$\begin{aligned} F_{\mathcal{B}}(x,y) = {} & ((t^3-2t^2p+t(q+p^2)+r-pq)/f^2)x^3 \\ & +((-3t^2+4tp-(p^2+q))/f)x^2y+(3t-2p)xy^2-fy^3 . \end{aligned}$$

Proof. (1) follows directly from the definitions, after noting that $d(K)$ is the square of the determinant of the matrix

$$\begin{pmatrix} 1 & \alpha_1 & \beta_1 \\ 1 & \alpha_2 & \beta_2 \\ 1 & \alpha_3 & \beta_3 \end{pmatrix} .$$

If we take one of the expressions on the right-hand side of the equalities in (1) as the definition of $F_{\mathcal{B}}$, there is a sign ambiguity. However, this will not matter in the sequel since $F_{\mathcal{B}}(-x,-y) = -F_{\mathcal{B}}(x,y)$ so that $F_{\mathcal{B}}$ and $-F_{\mathcal{B}}$ are $\mathrm{SL}_2(\mathbb{Z})$-equivalent.

(2). The definition of $F_{\mathcal{B}}$ shows that $F_{\mathcal{B}}$ is the restriction to $\mathbb{Q}\times\mathbb{Q}$ of a cubic form with coefficients in K^g. Furthermore, Galois theory shows that the coefficients of $F_{\mathcal{B}}(x,y)$ are invariant under the Galois group of K^g over $\mathbb{Q}$, so that in fact $F_{\mathcal{B}}$ is a rational cubic form. Note that the rationality of $F_{\mathcal{B}}(x,y)$ when x and y are rational also follows from the remark made before the proposition.

(3). The roots of $\sqrt{d(K)}F_{\mathcal{B}}$ (with any choice for the square root) are by assumption the $(\alpha_i-\alpha_j : \beta_i-\beta_j)$ for $1\le i<j\le 3$, and this respects the convention for the choice of roots made at the beginning. Hence one checks that

$$\begin{aligned} \operatorname{disc}(\sqrt{d(K)}F_{\mathcal{B}}) &= (((\alpha_1-\alpha_2)(\beta_2-\beta_3)-(\alpha_2-\alpha_3)(\beta_1-\beta_2)) \\ &\quad\cdot((\alpha_2-\alpha_3)(\beta_3-\beta_1)-(\alpha_3-\alpha_1)(\beta_2-\beta_3)) \\ &\quad\cdot((\alpha_3-\alpha_1)(\beta_1-\beta_2)-(\alpha_1-\alpha_2)(\beta_3-\beta_1)))^2 \\ &= d(K)^3 \ , \end{aligned}$$

so that $\operatorname{disc}(F_{\mathcal{B}}) = d(K)$, since by the explicit formula for $\operatorname{disc}(F_{\mathcal{B}})$ we have $\operatorname{disc}(\lambda F_{\mathcal{B}}) = \lambda^4 \operatorname{disc}(F_{\mathcal{B}})$.

(4). This is the longest part of the proof. Let $F_{\mathcal{B}} = (a,b,c,d)$ with a, b, c, d in $\mathbb{Q}$ by (2). First, we note that if x and y are in $\mathbb{Z}$, then $\gamma = \beta x - \alpha y$ is in $\mathbb{Z}_K$; hence by the remark made above, $F_{\mathcal{B}}(x,y) = f$, where $f = [\mathbb{Z}_K : \mathbb{Z}[\gamma]]$, so $F_{\mathcal{B}}(x,y) \in \mathbb{Z}$. Applying this to $(x,y) = (1,0)$ and $(x,y) = (0,1)$, we deduce that $a \in \mathbb{Z}$ and $d \in \mathbb{Z}$. Then applying this to $(x,y) = (1,1)$ and $(x,y) = (1,-1)$, we deduce that $b+c \in \mathbb{Z}$ and $b-c \in \mathbb{Z}$. It follows that b and c belong to $\frac{1}{2}\mathbb{Z}$ with $b \equiv c \pmod 1$. By (3), we know that $\operatorname{disc}(F_{\mathcal{B}}) = d(K)$ is an integer. By the explicit formula for $\operatorname{disc}(F_{\mathcal{B}})$ and using the fact that b and c belong to $\frac{1}{2}\mathbb{Z}$, we obtain $2\operatorname{disc}(F_{\mathcal{B}}) \equiv 2b^2c^2 \in \mathbb{Z}$. Since $b \equiv c \pmod 1$, we cannot have $b \equiv c \equiv \frac{1}{2} \pmod 1$; hence $b \equiv c \equiv 0 \pmod 1$, proving that $F_{\mathcal{B}}$ is integral.

Recall that by (1), if x and y are in $\mathbb{Z}$, we have $F_{\mathcal{B}}(x,y) = \pm f(\beta x - \alpha y)$, where for $\gamma \in \mathbb{Z}_K$, $f(\gamma) = [\mathbb{Z}_K : \mathbb{Z}[\gamma]]$. Let δ be the GCD of the coefficients of $F_{\mathcal{B}}$. Thus, for each x, y in $\mathbb{Z}$ we have $\delta \mid f(\beta x - \alpha y)$. But since $f(\gamma+n) = f(\gamma)$ for any $n \in \mathbb{Z}$ and since $(1,\alpha,\beta)$ is an integral basis (and not only a triplet of elements of $\mathbb{Z}_K$), it follows that for *any* $\gamma \in \mathbb{Z}_K$ we have $\delta \mid f(\gamma)$.

This means by definition that δ is an inessential discriminantal divisor. According to a theorem of Dedekind, for a cubic field this implies either that $\delta = 1$ or else that $\delta = 2$ and 2 is totally split in K (see Exercises 4 and 5). Assume the latter. We thus have $F_{\mathcal{B}} = 2G$ for an integral cubic form G, and $\operatorname{disc}(G) = \operatorname{disc}(F_{\mathcal{B}})/16 = d(K)/16$ by (3). However, since 2 is totally split in K, it is in particular unramified, hence $2 \nmid d(K)$, which is absurd since this implies that the integral cubic form G has a nonintegral discriminant. Thus we have $\delta = 1$, and hence $F_{\mathcal{B}}$ is primitive.

Finally, let us show that $F_{\mathcal{B}}$ is irreducible. Since we are in the cubic case, this means that $F_{\mathcal{B}}(x, y)$ has no linear factor in $\mathbb{Q}[x, y]$. Assume the contrary. By definition of $F_{\mathcal{B}}$, we may assume that such a linear factor is proportional to $((\beta_1 - \beta_2)x - (\alpha_1 - \alpha_2)y)$. It follows that there exist integers r and s not both zero such that $s(\alpha_1 - \alpha_2) = r(\beta_1 - \beta_2)$, so that $s\alpha_1 - r\beta_1 = s\alpha_2 - r\beta_2$. Taking conjugates, we see that we have $s\alpha_1 - r\beta_1 = s\alpha_2 - r\beta_2 = s\alpha_3 - r\beta_3$, so that the conjugates of $s\alpha - r\beta$ are equal. Hence by Galois theory, $s\alpha - r\beta \in \mathbb{Q}$, and since r and s are not both zero, this is in contradiction to $(1, \alpha, \beta)$ being an integral basis. It follows that $F_{\mathcal{B}}$ is irreducible.

(5). Let $\mathcal{B}' = (1, \alpha', \beta')$ be another integral basis. This means that there exist integers A, B, C, D, E, F such that

$$\begin{pmatrix} \alpha' \\ \beta' \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \begin{pmatrix} E \\ F \end{pmatrix}$$

with $AD - BC = \pm 1$, since $(\alpha' - E, \beta' - F)$ and (α, β) must generate the same lattice. It follows that

$$\begin{aligned} \operatorname{disc}(\beta' x - \alpha' y) &= \operatorname{disc}((C\alpha + D\beta + F)x - (A\alpha + B\beta + E)y) \\ &= \operatorname{disc}((C\alpha + D\beta)x - (A\alpha + B\beta)y) \\ &= \operatorname{disc}(\beta(Dx - By) - \alpha(-Cx + Ay)) \ , \end{aligned}$$

where the second equality follows from $\operatorname{disc}(x + n) = \operatorname{disc}(x)$ for all $x \in K$ and $n \in \mathbb{Q}$. Since $AD - BC = \pm 1$, it follows from (1) that $F_{\mathcal{B}'} = \pm F_{\mathcal{B}} \circ \gamma^{-1}$ with $\gamma = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$, so $F_{\mathcal{B}'}$ and $F_{\mathcal{B}}$ are equivalent.

(6). This follows from a straightforward but tedious computation. □

Conversely, given an integral, primitive, irreducible binary cubic form F, we can define a number field K_F associated to F by $K_F = \mathbb{Q}(\theta)$, where θ is any root of $F(x, 1)$. Since F is irreducible, θ is an algebraic number of degree exactly equal to 3, so K_F is a cubic field. Choosing another root of F gives an isomorphic (in fact, conjugate) field K_F; hence the isomorphism class of K_F is well-defined. Finally, if F and G are equivalent under $\mathrm{GL}_2(\mathbb{Z})$, K_F and K_G are again clearly conjugate.

It follows that if we let $\mathcal{C}$ be the set of isomorphism classes of cubic number fields, we have defined maps $\phi_{\mathcal{C}\Phi} : K \mapsto F_K$ from $\mathcal{C}$ to Φ and $\phi_{\Phi\mathcal{C}} : F \mapsto K$ from Φ to $\mathcal{C}$.

Proposition 8.2.2. *We have $\phi_{\Phi C} \circ \phi_{C\Phi} = 1$; hence $\phi_{C\Phi}$ is injective and $\phi_{\Phi C}$ is surjective.*

Proof. Let K be a cubic field and let $(1, \alpha, \beta)$ be an integral basis. We have (for example) $K_{F_K} = \mathbb{Q}((\alpha_2 - \alpha_3)/(\beta_2 - \beta_3)) \subset K^g$. If K is a cyclic cubic field, then $K_{F_K} = K$. Otherwise, K^g is a number field of degree 6, with Galois group isomorphic to S_3, and $(\alpha_2 - \alpha_3)/(\beta_2 - \beta_3)$ is fixed by the transposition (23) of order 2, hence belongs to a cubic subfield, and so K_{F_K} is isomorphic to K. In fact, by choosing the numbering such that $\alpha = \alpha_1$ and $\beta = \beta_1$, K_{F_K} is even *equal* to K, not only conjugate to it. Note also that we cannot have $(\alpha_2 - \alpha_3)/(\beta_2 - \beta_3) \in \mathbb{Q}$, because F_K would then be reducible. □

Let $I \subset \Phi$ be the image of $\phi_{C\Phi}$. It follows from this proposition that $\phi_{C\Phi}$ and the restriction of $\phi_{\Phi C}$ to I are inverse discriminant-preserving bijections between C and I, and this is the Davenport–Heilbronn correspondence that we are looking for. There now remains to determine the image I.

Before doing so, we will show that the form F_K determines the simple invariants of a cubic number field.

Proposition 8.2.3. *Let K be a cubic field, $F_K = (a, b, c, d)$ the associated cubic form, and θ a root of F_K such that $K = \mathbb{Q}(\theta)$ (we have seen above that such a θ exists). Then we have the following results.*

(1) $d(K) = \operatorname{disc}(F_K)$.

(2) $(1, a\theta, a\theta^2 + b\theta)$ *is an integral basis of* $\mathbb{Z}_K$.

(3) *A prime $p \in \mathbb{Z}$ decomposes in K as F decomposes in $\mathbb{F}_p[X, Y]$. More precisely, if*

$$F_K(X, Y) \equiv \prod_{1 \le i \le g} \overline{T_i}^{e_i}(X, Y) \pmod{p}$$

is a decomposition of F into irreducible homogeneous factors in $\mathbb{F}_p[X, Y]$, then we have

$$p\mathbb{Z}_K = \prod_{1 \le i \le g} \mathfrak{p}_i^{e_i} ,$$

where the $\mathfrak{p}_i$ are distinct prime ideals of $\mathbb{Z}_K$ given as follows. Call T_i any lift of $\overline{T_i}$ in $\mathbb{Z}[X, Y]$, and set $d_i = \deg(\overline{T_i})$.

a) If $p \nmid a$, then

$$\mathfrak{p}_i = p\mathbb{Z}_K + T_i(\theta, 1)\mathbb{Z}_K .$$

b) If $p \mid a$ but $p \nmid d$, then

$$\mathfrak{p}_i = p\mathbb{Z}_K + \frac{T_i(\theta, 1)}{\theta^{d_i}}\mathbb{Z}_K .$$

c) If $p \nmid a$, $p \mid d$, then if $p \neq 2$ or if $p = 2$ and $F(X, Y) \not\equiv X^2Y + XY^2 \pmod 2$, there exists $e \in \mathbb{Z}$ such that $e \not\equiv 0$ and $e \not\equiv -b/c \pmod p$ (any $e \not\equiv 0 \pmod p$ if $p \mid c$), and then

$$\mathfrak{p}_i = p\mathbf{Z}_K + \frac{T_i(\theta,1)}{(1-e\theta)^{d_i}}\mathbf{Z}_K \ .$$

d) Finally, if $p = 2$ *and* $F(X,Y) \equiv X^2Y + XY^2 \pmod 2$, *we can take*

$$\mathfrak{p}_1 = 2\mathbf{Z}_K + a\theta\mathbf{Z}_K \ , \quad \mathfrak{p}_2 = 2\mathbf{Z}_K + (a\theta^2 + b\theta + 1)\mathbf{Z}_K \ ,$$
$$\textit{and } \mathfrak{p}_3 = 2\mathbf{Z}_K + (a\theta^2 + (a+b)\theta)\mathbf{Z}_K \ .$$

Proof. Statement (1) has been proved in the preceding section.

(2). θ is a root of $a\theta^3 + b\theta^2 + c\theta + d = 0$. It follows from [Coh0, Exercise 15 of Chapter 4], and easily checked directly, that $\mathcal{O} = \{x + ya\theta + z(a\theta^2 + b\theta),\ x,y,z \in \mathbf{Z}\}$ is an order in K; in other words, it is an algebra and a $\mathbf{Z}$-module of finite type and in particular is a suborder of the maximal order $\mathbf{Z}_K$. If θ_i denotes the three roots of $F_K(x,1)$, an easy computation shows that

$$\operatorname{disc}(\mathcal{O}) = a^4 \prod_{1\le i<j\le 3} (\theta_i - \theta_j)^2 = \operatorname{disc}(F_K) = \operatorname{disc}(\mathbf{Z}_K) \ ;$$

hence $\mathcal{O} = \mathbf{Z}_K$.

(3). Assume first that $p \nmid a$. Set

$$f(X) = a^2 F_K\left(\frac{X}{a}, 1\right) = X^3 + bX^2 + acX + a^2d \ .$$

Then f is a monic irreducible polynomial over $\mathbf{Q}$ with a root $a\theta \in K$.

We have $\mathbf{Z}[a\theta] \subset \mathbf{Z}_K$ and

$$\operatorname{disc}(\mathbf{Z}[a\theta]) = a^6 \prod_{1\le i<j\le 3} (\theta_i - \theta_j)^2 = a^2 \operatorname{disc}(\mathbf{Z}_K) \ ;$$

hence $[\mathbf{Z}_K : \mathbf{Z}[a\theta]] = a$ (this also follows directly from (2)). Since $p \nmid a = [\mathbf{Z}_K : \mathbf{Z}[a\theta]]$, it follows, for example, from [Coh0, Theorem 4.8.13] that $p\mathbf{Z}_K = \prod_{1\le i\le g} \mathfrak{p}_i^{e_i}$ for $\mathfrak{p}_i = p\mathbf{Z}_K + U_i(a\theta)\mathbf{Z}_K$, where $f \equiv \prod_{1\le i\le g} \overline{U_i}^{e_i} \pmod p$ is an irreducible decomposition of f in $\mathbf{F}_p[X]$.

But then

$$F_K(X,Y) = Y^3 F_K\left(\frac{X}{Y}, 1\right) = \frac{Y^3}{a^2} f\left(\frac{aX}{Y}\right)$$
$$\equiv \prod_{1\le i\le g} \varepsilon_i \left(Y^{d_i} U_i\left(\frac{aX}{Y}\right)\right)^{e_i} \pmod p$$

for some $\varepsilon_i \in \mathbf{F}_p^*$, and so $T_i(X,Y) = \varepsilon_i Y^{d_i} U_i(aX/Y)$ and $T_i(\theta,1) = \varepsilon_i U_i(a\theta)$. Finally, we note that, for $\varepsilon \not\equiv 0 \pmod p$, we have

$$p\mathbf{Z}_K + \alpha\mathbf{Z}_K = p\mathbf{Z}_K + \varepsilon\alpha\mathbf{Z}_K \ ,$$

and the case $p \nmid a$ follows.

Assume now that $p \mid a$. If we can find a matrix $M = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) \in \mathrm{GL}_2(\mathbf{Z})$ such that $G = F_K \circ M = (a', b', c', d')$ is such that $p \nmid a'$, we can apply the preceding case, since K is also generated by a root of G.

One easily checks that if $p \mid a$ but $p \nmid d$ we can take $M = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, and if $p \mid a$, $p \mid d$ but either $p \neq 2$ or $p = 2$ and $F_K(X,Y) \not\equiv X^2Y + XY^2 \pmod 2$, then there exists $e \in \mathbf{Z}$ such that $e \not\equiv 0$ and $e \not\equiv -b/c \pmod p$ (any $e \not\equiv 0 \pmod p$ if $p \mid c$), and we then take $M = \left(\begin{smallmatrix} 1 & 0 \\ e & 1 \end{smallmatrix}\right)$. This immediately gives the formulas of the proposition.

Finally, if $p = 2$ and $F_K(X,Y) \equiv X^2Y + XY^2 \pmod 2$, then 2 divides the coefficient of x^3 of any form equivalent to F_K, and from the definition of F_K this means that 2 divides the index of any $\alpha \in \mathbf{Z}_K$; in other words, 2 is an inessential discriminantal divisor. We then know that 2 is totally split, hence 2 still factors as $F_K \equiv XY(X+Y)$ modulo 2. To find the factors explicitly, we must split the étale algebra $\mathcal{A} = \mathbf{Z}_K/2\mathbf{Z}_K$. Since $\mathcal{A} \simeq (\mathbf{Z}/2\mathbf{Z})^3$, all its elements are idempotents. If we set $e_1 = 1$, $e_2 = a\theta$, and $e_3 = a\theta^2 + b\theta$ considered as elements of $\mathcal{A}$ (they are in $\mathbf{Z}_K$; see above), we check that $e_2e_3 = a^2\theta^3 + ab\theta^2 = -ac\theta - ad = a\theta = e_2$ in $\mathcal{A}$ since c is odd and a and d are even. It follows that e_2, $e_1 + e_3$, and $e_2 + e_3$ are orthogonal idempotents of sum 1, thus giving the desired splitting of $\mathcal{A}$, hence of $p\mathbf{Z}_K$. □

8.3 Algorithmic Characterization of the Set U

We will now introduce a set U of cubic forms and study some of its properties. In the next section, we will prove that U is the image I of the Davenport–Heilbronn correspondence.

We first need some notation. For a prime p, we let V_p be the set of $F \in \Phi$ such that $p^2 \nmid \mathrm{disc}(F)$ if $p \neq 2$, or $\mathrm{disc}(F) \equiv 1 \pmod 4$ or $\mathrm{disc}(F) \equiv 8$ or $12 \pmod{16}$ if $p = 2$. If other words, if $\mathrm{disc}(F) = d_k f^2$ with d_k a fundamental discriminant, $F \in V_p$ if and only if $p \nmid f$.

In particular, $F \in \bigcap_p V_p$ if and only if $\mathrm{disc}(F)$ is a fundamental discriminant.

Furthermore, let U_p be the set of $F \in \Phi$ such that either $F \in V_p$, or else

$$F(x,y) \equiv \lambda(\delta x - \gamma y)^3 \pmod p$$

for some $\lambda \in \mathbf{F}_p^*$ and x, y in $\mathbf{F}_p$ not both zero, and in addition $F(\gamma,\delta) \not\equiv 0 \pmod{p^2}$.

We will summarize the condition $F(x,y) \equiv \lambda(\delta x - \gamma y)^3 \pmod p$ for some $\lambda \in \mathbf{F}_p^*$ by saying that F has three identical roots in $\mathbf{F}_p$, and we will write $(F,p) = (1^3)$.

Finally, we set $U = \bigcap_p U_p$. The Davenport–Heilbronn theorem states that $U = I$, the image of the map $\phi_{C\Phi}$ that we are looking for. Before proving this theorem, we must study in detail the set U.

For this, we will use the *Hessian* of a form F, introduced in Corollary 8.1.4, which we divide by -4 to avoid useless constants.

Definition 8.3.1. *Let $F = (a,b,c,d)$ be a cubic form. We define the Hessian of F and numbers P, Q, and R by the formula*

$$H_F(x,y) = -\frac{1}{4}\begin{vmatrix} \frac{\partial^2 F}{\partial x \partial x} & \frac{\partial^2 F}{\partial x \partial y} \\ \frac{\partial^2 F}{\partial y \partial x} & \frac{\partial^2 F}{\partial y \partial y} \end{vmatrix} = Px^2 + Qxy + Ry^2 \ .$$

We have $P = b^2 - 3ac$, $Q = bc - 9ad$, and $R = c^2 - 3bd$.

Proposition 8.3.2. *Let $F = (a,b,c,d)$ be a cubic form and $H_F = (P,Q,R)$ its Hessian.*

(1) *For any $M \in \mathrm{GL}_2(\mathbb{C})$ we have $H_{F\circ M} = \det(M)^2 H_F \circ M$. In particular, if $M \in \mathrm{GL}_2(\mathbb{Z})$, we have $H_{F\circ M} = H_F \circ M$.*
(2) $\mathrm{disc}(H_F) = -3\,\mathrm{disc}(F)$.
(3) $F(Q,-2P) = (2bP - 3aQ)\,\mathrm{disc}(F)$.
(4) $3dP - cQ + bR = 3aR - bQ + cP = 0$.

Formula (1) simply states that the Hessian is a covariant of weight 2, which we have already proved (Corollary 8.1.4). The other formulas are easily proved by a direct computation. □

We are now ready to give an algorithmic description of the set U. We need two propositions.

Proposition 8.3.3. *Let $F = (a,b,c,d)$ be a* primitive *form, and let $H_F = (P,Q,R)$ be its Hessian. Recall that we write $(F,p) = (1^3)$ if F has a triple root in $\mathbb{F}_p$. Then we have the following results.*

(1) *$p \mid \mathrm{disc}(F)$ if and only if F has at least a double root in $\overline{\mathbb{F}_p}$, and if this is the case, all the roots of F are in fact in $\mathbb{F}_p$ itself.*
(2) *$(F,p) = (1^3)$ if and only if $p \mid \gcd(P,Q,R)$.*
(3) *If $(F,p) = (1^3)$ and $p \neq 3$, then $F \in U_p$ if and only if $p^3 \nmid \mathrm{disc}(F)$.*
(4) *If $(F,3) = (1^3)$ and $F \in U_3$, then $3^6 \nmid \mathrm{disc}(F)$.*
(5) *If $(F,3) = (1^3)$ then we have the following:*
   a) if $3 \mid a$, then $F \in U_3 \iff 9 \nmid a$ and $3 \nmid d$;
   b) if $3 \nmid a$ but $3 \mid d$, then $F \in U_3 \iff 9 \nmid d$;
   c) if $3 \nmid a$ and $3 \nmid d$, then there exists $\varepsilon = \pm 1$ such that $3 \mid (a - \varepsilon d)$, and then $F \in U_3 \iff 9 \nmid ((a+c) - \varepsilon(b+d))$.

Proof. (1). Assume that $p \mid \mathrm{disc}(F)$. We know that any nonzero polynomial in one variable over $\mathbb{F}_p$ can be written as $\prod_{i\geq 1} A_i^i$, where $A_i \in \mathbb{F}_p[X]$ are pairwise coprime and squarefree polynomials, and essentially in a unique manner (up to multiplication of each A_i by suitable constants). This result can be homogenized and transformed into an identical one for homogeneous polynomials in two variables.

Since F is primitive, it is nonzero modulo p. Since $p \mid \mathrm{disc}(F)$, by definition of the discriminant this means that F has at least a double root in

$\mathbb{P}_1(\overline{\mathbb{F}_p})$. In other words, in the decomposition above there exists $i > 1$ such that A_i is not equal to a constant. Since F is of degree 3, this means that $F = A_3^3$ or else $F = A_1 A_2^2$ with A_1, A_2, and A_3 of degree 1. It follows in particular that all the roots of F modulo p are in $\mathbb{F}_p$ itself.

(2). We have just seen that if $p \mid \mathrm{disc}(F)$ then all the roots of F are in $\mathbb{F}_p$ and there is at least a double root. Hence write

$$F(x,y) \equiv (\delta x - \gamma y)^2(\beta x - \alpha y) \pmod{p} \ .$$

Then we find that

$$H(x,y) \equiv (\delta x - \gamma y)^2(\alpha\delta - \beta\gamma)^2 \pmod{p} \ .$$

Since F is primitive, γ and δ cannot both be zero modulo p; hence

$$H(x,y) \equiv 0 \pmod{p} \iff \alpha\delta - \beta\gamma \equiv 0 \pmod{p} \iff (F,p) = (1^3) \ .$$

On the other hand, if $p \nmid \mathrm{disc}(F)$, then we cannot have $p \mid \gcd(P,Q,R)$ since otherwise $p^2 \mid \mathrm{disc}(H_F) = -3\,\mathrm{disc}(F)$, and so $p \mid \mathrm{disc}(F)$, which is absurd.

(3). From now on we assume that $(F,p) = (1^3)$. Replacing F by an equivalent form G, we may assume that the triple root of G modulo p is at $(0:1)$, so that $G = (A,B,C,D) \equiv (A,0,0,0) \pmod{p}$ for some $A \in \mathbb{Z}$. This implies that $\mathrm{disc}(G) \equiv -27A^2D^2 \pmod{p^3}$. Since G is primitive, we have $p \nmid A$.

Assume first that $p \neq 3$. We thus have

$$p^3 \mid \mathrm{disc}(G) \iff p^2 \mid D \iff p^2 \mid G(0,1) \iff G \notin U_p$$

by definition of U_p (note that $p^3 \mid \mathrm{disc}(G)$ implies that $G \notin V_p$ when $p \neq 3$; see Exercise 6).

We could also have written $F(x,y) = \lambda(\delta x - \gamma y)^3 + pF_1(x,y)$ for an integral form F_1, from which we obtain $\mathrm{disc}(F) \equiv -27\lambda^2 F_1^2(\gamma,\delta)p^2 \pmod{p^3}$, which immediately implies the result.

(4). Assume now that $p = 3$ and that $F \in U_3$ or, equivalently, that $9 \nmid D$. Then $\mathrm{disc}(G) \equiv -4AC^3 \pmod{3^4}$, hence $3^2 \mid C$, so $\mathrm{disc}(G) \equiv -4B^3D \pmod{3^5}$, hence $3^2 \mid B$, so finally $\mathrm{disc}(G) \equiv -27A^2D^2 \not\equiv 0 \pmod{3^6}$.

(5). Assume that $p = 3$ and that $(F,3) = (1^3)$. Since $F(x,y) \equiv \lambda(\delta x - \gamma y)^3 \pmod{3}$, we see that $3 \mid b$ and $3 \mid c$. It follows that

$$F(-d,a) \equiv -ad^3 + da^3 \equiv -ad + da \equiv 0 \pmod{3} \ ,$$

and since F is primitive we cannot have $3 \mid a$ and $3 \mid d$, so $(-d:a)$ is a root of F modulo 3, hence *the* root of F modulo 3. Therefore, $F \in U_3$ if and only if $F(-d,a) \not\equiv 0 \pmod{9}$. Since b and c are divisible by 3, the value of $F(x,y)$ modulo 9 depends only on x and y modulo 3, so

$$F(-d,a) \equiv ad(-d^2 + bd - ca + a^2) \pmod{9} \ ,$$

and the result follows by separately considering the three cases of (5). □

Corollary 8.3.4. *Let* $F = (a,b,c,d)$ *be a* primitive *form, and let* $H_F = (P,Q,R)$ *be its Hessian. Then* $F \notin U_2$ *if and only if* $\operatorname{disc}(F) \equiv 0 \pmod{16}$ *or* $\operatorname{disc}(F) \equiv 4 \pmod{16}$ *and* P *or* R *is odd.*

Proof. The proof is trivial and is left to the reader (Exercise 7). □

Proposition 8.3.5. *Let* F *be a primitive form, and write* $\operatorname{disc}(F) = d_k f^2$ *with* d_k *a fundamental discriminant. Then* $p \mid f$ *if and only if either* $(F,p) = (1^3)$ *or*

$$F(x,y) \equiv (\delta x - \gamma y)^2(\beta x - \alpha y) \pmod{p} \quad \textit{and} \quad F(\gamma,\delta) \equiv 0 \pmod{p^2} \ .$$

Proof. Assume first that $p \mid f$. Then $p \mid \operatorname{disc}(F)$ and so we can write

$$F(x,y) = (\delta x - \gamma y)^2(\beta x - \alpha y) + pF_1(x,y)$$

with F_1 integral.

Assume $p \neq 2$. A computation shows that

$$\operatorname{disc}(F) \equiv 4p(\alpha\delta - \beta\gamma)^3 F_1(\gamma,\delta) \pmod{p^2} \ .$$

Since $p \mid f$, we have $p^2 \mid \operatorname{disc}(F)$. Hence, if $p \neq 2$, either $p \mid (\alpha\delta - \beta\gamma)$ — in other words, $(F,p) = (1^3)$ — or $p \mid F_1(\gamma,\delta)$ — in other words, $F(\gamma,\delta) \equiv 0 \pmod{p^2}$.

Assume now that $p = 2$. If $F_1 = (a_1,b_1,c_1,d_1)$, a computation shows that

$$\operatorname{disc}(F) \equiv 8(\alpha\delta - \beta\gamma)^3 F_1(\gamma,\delta) + 4((a_1\alpha + b_1\beta)\gamma^2 + (c_1\alpha + d_1\beta)\delta^2)^2 \pmod{16} \ .$$

Since $2 \mid f$, $\operatorname{disc}(F) = d_k f^2 \equiv 0$ or $4 \pmod{16}$. Since the square of an integer is congruent to 0 or 1 modulo 4, it follows that

$$8(\alpha\delta - \beta\gamma)^3 F_1(\gamma,\delta) \equiv -4,\ 0,\ \text{or } 4 \pmod{16} \ ;$$

in other words, $(\alpha\delta - \beta\gamma)^3 F_1(\gamma,\delta) \equiv 0 \pmod{2}$. So once again, either $(F,2) = (1^3)$ or $2 \mid F_1(\gamma,\delta)$, as before.

Conversely, assume that either $(F,p) = (1^3)$ or $F(\gamma,\delta) \equiv 0 \pmod{p^2}$. If $(F,p) = (1^3)$, then since F is primitive, by Proposition 8.3.3 (2), we have $p \mid \gcd(P,Q,R)$.

Assume first that $p > 3$. Then $p^2 \mid \operatorname{disc}(H_F) = -3\operatorname{disc}(F)$, hence $p^2 \mid \operatorname{disc}(F)$, and so $p \mid f$.

Assume now that $p = 2$. Then $H_F = 2H'$ for some other quadratic form H', thus $\operatorname{disc}(H_F)/4 \equiv 0$ or $1 \pmod 4$, so the same is true for $\operatorname{disc}(F) = \operatorname{disc}(H_F)/(-3)$, and hence $2 \mid f$.

Finally, assume that $p = 3$. From the explicit formulas, $p \mid \gcd(P,Q,R)$ is equivalent to $3 \mid b$ and $3 \mid c$, from which it follows by the formula for the discriminant that $27 \mid \operatorname{disc}(F)$ and, in particular, that $3 \mid f$.

Assume now that $F(x,y) \equiv (\delta x - \gamma y)^2(\beta x - \alpha y) \pmod{p}$ and $F(\gamma,\delta) \equiv 0 \pmod{p^2}$. We may assume that $\alpha\delta - \beta\gamma \not\equiv 0 \pmod{p}$; otherwise, we are in the case $(F,p) = (1^3)$ that we just considered.

By the same reasoning as before, writing $F(x,y) = (\delta x - \gamma y)^2(\beta x - \alpha y) + pF_1(x,y)$, we have

$$\operatorname{disc}(F) \equiv 4p(\alpha\delta - \beta\gamma)^3 F_1(\gamma,\delta) \pmod{p^2} \ .$$

Therefore, $F(\gamma,\delta) \equiv 0 \pmod{p^2}$ is equivalent to $F_1(\gamma,\delta) \equiv 0 \pmod{p}$, which implies $\operatorname{disc}(F) \equiv 0 \pmod{p^2}$; hence, if $p > 2$, we have $p \mid f$.

For $p = 2$ and $F_1 = (a_1, b_1, c_1, d_1)$, we have as before

$$\operatorname{disc}(F) \equiv 8(\alpha\delta - \beta\gamma)^3 F_1(\gamma,\delta) + 4((a_1\alpha + b_1\beta)\gamma^2 + (c_1\alpha + d_1\beta)\delta^2)^2 \pmod{16} \ ,$$

and since $F_1(\gamma,\delta) \equiv 0 \pmod{2}$, we deduce that

$$\frac{\operatorname{disc}(F)}{4} \equiv ((a_1\alpha + b_1\beta)\gamma^2 + (c_1\alpha + d_1\beta)\delta^2)^2 \pmod{4} \ ,$$

and hence $2 \mid f$, thus finishing the proof of the proposition. □

Corollary 8.3.6. *Let F be a primitive cubic form and p be a prime. Then $F \notin U_p$ if and only if F has at least a double root $(\gamma : \delta)$ modulo p, and $F(\gamma,\delta) \equiv 0 \pmod{p^2}$.*

Proof. If $F \notin U_p$, then in particular $F \notin V_p$, hence $p \mid f$, and so by Proposition 8.3.5, either $F(\gamma,\delta) \equiv 0 \pmod{p^2}$ if F has a double root, or $(F,p) = (1^3)$, but then by definition of U_p, we again have $F(\gamma,\delta) \equiv 0 \pmod{p^2}$. Conversely, if F has at least a double root $(\gamma : \delta)$ modulo p, and $F(\gamma,\delta) \equiv 0 \pmod{p^2}$, then if it is a triple root, by definition $F \notin U_p$. If it is only a double root, by Proposition 8.3.5 we have $p \mid f$, so $F \notin V_p$, and hence $F \notin U_p$. When (γ,δ) is at least a double root modulo p, it is easily checked that the condition $F(\gamma,\delta) \equiv 0 \pmod{p^2}$ depends only on γ and δ modulo p. □

8.4 The Davenport–Heilbronn Theorem

We first need the following well-known results about cubic fields (note also the generalization given in Theorem 9.2.6) whose proofs are given in Section 10.1.5.

Proposition 8.4.1. *Let K be a cubic number field of discriminant $d(K)$, and write $d(K) = d_k f^2$, where d_k is a fundamental discriminant (including 1). Then*

(1) *$p \mid f$ if and only if p is totally ramified; in other words, if and only if $p\mathbb{Z}_K = \mathfrak{p}^3$,*

(2) $p \mid (d_k, f)$ *implies* $p = 3$,
(3) $p^2 \mid f$ *implies* $p = 3$.

We can now state and prove the Davenport–Heilbronn theorem.

Theorem 8.4.2. *We have* $I = \mathrm{Im}(\phi_{C\Phi}) = U$. *In other words, the maps* $\phi_{C\Phi}$ *and* $\phi_{\Phi C}$ *are discriminant-preserving inverse bijections between isomorphism classes of cubic fields and binary cubic forms belonging to* U.

Proof. Let K be a cubic number field, and let F_K be the image of K by $\phi_{C\Phi}$. We will first show that $F_K \in U$. As in Proposition 8.4.1, we write $\mathrm{disc}(F_K) = d(K) = d_k f^2$ with d_k a fundamental discriminant.

Let p be a prime. If $p \nmid f$, then $F_K \in V_p \subset U_p$, so $F_K \in U_p$. Hence we now assume that $p \mid f$.

By Proposition 8.4.1 (1), it follows that p is totally ramified. By Proposition 8.2.3, this means that F_K splits as the cube of a linear form, so that $(F_K, p) = (1^3)$. We consider three cases:

- $p > 3$. In this case, it follows from Proposition 8.4.1 (2) and (3) that $p^3 \nmid d(K)$ and hence that $F_K \in U_p$ by Proposition 8.3.3 (3).
- $p = 2$. Since $2 \mid f$, we have $2 \nmid d_k$ by Proposition 8.4.1 (2). Therefore, by Proposition 8.4.1 (3) we deduce that $2^3 \nmid d(K)$, hence $F_K \in U_2$, as before.
- $p = 3$. This is the only difficult case. Since 3 is totally ramified, write $3\mathbb{Z}_K = \mathfrak{p}^3$ for a prime ideal $\mathfrak{p}$ of degree 1, and let $\gamma \in \mathfrak{p} \setminus \mathfrak{p}^2$. According to Lemma 10.1.2, $\mathbb{Z}[\gamma]$ is a 3-maximal order in $\mathbb{Z}_K$. In particular, $v_3(d(K)) = v_3(\mathrm{disc}(\gamma))$. On the other hand, we clearly have

$$\mathrm{disc}(\gamma^2) = \mathrm{disc}(\gamma)((\gamma+\gamma')(\gamma+\gamma'')(\gamma'+\gamma''))^2 = \mathrm{disc}(\gamma)\,\mathcal{N}^2(\mathrm{Tr}(\gamma)-\gamma)\ ,$$

where γ, γ', and γ'' are the conjugates of γ. Since $\mathfrak{p}$ is totally ramified, $\mathrm{Tr}(\gamma) \in \mathfrak{p} \cap \mathbb{Z} = 3\mathbb{Z}$, and since $\mathcal{N}(\mathfrak{p}) = 3$,

$$\mathcal{N}(\mathrm{Tr}(\gamma)-\gamma) \equiv -\mathcal{N}(\gamma) \equiv \pm 3 \pmod 9\ ;$$

hence it follows that

$$f(\gamma^2) = \left(\frac{\mathrm{disc}(\gamma^2)}{d(K)}\right)^{1/2} \equiv \pm 3 \pmod 9\ .$$

Let $(1, \alpha, \beta)$ be the integral basis used to define the form F_K. Thus $\gamma^2 = \beta u - \alpha v + w$ for some u, v, and w in $\mathbb{Z}$, and by definition we have $F_K(u,v) \equiv \pm 3 \pmod 9$. We have the following lemma.

Lemma 8.4.3. *Assume that* F *is a primitive cubic form and* p *a prime such that* $(F,p) = (1^3)$. *Then* $F \in U_p$ *if and only if there exists* $(u,v) \in \mathbb{Z}^2$ *such that* $F(u,v) = ep$ *with* $p \nmid e$.

Assuming this lemma for a moment, we see that since $F_K(x,y)$ represents an integer congruent to ± 3 modulo 9, we have $F_K \in U_3$, thus finishing the proof that $F \in U$.

Let us prove the lemma. By lifting the condition $(F,p) = (1^3)$ to $\mathbb{Z}$, we can write

$$F(x,y) = \lambda(\delta x - \gamma y)^3 + pG(x,y) \ ,$$

with λ, γ, δ in $\mathbb{Z}$ and G an integral cubic form. Then $F \in U_p$ if and only if $p \nmid G(\gamma,\delta)$. Thus if $F \in U_p$, we have $F(\gamma,\delta) = ep$ with $e = G(\gamma,\delta) \not\equiv 0 \pmod{p}$.

Conversely, assume that there exists $(u,v) \in \mathbb{Z}^2$ such that $F(u,v) = ep$ with $p \nmid e$. Then $F(u,v) \equiv 0 \pmod{p}$, hence $\lambda(\delta u - \gamma v) \equiv 0 \pmod{p}$. Since F is primitive, $p \nmid \lambda$ and hence $\delta u - \gamma v \equiv 0 \pmod{p}$. Again, since F is primitive, γ and δ cannot both be divisible by p, from which it follows that there exists μ such that $\gamma \equiv \mu u \pmod{p}$ and $\delta \equiv \mu v \pmod{p}$. Also, since $p \nmid e$, we have $p \nmid \mu$. But then $ep = F(u,v) \equiv pG(u,v) \pmod{p^3}$, hence $G(u,v) \equiv e \pmod{p^2}$, and so $G(\gamma,\delta) \equiv \mu^3 G(u,v) \equiv \mu^3 e \not\equiv 0 \pmod{p}$, so $F \in U_p$, proving the lemma. □

To finish the proof of Davenport–Heilbronn's Theorem 8.4.2, we must now prove that if $F \in U$, there exists a cubic field K such that F is equivalent to F_K. For this, we introduce a definition.

Definition 8.4.4. *We will say that two cubic forms F_1 and F_2 are* rationally equivalent *if there exists $M \in \mathrm{GL}_2(\mathbb{Q})$ such that $F_1 \circ M = \mu F_2$ for some $\mu \in \mathbb{Q}^*$.*

We first show the following lemma.

Lemma 8.4.5. *Let F be any form in Φ (in other words, primitive and irreducible). Then there exists a number field K such that F is rationally equivalent to F_K.*

Proof. In some algebraic closure of $\mathbb{Q}$, write $F = a(x - \lambda y)(x - \lambda' y)(x - \lambda'' y)$. Since F is irreducible, λ is a cubic irrationality, and we will take $K = \mathbb{Q}(\lambda)$. Write $F_K = a_K(x - \nu y)(x - \nu' y)(x - \nu'' y)$ so that $\nu \in K$ (we saw above that this is always possible). Since K is a $\mathbb{Q}$-vector space of dimension 3, there exist four integers k, l, m, and n not all zero such that $l + k\lambda - n\nu - m\lambda\nu = 0$. Taking conjugates, we obtain the same equality with $'$ and $''$. Using $\lambda = (\nu n - l)/(k - \nu m)$, we obtain

$$F_K(kx + ly, mx + ny) = \rho(x - \lambda y)(x - \lambda' y)(x - \lambda'' y) = (\rho/a)F$$

with $\rho = a_K \mathcal{N}(k - \nu m) \in \mathbb{Q}$. Furthermore, the determinant $kn - lm$ is nonzero since otherwise either λ or ν would be in $\mathbb{Q}$. Thus, F_K is rationally equivalent to F. □

Finally, we have the following lemma.

Lemma 8.4.6. *Let F_1 and F_2 be two forms belonging to U. If the forms F_1 and F_2 are rationally equivalent, they are equivalent.*

Since F_K belongs to U, it follows from Lemmas 8.4.5 and 8.4.6 that any form in U is equivalent to F_K for a certain K, and this finishes the proof of Davenport–Heilbronn's Theorem 8.4.2.

Proof. Assume that $F_1 \circ M = \mu F_2$ for some $M \in \mathrm{GL}_2(\mathbb{Q})$ and $\mu \in \mathbb{Q}^*$. Since we want to show that F_1 and F_2 are equivalent, we can replace them by equivalent forms. In other words, without changing the equivalence classes of F_1 and F_2, we may replace M by any matrix of the form UMV with U and V in $\mathrm{GL}_2(\mathbb{Z})$. The elementary divisor theorem (or the existence of the Smith normal form) tells us that we can choose U and V so that

$$UMV = \begin{pmatrix} \alpha m & 0 \\ 0 & \alpha \end{pmatrix} \quad (\alpha \in \mathbb{Q}^*,\ m \in \mathbb{Z}_{>0}) \ .$$

Replacing μ by $\alpha^3\mu$, we may assume that $\alpha = 1$. To summarize, by replacing F_1 and F_2 by equivalent forms and modifying μ and M, we may assume that $M = \left(\begin{smallmatrix} m & 0 \\ 0 & 1 \end{smallmatrix}\right)$ with $m \in \mathbb{Z}_{>0}$.

If $m = 1$, then $F_1 = \mu F_2$, and since F_1 and F_2 are in U they are primitive. Thus, $\mu = \pm 1$, so F_1 and F_2 are equivalent.

Otherwise, there exists a prime p such that $p \mid m$. Write $m = p^k m_0$ and $\mu = p^l \mu_0$ with m_0 and μ_0 having zero p-adic valuation. Since $p \mid m$, we have $k \geq 1$. Writing $F_i = (a_i, b_i, c_i, d_i)$ for $i = 1$ and 2, we see that the equality $F_1(p^k m_0 x, y) = p^l \mu_0 F_2(x, y)$ is equivalent to

$$\begin{cases} a_1 = \tau_a p^{l-3k} a_2 \ , \\ b_1 = \tau_b p^{l-2k} b_2 \ , \\ c_1 = \tau_c p^{l-k} c_2 \ , \\ d_1 = \tau_d p^l d_2 \ , \end{cases}$$

where the τ_x are rational numbers having zero p-adic valuation.

Assume that $l - k > 0$, hence that $l > k \geq 1$. Then $p \mid c_1$ and $p^2 \mid d_1$. If, on the other hand, $l - k \leq 0$, then $l \leq k$ so that $p \mid b_2$ and $p^2 \mid a_2$. Replacing $(F_1(x, y), F_2(x, y))$ by $(F_2(y, x), F_1(y, x))$, we again obtain $p \mid c_1$ and $p^2 \mid d_1$.

The formula for the discriminant implies $p^2 \mid \mathrm{disc}(F_1)$. If $p > 2$, this immediately implies $F_1 \notin V_p$. If $p = 2$, the formula for the discriminant shows that $\mathrm{disc}(F_1)/4 \equiv b_1^2 (c_1/2)^2 \pmod 4$ and this is congruent to 0 or 1 modulo 4, so $F_1 \notin V_2$.

Since $F_1 \in U_p$, we must have $(F_1, p) = (1^3)$, and $F_1(x, y) \equiv \lambda(\delta x - \gamma y)^3 \pmod p$ with $F_1(\gamma, \delta) \not\equiv 0 \pmod{p^2}$. Since $p \nmid \lambda$, $p \mid d_1$ means that $p \mid \gamma$; hence $p \mid b_1 \equiv -3\lambda\delta^2\gamma \pmod p$. But then, since $p \mid \gamma$, $F_1(\gamma, \delta) \equiv 0 \pmod{p^2}$, which is a contradiction. This finishes the proof of Lemma 8.4.6 and hence of Davenport–Heilbronn's Theorem 8.4.2. □

As a consequence of this theorem and of the results of the previous section, we now have an efficient algorithm to test whether a cubic form corresponds to the image of a cubic field by the Davenport–Heilbronn map as follows.

Algorithm 8.4.7 (Cubic Form Test, Version 1). Let $F = (a, b, c, d)$ be a cubic form. This algorithm outputs `true` or `false` according to whether or not F corresponds to the image of a cubic field by the Davenport–Heilbronn map.

1. [Check irreducible] If F is not irreducible, return `false`.
2. [Check primitive] If F is not primitive, return `false`.
3. [Compute Hessian] Let $P \leftarrow b^2 - 3ac$, $Q \leftarrow bc - 9ad$, and $R \leftarrow c^2 - 3bd$. Let $H_F \leftarrow (P, Q, R)$ be the Hessian of F, and set $f_H \leftarrow \gcd(P, Q, R)$ and $D \leftarrow 4PR - Q^2$ (so $D = 3\,\mathrm{disc}(F)$).
4. [Check conditions at 2 and 3] Using Proposition 8.3.3 or Corollary 8.3.6, check that $F \in U_2$ and $F \in U_3$. If this is not the case, return `false`.
5. [Check f_H almost squarefree] If $p^2 \mid f_H$ for some $p > 3$, return `false`.
6. [Check f_H and D/f_H^2 almost coprime] Set $t \leftarrow D/f_H^2$. Remove all powers of 2 and 3 from t (in fact, at most 2^3 and 3^2), and again let t be the result. If $\gcd(t, f_H) > 1$, return `false`.
7. [Check D/f_H^2 almost squarefree] If t is squarefree, return `true`; otherwise, return `false`.

In the comments to this algorithm, "almost" means outside the primes 2 and 3.

Proof. Steps 1, 2, 3, and 4 are clear since F_K is irreducible, is primitive, and belongs to U_p for all p.

If $p^2 \mid f_H$, then $p^4 \mid D = 3\,\mathrm{disc}(F)$. If $p \mid f_H$ and $p \mid D/f_H^2$, we have $p^3 \mid D$. It follows in both cases that if $p > 3$, we have $p^3 \mid \mathrm{disc}(F)$, so $F \notin U_p$ by Proposition 8.3.3, thus proving steps 5 and 6.

Assume that for every prime $p > 3$, we have $p^2 \nmid f_H$ and $p \nmid \gcd(t, f_H)$ (which is the situation at the beginning of step 7). Then t is not squarefree if and only if there exists $p > 3$ such that $p^2 \mid t$ and hence $p \nmid f_H$. By Proposition 8.3.3, we cannot have $(F, p) = (1^3)$. On the other hand, since $p > 3$, we have $p^2 \mid \mathrm{disc}(F)$, so $F \notin V_p$, and hence $F \notin U_p$. Thus t is squarefree if and only if for all $p > 3$ we have $F \in U_p$, proving the algorithm's validity. □

Remarks

(1) Step 1 will in practice not be necessary since we will always use this algorithm with *reduced forms* (see later), which are irreducible.

(2) Although step 2 seems to be necessary, this is in fact not the case since nonprimitive forms will be excluded in the subsequent steps. Indeed, let p be a prime dividing all the coefficients of F. Then clearly $p^2 \mid f_H$. If $p > 3$, step 4 will return `false`. Assume now that $p = 2$ or $p = 3$. Then $F \notin V_p$ and F has at least a double root modulo p. If F has only a double

and not a triple root, then $F \notin U_p$. However, if F has a triple root, so that $(F,p) = (1^3)$, then $F(x,y) = \lambda(\delta x - \gamma y)^3$ with $\lambda \in \mathbb{F}_p^*$ and hence $p \mid \gcd(\delta^3, \gamma^3)$. Thus, γ and δ are equal to zero modulo p, and hence $F \notin U_p$.

(3) Even though it may seem useful to include the unnecessary step 2, it can be shown that on average it *slows down* the algorithm, so thanks to the preceding remark, in the final form we will suppress it.

(4) It may seem surprising that the slowest part of the algorithm is by far the squarefreeness test in steps 5 and 7. We will see later how to speed this up in practice.

We end this section by giving the following proposition.

Proposition 8.4.8. *Let K be a cubic number field, and as before write $d(K) = d_k f^2$, where d_k is a fundamental discriminant. Let F_K be the cubic form associated to K by the Davenport–Heilbronn map, and let $H_K = (P,Q,R)$ be its Hessian. Finally, set $f_H = \gcd(P,Q,R)$. Then*

(1) *$f_H = f$ or $f_H = 3f$ and apart from powers of 3, f and f_H are squarefree,*
(2) *if $f_H = 3f$, then $3 \mid f$,*
(3) *if $3 \| f$, then $f_H = 3f$,*
(4) *in particular, $f \mid f_H$, f and f_H have the same prime divisors, and we can have $(v_3(f), v_3(f_H))$ only equal to $(0,0)$, $(1,2)$, $(2,2)$, and $(2,3)$.*

Proof. By Proposition 8.3.3 (2), we have $p \mid f_H$ if and only if $(F,p) = (1^3)$, and by Proposition 8.2.3 (3) this is true if and only if p is totally ramified, hence by Proposition 8.4.1 (1), if and only if $p \mid f$. Hence f and f_H have the same prime divisors. Let p be such a prime divisor. By Davenport–Heilbronn's theorem, $F_K \in U$. Hence by Proposition 8.3.3 (3) we have $p^3 \nmid d(K)$ if $p \neq 3$, so if $p \neq 3$, we have $p^2 \nmid f$ and $p \nmid d_k$, and so up to powers of 3, f is squarefree. Since $f_H^2 \mid 3d(K) = d_k f^2$, we have $v_p(f_H^2) \leq 2$, and thus up to powers of 3, f_H is also squarefree.

Assume now that $p = 3$ and that p divides f (hence also f_H). By Proposition 8.3.3 (3) we have $v_3(f) \leq 2$ and since $f_H^2 \mid 3d(K)$, we have $v_3(f_H) \leq 3$. Furthermore, if $v_3(f) = 1$, we have $v_3(f_H^2) \leq 3^4$, and so $v_3(f_H) \leq 2$. Finally, since $3 \mid (P,Q,R)$, using the explicit formulas in terms of the coefficients of the form F_K we see that $3 \mid b$ and $3 \mid c$, which implies $9 \mid (P,Q,R)$, hence $v_3(f_H) \geq 2$. This proves all the assertions of the proposition. □

8.5 Real Cubic Fields

We would now like to single out a unique representative of a cubic form $F \in U$, which we will call "reduced". For this purpose, as in the quadratic case, we must distinguish according to the signature of the corresponding cubic field. In this section, we assume that the field $K = K_F$ is totally real or, equivalently, that $d(K) = \mathrm{disc}(F) > 0$.

The Hessian H_F satisfies $\mathrm{disc}(H_F) = -3\,\mathrm{disc}(F) < 0$, hence is a (positive or negative) definite quadratic form for which the notion of reduction is well-defined. We will essentially define F to be reduced when H_F is, but for this we must make a few technical modifications to the usual definitions.

Definition 8.5.1. *Let $H = (P,Q,R)$ be a quadratic form with real coefficients. We will say that H is* reduced *if $|Q| \le P \le R$ and $R > 0$ (to exclude the trivial case of the zero form).*

Note that this is not quite the same as the usual definition, which would be $|Q| \le P \le R$ with $Q \ge 0$ when one of the inequalities is an equality. The reason for the modification to the usual definition is that we must work with forms modulo $\mathrm{GL}_2(\mathbb{Z})$ and not only $\mathrm{SL}_2(\mathbb{Z})$.

As usual, if $H = (P,Q,R)$, we will set $H^{-1} = (P,-Q,R)$, and we will denote by $\mathrm{Aut}(H)$ the set of elements $M \in \mathrm{GL}_2(\mathbb{Z})$ stabilizing H. Finally we set $\sigma = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$.

The following lemma shows that the above definition works.

Lemma 8.5.2. *Let $H = (P,Q,R)$ and $H' = (P',Q',R')$ be two reduced, definite, integral, binary quadratic forms such that there exists $M \in \mathrm{GL}_2(\mathbb{Z})$ with $H' = H \circ M$. Then, either $H' = H$ and $M \in \mathrm{Aut}(H)$, or $H' = H^{-1}$ and $M \in \mathrm{Aut}(H)\sigma$. Furthermore, the only elements of $\mathrm{Aut}(H)$ are $\pm I_2$, except in the following special cases that can occur simultaneously:*

$$\begin{aligned}
&\text{if } P = R, && \text{add } \pm\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix};\\
&\text{if } Q = 0, && \text{add } \pm\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix};\\
&\text{if } P = R \text{ and } Q = 0, && \text{add } \pm\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix};\\
&\text{if } P = \varepsilon Q, && \text{add } \pm\begin{pmatrix} 1 & \varepsilon \\ 0 & -1 \end{pmatrix};\\
&\text{if } P = \varepsilon Q = R, && \text{add } \pm\begin{pmatrix} -1 & 0 \\ \varepsilon & 1 \end{pmatrix}, \pm\begin{pmatrix} 0 & -1 \\ 1 & \varepsilon \end{pmatrix}, \pm\begin{pmatrix} \varepsilon & 1 \\ -1 & 0 \end{pmatrix},
\end{aligned}$$

where in the last two cases $\varepsilon = \pm 1$.

Proof. Since H and H' are equivalent, they have the same discriminant and represent the same integers. To say that H is reduced implies that P is the minimum of H on $\mathbb{Z}^2 - \{(0,0)\}$ and that R is the second minimum; hence $P = P'$ and $R = R'$. Equality of discriminants implies $Q = \pm Q'$. Hence $H' = H$ or $H' = H^{-1} = H \circ \sigma$, so we need only to compute $\mathrm{Aut}(H)$.

Let $T = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ and $S = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ be the usual generators of the modular group $\mathrm{PSL}_2(\mathbb{Z})$, and let $\mathcal{F}$ be the usual compact fundamental domain for the modular group in the upper half-plane $\mathcal{H}$.

Let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ be an automorphism of $H = (P, Q, R)$.

If $M \in \mathrm{SL}_2(\mathbb{Z})$, then it fixes $\tau \in \mathcal{F}$. Thus, by the usual theory, it is either $\pm I$, $\pm S$ (if $H = (P, 0, P)$), $\pm ST$ or $\pm(ST)^2$ (if $H = (P, P, P)$), $\pm TS$ or $\pm(TS)^2$ (if $H = (P, -P, P)$).

If $M \notin \mathrm{SL}_2(\mathbb{Z})$, then M swaps the two complex roots τ and $\bar{\tau}$ of $H(x, 1)$, so that

$$a\tau + b = c\tau\bar{\tau} + d\bar{\tau} \ .$$

Taking imaginary parts, we get $a = -d$, and taking real parts and using $\tau + \bar{\tau} = -Q/P$ and $\tau\bar{\tau} = R/P$, we obtain $bP = aQ + cR$. Finally, the determinant condition gives $a^2 + bc = 1$. One easily checks that the three conditions $a = -d$, $a^2 + bc = 1$, and $bP = aQ + cR$ are necessary and sufficient conditions for M to be an automorphism of H.

Thus, $H(a, c) = Pa^2 + Qac + Rc^2 = Pa^2 + Pbc = P$, and $H(a, c) \geq (P - Q + R)\min(a^2, c^2)$, and since H is reduced we have $|Q| \leq P \leq R$. We thus obtain the following.

- If $ac \neq 0$, then $a^2 = c^2 = 1$ and $P = |Q| = R$, so $b = 0$, $d = -a = \pm 1$. If $P = \varepsilon Q$ with $\varepsilon = \pm 1$, we have $a = -\varepsilon c$.
- If $a = 0$, then $bc = 1$ and $Rc^2 = P$, so $R = P$ and $c = \pm 1$, and hence $b = c$ and $d = 0$.
- If $c = 0$, then $a^2 = 1$ and $bP = aQ$. This implies that either $b = 0$ and $Q = 0$, or $b = \varepsilon a$ and $P = \varepsilon Q$ with $\varepsilon = \pm 1$.

This finishes the proof of the lemma. Note that it follows from this result that the group G of automorphisms of H is always isomorphic to a group of the form $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$, with $m = 1$, $m = 2$ (when $Q = 0$ or $P = R$ or $P = \varepsilon Q$ are the only equalities), $m = 4$ (when $P = R$ and $Q = 0$), or $m = 6$ (when $P = \varepsilon Q = R$). □

We can now give the definition of a reduced cubic form in the case of a positive discriminant.

Definition 8.5.3. *Let $F = (a, b, c, d)$ be an integral binary cubic form of* positive *discriminant. We will say that F is* reduced *if its Hessian H_F is reduced in the above sense and if, in addition:*

(1) $a > 0$, $b \geq 0$, *and* $d < 0$ *if* $b = 0$;
(2) *if* $P = Q$, *then* $b < |3a - b|$;
(3) *if* $P = R$, *then* $a \leq |d|$ *and* $b < |c|$ *if* $|d| = a$.

Note that there is no extra condition for $Q = 0$ or for $P = -Q$.

With this definition, we have the following.

Proposition 8.5.4. (1) *Two equivalent, reduced, real cubic forms are equal.*
(2) *A reduced real cubic form belonging to U is irreducible.*
(3) *Any irreducible real cubic form is equivalent to a unique reduced form.*

Proof. (1). Let F and F' be two reduced cubic forms and $M \in GL_2(\mathbb{Z})$ such that $F' = F \circ M$. Then $H_{F'} = H_F \circ M$ and since H_F and $H_{F'}$ are reduced, it follows from Lemma 8.5.2 that $H_{F'} = H_F$ or $H_{F'} = H_F^{-1}$, so M belongs either to $\mathrm{Aut}(H)$ or to $\mathrm{Aut}(H)\sigma$. To simplify notation, we will write

$$N = \pm \begin{pmatrix} a & b\eta \\ c & d\eta \end{pmatrix} \in \mathrm{Aut}(H)$$

and

$$M = \pm \begin{pmatrix} a & b\eta \\ c & d\eta \end{pmatrix} ,$$

where $\eta = 1$ if $M = N$, or $\eta = -1$ if $M = N\sigma$.

We also write $H_F = (P, Q, R)$ and $H_{F'} = (P', Q', R')$, so that $P' = P$, $R' = R$, and $Q' = \eta Q$.

We have only a finite number of possibilities (exactly 16) to test for N. Let $F = (a, b, c, d)$.

First, if $N = \pm I$, hence $M = \pm \left(\begin{smallmatrix} 1 & 0 \\ 0 & \eta \end{smallmatrix}\right)$, we have $F' = \pm(a, b\eta, c, d\eta)$ and therefore the conditions $a > 0$ on F and F' imply the $+$ sign, and the conditions $b \geq 0$ and $d < 0$ if $b = 0$ on F and F' imply $\eta = 1$, so $M = I$ and $F = F'$.

If $P = R$ and $M = \pm \left(\begin{smallmatrix} 0 & \eta \\ 1 & 0 \end{smallmatrix}\right)$, we have $F' = \pm(d, c\eta, b, a\eta)$; hence the conditions $a \leq |d|$ on F and F' imply $a = |d|$, so the additional conditions $b < |c|$ on F and F' give a contradiction.

In the two special cases corresponding to $Q = 0$, we have $M = \pm \left(\begin{smallmatrix} 1 & 0 \\ 0 & -\eta \end{smallmatrix}\right)$ or $M = \mp \left(\begin{smallmatrix} 0 & -\eta \\ 1 & 0 \end{smallmatrix}\right)$, which have been considered above.

If $P = \varepsilon Q$ and $M = \pm \left(\begin{smallmatrix} 1 & \varepsilon\eta \\ 0 & -\eta \end{smallmatrix}\right)$, we have $F' = \pm(a, \eta(3a\varepsilon - b), 3a - 2b\varepsilon + c, \eta(a\varepsilon - b + c\varepsilon - d))$; hence the conditions $a > 0$ on F and F' imply the $+$ sign. Assume first that $\varepsilon = 1$. We have $3a - b > 0$ since otherwise, the condition $b < |3a - b|$ on F implies $a < 0$, which is a contradiction. Condition $b \geq 0$ on F' implies that $\eta = \mathrm{sign}(3a - b) = 1$, hence $P' = Q'$, so the condition $b < |3a - b|$ on F' implies $0 \leq 3a - b < b$, again a contradiction.

Assume now that $\varepsilon = -1$. Then condition $b \geq 0$ on F' implies that $\eta = -1$, hence $P' = Q'$, so the condition $b < |3a - b|$ on F' implies that $3a + b < |-b| = b$, another contradiction.

Finally, assume that $P = \varepsilon Q = R$ and that N is one of the six matrices given by Lemma 8.5.2. Then an easy computation shows that F is of the form $F = (a, b, \varepsilon b - 3a, -\varepsilon a)$. If $\varepsilon = 1$, the reducedness of F is equivalent to $a > 0$, $b \geq 0$, $b < |3a - b|$; while if $\varepsilon = -1$, the reducedness of F is equivalent to $a > 0$ and $b > 0$.

For the six matrices N of Lemma 8.5.2, we have $F' = \pm(a, \eta(3\varepsilon a - b), -\varepsilon b, -\varepsilon\eta a)$, $F' = \pm(-\varepsilon a, -\varepsilon\eta b, 3\varepsilon a - b, a\eta)$, and $F' = \pm(-\varepsilon a, -\varepsilon\eta b, 3\varepsilon a - b, a\eta)$, respectively.

For the first two, we are in a special case of the case $P = \varepsilon Q$ considered above, so we obtain a contradiction.

The last four, together with the conditions $a > 0$ on F and F', imply that $\pm 1 = -\varepsilon$, hence that $F' = (a, \eta b, \varepsilon b - 3a, -\varepsilon\eta a)$. If $\eta = -1$, then the conditions $b \geq 0$ on F and F' imply that $b = 0$, in which case the conditions $d < 0$ on F and F' give a contradiction. Thus $\eta = 1$, so $F' = F$, and the matrices $\left(\begin{smallmatrix} 0 & \varepsilon \\ -\varepsilon & -1 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} -1 & -\varepsilon \\ \varepsilon & 0 \end{smallmatrix}\right)$ together with the identity form the automorphism group of F. Thus in this case, and only in this case, the automorphism group of F is nontrivial and is cyclic of order 3 generated by one of the above two matrices. This finishes the proof of (1).

(2). Let $F = (a, b, c, d) \in U$ be a reduced real cubic form. We will successively replace F by equivalent forms, which we will still denote (a, b, c, d) by abuse of notation, until we can conclude. Note that all forms equivalent to F still belong to U. Set $\Delta = \mathrm{disc}(F)$. Assume by contradiction that F is reducible. Then by transforming F by a suitable element of $\mathrm{GL}_2(\mathbb{Z})$, we may assume that $a = 0$ and hence $b > 0$ since $\Delta \neq 0$. Changing (x, y) into $(x - ky, y)$ for a suitable k, we may assume that $|c| \leq b$, and finally by changing (x, y) into $(-x, y)$ if necessary, we may assume that $0 \leq c \leq b$. We thus have $\Delta = b^2c^2 - 4b^3d$ and $(P, Q, R) = (b^2, bc, c^2 - 3bd)$. Let us show that $b = 1$. First, let p be an odd prime dividing b. Since $b^2 \mid \Delta$, we have $F \notin V_p$, but since $F \in U_p$, by definition this means that $(F, p) = (1^3)$ plus an additional condition. In particular, by Proposition 8.3.3 (2) we have $p \mid \gcd(P, Q, R)$; hence $p \mid (c^2 - 3bd)$, so $p \mid c$, and hence $p^3 \mid \Delta$, which leads to a contradiction unless $p = 3$ by Proposition 8.3.3 (3). But if $p = 3$, we have $3 \mid a$ and $9 \mid a$ since $a = 0$, so Proposition 8.3.3 (5) implies that $F \notin U_p$, again a contradiction.

Assume now that $2 \mid b$. We then have $\Delta \equiv b^2c^2 \pmod{16}$, hence $(\Delta/4) \equiv ((b/2)c)^2 \equiv 0$ or $1 \pmod 4$, and so $F \notin V_2$. Since $F \in U_2$, we have $(F, 2) = (1^3)$. We conclude as before by Proposition 8.3.3 (2) that $2 \mid (c^2 - bd)$; hence $2 \mid c$, so $16 \mid \Delta$, in contradiction with $8 \nmid \Delta$ which comes from Proposition 8.3.3 (3).

We have thus shown that $b = 1$, and hence $c = 0$ or $c = 1$. Thus, if we call G the final cubic form that we have obtained, we have $F = G \circ M$ for a certain $M \in GL_2(\mathbb{Z})$, and $G = (0, 1, c, d)$ with $c = 0$ or 1, hence $H_G = (1, 0, -3d)$ or $H_G = (1, 1, 1 - 3d)$. Since $\Delta = c^2 - 4d > 0$, we must have $d < 0$. Thus, these two quadratic forms are reduced and are thus equal to H_F or to $H_F \circ \sigma$. Since H_F is a covariant, we have $H_F = H_G \circ M$. It follows that M or $M\sigma$ belongs to $\mathrm{Aut}(H_G)$. However, an examination of the special cases of Lemma 8.5.2 shows that the elements of $\mathrm{Aut}(H_G)$ and $\mathrm{Aut}(H_G)\sigma$ fix a up to sign. Since the final a that we have obtained is equal to 0, we deduce that the initial a is also equal to 0, which is forbidden for a reduced cubic form.

(3). The uniqueness statement follows from (1). Let $F = (a, b, c, d)$ be an irreducible real cubic form and H_F its Hessian. By the usual theory of reduction of quadratic forms, we can find $M \in \mathrm{GL}_2(\mathbb{Z})$ such that $H_F \circ M$ is reduced; hence by changing F into $F \circ M$, we may assume that $H_F = (P, Q, R)$ is reduced.

For the special cases of Definition 8.5.3, we must now check that the use of the 16 matrices of Lemma 8.5.2 will lead to a reduced form, assuming that F is irreducible. For the sake of completeness, we give the details.

First note that $a \neq 0$ and $d \neq 0$, since otherwise F would be reducible. Hence, using $-I_2$, we may assume that $a > 0$; using σ, we may assume that $b \geq 0$. Furthermore, if $b = 0$, using σ we may assume that $d < 0$. Using these matrices, we have either fixed H_F or changed it into $H_F \circ \sigma = H_F^{-1}$, which is also reduced.

Assume $P = Q$ and $b \geq |3a - b|$. In this case, by Lemma 8.5.2 the matrix $\left(\begin{smallmatrix} 1 & 1 \\ 0 & -1 \end{smallmatrix}\right)$ belongs to $\mathrm{Aut}(H_F)$ and transforms F into $F' = (a, 3a - b, 3a - 2b + c, a - b + c - d)$.

Assume first that $b < 3a$. Since $b \geq |3a - b| = 3a - b$, we have $b \geq 3a/2$, so $0 < 3a - b \leq |3a - (3a - b)| = b$. An easy computation shows that if $b = 3a/2$, then $F = (2c - 4d, 3c - 6d, c, d)$ is divisible by $2x + y$, and hence is reducible. Thus $b > 3a/2$, so $3a - b < |3a - (3a - b)|$, as desired.

Assume now that $b = 3a$. Thus $F' = (a, 0, c', d')$ for certain values of c' and d'. Using σ if necessary (which will change Q into $-Q$, for which there is no extra reducedness condition), we may assume that $d' < 0$; hence F' is reduced.

Finally, if $b > 3a$, we have $F'' = F' \circ \sigma = (a, b - 3a, c'', d'')$ and $P'' = -Q''$, so F'' is reduced.

Consider now the case $P = R$ and $a > |d|$ or $a = |d|$ and $b \geq c$. Let $s = \mathrm{sign}(d)$ (recall that $d \neq 0$ since F is irreducible) and $t = \mathrm{sign}(c)$ if $c \neq 0$, otherwise $t = -\mathrm{sign}(a)$. By Lemma 8.5.2, the matrix $\left(\begin{smallmatrix} 0 & t \\ s & 0 \end{smallmatrix}\right)$ belongs to $\mathrm{Aut}(H_F)$ or to $\mathrm{Aut}(H_F)\sigma$ and transforms F into $F' = (a', b', c', d') = (|d|, |c|, sb, ta)$, and $a' > 0$, $b' \geq 0$, and $d' < 0$ if $b' = 0$. In addition, we have $a' < |d'|$ or $a' = |d'|$ and $b' \leq |c'|$. If $b = |c|$, however, we cannot have $c = 0$; otherwise, $b = 0$, hence $P = R = 0$, so $\mathrm{disc}(F) = 0$, and an easy computation shows that (still with $t = \mathrm{sign}(c)$) $F = (a, b, tb, ta)$ is divisible by $x + ty$ and thus is reducible. Thus, $b > |c|$, hence $b' < |c'|$, as desired, finishing the proof of the proposition. □

To be able to produce all reduced binary cubic forms of discriminant bounded by X, we must be able to give bounds on the coefficients of a reduced form. Such a result is as follows.

Proposition 8.5.5. *Let $F = (a, b, c, d)$ be a reduced form such that $0 < \mathrm{disc}(F) \leq X$. We have the following inequalities.*

(1)

$$1 \leq a \leq \frac{2}{3\sqrt{3}} X^{1/4} .$$

(2) *If* $a \le X^{1/4}/3$, *we have*

$$0 \le b \le \frac{3a}{2} + \sqrt{\sqrt{X} - \frac{27a^2}{4}} \ ,$$

while if $X^{1/4}/3 < a \le 2X^{1/4}/(3\sqrt{3})$, *we have*

$$\frac{3a}{2} - \sqrt{\sqrt{X} - \frac{27a^2}{4}} \le b \le \frac{3a}{2} + \sqrt{\sqrt{X} - \frac{27a^2}{4}} \ .$$

(3) *If* $a > X^{1/4}/3$ *or* $a \le X^{1/4}/3$ *and* $b \ge -3a/2 + \sqrt{\sqrt{X} - 27a^2/4}$, *we have*

$$\frac{b^2 - \sqrt{X}}{3a} \le c \le b - 3a \ ,$$

while if $a \le X^{1/4}/3$ *and* $b \le -3a/2 + \sqrt{\sqrt{X} - 27a^2/4}$, *we have*

$$\frac{b^2 - P_2}{3a} \le c \le b - 3a \ ,$$

where P_2 *is the unique positive solution of the equation*

$$4P_2^3 - (3a + 2b)^2 P_2^2 - 27a^2 X = 0 \ .$$

Proof. Let $H = (P, Q, R)$ be the Hessian of F, and set $\Delta = \mathrm{disc}(F)$ so that $4PR - Q^2 = 3\Delta$. Since F is reduced, we have

$$3X \ge 3\Delta \ge 4PR - P^2 \ge 3P^2 \ .$$

However, it is easily checked that

$$Pb^2 - 3Qab + 9Ra^2 - P^2 = 0 \ .$$

This is a quadratic equation in b, which therefore must have a nonnegative discriminant. As its discriminant is equal to

$$9a^2(Q^2 - 4PR) + 4P^3 = 4P^3 - 27a^2\Delta \ ,$$

we thus have

$$a^2 \le \frac{4P^3}{27\Delta} \le \frac{4P}{27} \le \frac{4\sqrt{X}}{27} \ ,$$

proving the inequality for a.

For b and c, we note that $P = b^2 - 3ac \le \sqrt{\Delta} \le \sqrt{X}$; hence the lower bound $c \ge (b^2 - \sqrt{X})/(3a)$ is clear. Furthermore, the inequality $Q \le P$ gives $bc - 9ad \le b^2 - 3ac$, hence $9ad \ge (b + 3a)c - b^2$. The inequality $P \le R$ thus gives

$$b^2 - 3ac \le c^2 - 3bd \le c^2 - \frac{b}{3a}((b+3a)c - b^2) \ ,$$

or in other words the quadratic inequality

$$c^2 + c\left(3a - \frac{b^2}{3a} - b\right) + \frac{b^3}{3a} - b^2 \ge 0 \ .$$

The roots of the polynomial in c are $b^2/(3a)$ and $b - 3a$, and since $b - 3a \le b^2/(3a)$ (the corresponding quadratic equation having a negative discriminant), we have $c \le b - 3a$ or $c \ge b^2/(3a)$. The latter is impossible, however, since it would imply that $P \le 0$. Thus $c \le b - 3a$, proving the inequalities $(b^2 - \sqrt{X})/(3a) \le c \le b - 3a$. In particular, it implies that $b^2 - 3ab + 9a^2 - \sqrt{X} \le 0$, hence b lies between the roots of this quadratic equation; in other words,

$$\frac{3a}{2} - \sqrt{\sqrt{X} - \frac{27a^2}{4}} \le b \le \frac{3a}{2} + \sqrt{\sqrt{X} - \frac{27a^2}{4}} \ ,$$

as claimed. It is immediately checked that this lower bound for b is sharper than the trivial lower bound $b \ge 0$ if and only if $a > X^{1/4}/3$.

The upper bounds for a and b are sharp, since they are reached for $P = Q = R$.

To finish the proof of the proposition, we must prove the other lower bound for c. By Proposition 8.3.2, we have $3aR - bQ + cP = 0$. Since $R = (3\Delta + Q^2)/(4P)$, this gives

$$9a\Delta + 3aQ^2 - 4bPQ + 4cP^2 = 0 \ ,$$

and solving this quadratic equation in Q gives

$$Q = \frac{2bP + \varepsilon\sqrt{4P^3 - 27a^2\Delta}}{3a}$$

for some $\varepsilon = \pm 1$. From the inequalities $|Q| \le P \le R = (3\Delta + Q^2)/(4P)$, we obtain $\sqrt{4P^2 - 3\Delta} \le |Q| \le P$. Hence

$$3a\sqrt{4P^2 - 3\Delta} \le \left|2bP + \varepsilon\sqrt{4P^3 - 27a^2\Delta}\right| \le 3aP \ ,$$

where it is understood that the lower inequality holds only when $4P^2 - 3\Delta \ge 0$.

To simplify the algorithm, we will take into account only the upper inequality. Taking into account the lower inequality (for $4P^2 - 3\Delta \ge 0$) would give an additional restriction on c, but the gain would be marginal compared to the expense of computing the precise necessary bounds.

We thus obtain

$$-P(3a + 2b) \le \varepsilon\sqrt{4P^3 - 27a^2\Delta} \le P(3a - 2b) \ ,$$

and an easy calculation shows that this is equivalent to

$$4P^3 - P^2(3a - 2\varepsilon b)^2 - 27a^2\Delta \le 0 \ .$$

Since $-P^2(3a+2b)^2 \le -P^2(3a-2b)^2$, the existence of $\varepsilon = \pm 1$ satisfying the above inequality is equivalent to the single inequality for $\varepsilon = -1$ — that is, to $4P^3 - P^2(3a+2b)^2 - 27a^2\Delta \le 0$, which of course implies $4P^3 - P^2(3a+2b)^2 - 27a^2X \le 0$.

The cubic function $f(P) = 4P^3 - P^2(3a+2b)^2 - 27a^2X$ satisfies $f(0) < 0$, $f'(0) = 0$, $f''(0) < 0$ and tends to $\pm\infty$ when P tends to $\pm\infty$. It follows that f has a unique real root P_2 that is larger than the nonzero root $(3a+2b)^2/6$ of f' and, in particular, is positive. Thus we must have $P \le P_2$, hence $c = (b^2 - P)/(3a) \ge (b^2 - P_2)/(3a)$. Furthermore, it is easily checked that this inequality for c is sharper than the simpler inequality $c \ge (b^2 - \sqrt{X})/(3a)$ if and only if $a \le X^{1/4}/3$ and $b \le -3a/2 + \sqrt{\sqrt{X} - 27a^2/4}$, finishing the proof of the proposition. □

With the inequalities of Proposition 8.5.5, it is not difficult to show that the number of quadruples (a,b,c,d) that will have to be checked is *linear* in X. In an actual implementation, we will first loop on a, then on b, then on c, and finally on d, satisfying the inequalities coming from $|Q| \le P \le R$ and from $\mathrm{disc}(F) \le X$; in other words, the inequalities

$$|bc - 9ad| \le b^2 - 3ac \le c^2 - 3bd$$

and

$$(-27a^2)d^2 + 2(9abc - 2b^3)d + (b^2c^2 - 4ac^3 - X) \le 0 \ .$$

The total number of triplets (a,b,c) satisfying the inequalities of the proposition is $O(X^{3/4})$. For some of these triplets the loop on d will be empty, and for the others we will have to examine approximately the number $H_3^+(X)$ of reduced real binary cubic forms of discriminant up to X (only approximately, because special cases have to be considered, but they add a negligible number of forms). Hence the total number of cases to be examined will be $H_3^+(X) + O(X^{3/4})$. We have the following important theorem concerning this quantity.

Theorem 8.5.6. *Let $H_3^+(X)$ (resp., $N_3^+(X)$) be the number of equivalent real cubic forms (resp., of isomorphism classes of real cubic fields) of discriminant less than or equal to X. Then as $X \to \infty$, we have*

$$H_3^+(X) = \frac{\pi^2}{72}X + C^+ \cdot X^{5/6} + O(X^{2/3+\epsilon}) \sim 0.137\ldots X \ ,$$

$$N_3^+(X) = \frac{1}{12\zeta(3)}X + O\left(Xe^{-c\sqrt{\log X \log\log X}}\right) \sim 0.0693\ldots X$$

for a known constant C^+ and any $c < 1/\sqrt{24}$.

Remarks

(1) The precise form for the remainder term for $H_3^+(X)$ is due to T. Shintani (see [Shin]), improving on Davenport's original result. The remainder term for $N_3^+(X)$ was proved by K. Belabas in [Bel2].
(2) For fascinating new developments in the theory of L-functions associated to binary cubic forms initiated by Shintani's work, see [Nak1].
(3) It has been conjectured by several authors (see [Rob], [Wri2]), that $N_3^+(X) = X/(12\zeta(3)) + C_N^+ X^{5/6} + o(X^{5/6})$ for an explicit constant C_N^+ (and similarly in the complex case, see Theorem 8.6.5 below). This is excellent agreement with the tables computed at the end of this chapter by K. Belabas.

It follows that the total number of steps in our algorithm will be linear in X and, in fact, approximatebly $0.137/.0693 \sim 1.977$ times more than the number of fields that we have to find; hence there will be very little waste.

We could still try to gain a little by avoiding the $O(X^{3/4})$ empty loops that we have mentioned. For this, we would need to find the exact range of values of c, given a and b. The result involves several cases and algebraic equations of degree even larger than 3, and this would probably slow down the final algorithm.

We can also easily characterize subclasses of real cubic fields. For example, we have the following.

Proposition 8.5.7. *Let K be a totally real cubic number field, F_K the unique reduced form associated to K, and H_K its Hessian. Then*

(1) *K is cyclic (that is, $d(K) = f^2$) if and only if $H_K = f_H(1, \pm 1, 1)$.*
(2) *$d(K) = 5f^2$ if and only if $H_K = f_H(1, \pm 1, 4)$ or $H_K = f_H(2, \pm 1, 2)$.*
(3) *$d(K) = 8f^2$ if and only if $H_K = f_H(1, 0, 6)$ or $H_K = f_H(2, 0, 3)$.*
(4) *$d(K) = 12f^2$ if and only if $H_K = f_H(1, 0, 9)$, $H_K = f_H(2, \pm 2, 5)$, or $H_K = f_H(1, 0, 1)$.*
(5) *Let $\Delta > 0$ be a fundamental discriminant. Then $d(K) = \Delta f^2$ if and only if H_K is a multiple of a primitive, reduced, positive definite, quadratic form of discriminant -3Δ (case $f_H = f$) or $-\Delta/3$ (case $f_H = 3f$ and $3 \mid f$).*

Proof. Using Proposition 8.4.8, proving this is just a matter of listing reduced quadratic forms. We leave it to the reader (see Exercise 8). □

8.6 Complex Cubic Fields

We now consider the case where $\mathrm{disc}(F) < 0$, hence when F corresponds to a complex cubic field. In this case the Hessian is an indefinite quadratic form, and in general there will be many reduced quadratic forms equivalent to it.

Instead of using the Hessian, we will use an old idea due to Matthews and Berwick. If $\mathrm{disc}(F) < 0$, then F has a unique real root θ, and $\theta \notin \mathbb{Q}$ if F is irreducible, so if we factor F in $\mathbb{R}[X,Y]$ as

$$F(x,y) = (x - \theta y)(Ax^2 + Bxy + Cy^2) \ ,$$

the quadratic form $H_F = (A,B,C)$ will be definite (that is, $B^2 - 4AC < 0$) but with real nonrational coefficients. We are going to show that the form (A,B,C) has many of the properties of the Hessian. An even better idea due to G. Julia and J. Cremona is presented in Exercise 9; see [Cre].

An easy computation gives

$$\mathrm{disc}(F) = (B^2 - 4AC)(A\theta^2 + B\theta + C)^2 \ .$$

By changing (x,y) to $(-x,-y)$, which changes F into $-F$, we may assume that $A \geq 0$. If $F = (a,b,c,d)$, we have

$$A = a, \quad B = a\theta + b, \quad C = a\theta^2 + b\theta + c \ ,$$

hence

$$a = A, \quad b = B - \theta A, \quad c = C - \theta B, \text{ and } \quad d = -\theta C \ .$$

If $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{GL}_2(\mathbb{Z})$, a simple computation shows that

$$H_{F \circ M} = |a - \theta c| \cdot H_F \circ M \ ,$$

where the absolute value sign comes from the choice $A \geq 0$.

Definition 8.6.1. *Let $F = (a,b,c,d)$ be an integral, binary, complex cubic form, and let (A,B,C) be defined as above. We say that F is* reduced *if $0 < |B| < A < C$ and if, in addition, $a > 0$, $b \geq 0$, $d \neq 0$, and $d > 0$ if $b = 0$.*

Note that when F is irreducible, θ is irrational; hence the special cases $B = 0$, $A = |B|$, or $A = C$ that occurred in the real case cannot occur here. Another nice fact is that we do not need to compute the irrational numbers A, B, and C at all:

Lemma 8.6.2. *Let $F = (a,b,c,d)$ be a complex cubic form. Then F is reduced if and only if*

$$d^2 - a^2 + ac - bd > 0 \ ,$$

$$-(a-b)^2 - ac < ad - bc < (a+b)^2 + ac \ ,$$

$$a > 0, \quad b \geq 0, \quad d \neq 0, \text{ and } \quad d > 0 \text{ if } b = 0 \ .$$

Proof. The condition $B < A$ is equivalent to $a\theta + b < a$ and hence to $\theta < (a-b)/a$ since $a > 0$. Since F has only one real root (and again $a > 0$), this is equivalent to $F(a-b,a) > 0$, which gives $-(a-d)^2 - ac < ad - bc$. Similarly, the condition $-B < A$ is equivalent to $-a\theta - b < a$, hence to $\theta > -(a+b)/a$, so $F(-(a+b),a) < 0$, which gives $ad - bc < (a+b)^2 + ac$.

Finally, the condition $A < C$ is equivalent to $a\theta^2 + b\theta + (c-a) > 0$, and this is equivalent to $\mathcal{R}(F_1, G) > 0$, where $\mathcal{R}$ is the resultant, $F_1(X) = F(X, 1)$ and $G(X) = aX^2 + bX + (c-a)$. An immediate computation shows that $\mathcal{R}(F_1, G) = a^3(d^2 - a^2 + ac - bd)$, and this gives the last unproved condition. □

For this notion to be useful, we must have the analog of Proposition 8.5.4.

Proposition 8.6.3. (1) *Two equivalent, irreducible, reduced, complex cubic forms are equal.*
(2) *A reduced complex cubic form belonging to U is irreducible.*
(3) *Any irreducible complex cubic form is equivalent to a unique reduced form.*

Proof. (1). Let $F' = F \circ M$, where F and F' are reduced and $M \in \mathrm{GL}_2(\mathbb{Z})$. Then by the formula proved above, there exists $\lambda > 0$ such that $H_{F'} = \lambda H_F \circ M$. As before, we deduce from the inequalities $|B| < A < C$ and $|B'| < A' < C'$ that $H_{F'} = \lambda H_F$ and hence that M is an automorphism of H_F. The proof then terminates as in the real case, except that there are no special cases to consider since the forms are irreducible.

(2). As in the real case, a complex reducible form F belonging to U must be equivalent to $G = y(x^2 + \delta y^2)$ or to $G = y(x^2 + xy + \delta y^2)$ with $\delta \geq 1$. If $F = (x - \theta y)(Ax^2 + Bxy + Cy^2)$ and $F = G \circ M$ with $M \in \mathrm{GL}_2(\mathbb{Z})$, then $H_F = \lambda H_G \circ M$ for some $\lambda > 0$. Hence the reduced form H_F is equivalent to a multiple of H_G with $H_G = (1, 0, \delta)$ or $H_G = (1, 1, \delta)$, respectively, which are also reduced, and it is therefore equal to that multiple or its inverse. Hence we have either $B = 0$ or $B = \pm A$ for the form F, which are both excluded from the definition of a reduced form in the complex case.

(3). We can reduce first H_F by an element of $\mathrm{GL}_2(\mathbb{Z})$ so that it satisfies $0 < |B| < A < C$, the strict inequalities being guaranteed by the irreducibility of F. We must have $d \neq 0$ as F would otherwise be reducible, and since $A = a$, we have $a > 0$. Changing (x, y) into $(x, -y)$ if necessary (which changes B into $-B$ and leaves A and C unchanged), we may also assume that $b \geq 0$. Finally, if $b = 0$, again changing (x, y) into $(x, -y)$ if necessary, we may assume $d > 0$ since $d \neq 0$. □

We have bounds on the coefficients of a reduced complex cubic form as follows.

Proposition 8.6.4. *Let $F = (a, b, c, d)$ be a reduced form such that $-X \leq \mathrm{disc}(F) < 0$. We have the following inequalities:*

(1)

$$1 \leq a \leq \frac{2X^{1/4}}{3^{3/4}} \ ;$$

(2)
$$0 \le b \le \frac{3a}{2} + \sqrt{\left(\frac{X}{3}\right)^{1/2} - \frac{3a^2}{4}} \ ;$$

(3)
$$1 - b \le c \le \left(\frac{X}{4a}\right)^{1/3} + U(a,b) \ ,$$

where $U(a,b) = b^2/(3a)$ *if* $b \le 3a/2$, *while* $U(a,b) = b - 3a/4$ *if* $b > 3a/2$.

Proof. Set $\Delta = |\mathrm{disc}(F)|$ and $3D = 4AC - B^2$. The inequalities $|B| < A < C$ imply as usual $A^2 < D$ or, equivalently, $a^2 < D$. In addition, by a computation made above, we have $\Delta = 3D(A\theta^2 + B\theta + C)$. Solving this as a quadratic equation in θ, we obtain

$$2a\theta = -B \pm \sqrt{4a\left(\frac{\Delta}{3D}\right)^{1/2} - 3D} \ .$$

Since the expression under the square root is nonnegative, we obtain $16a^2\Delta \ge 27D^3 \ge 27a^6$, proving (1) since $\Delta \le X$.

From the expression for $2a\theta$, we also obtain

$$b = B - a\theta = \frac{3B}{2} \pm \sqrt{a\left(\frac{\Delta}{3D}\right)^{1/2} - \frac{3D}{4}} \ .$$

Since the expression under the square root is a decreasing function of D and since $D \ge a^2$ and $|B| \le A = a$, we obtain

$$b \le \frac{3a}{2} + \sqrt{\left(\frac{\Delta}{3}\right)^{1/2} - \frac{3a^2}{4}} \ ,$$

proving the inequality for b.

We have $c = C - \theta B > A - |\theta|\, A = a - |a\theta|$. But $|a\theta| = |b + a\theta - b| \le b + |B| \le b + A = a + b$; hence $c > -b$ so $c \ge 1 - b$, as claimed, since b and c are integers.

For the upper bound, we check that

$$4ac = -3B^2 + 4bB + 3D \ .$$

This is a quadratic in B whose derivative is positive for $B < 2b/3$ and negative for $B > 2b/3$. Since we must have $-a < B < a$, it follows that the maximum of this quadratic is attained for $B = 2b/3$ when $2b/3 \le a$ and for $B = a$ when $2b/3 > a$. The upper bound for c follows immediately. □

In the actual implementation we will proceed essentially as in the real case. The analog of Theorem 8.5.6 is as follows.

Theorem 8.6.5. *Let $H_3^-(X)$ (resp., $N_3^-(X)$) be the number of equivalent complex cubic forms (resp., of isomorphism classes of complex cubic fields) of discriminant greater than or equal to $-X$. Then as $X \to \infty$, we have*

$$H_3^-(X) = \frac{\pi^2}{24}X + C^- \cdot X^{5/6} + O(X^{2/3+\epsilon}) \sim 0.411\ldots X \ ,$$
$$N_3^-(X) = \frac{1}{4\zeta(3)}X + O\left(Xe^{-c\sqrt{\log X \log\log X}}\right) \sim 0.208\ldots X$$

for a known constant C^- and any $c < 1/\sqrt{24}$.

Once again, we see that the algorithm will be linear in X, and the number of loops will be approximately 1.977 times the number of cubic fields found (see also the remarks following Theorem 8.5.6).

8.7 Implementation and Results

8.7.1 The Algorithms

We are now ready to give the complete algorithms for computing tables of cubic fields.

First, we give the final version of the algorithm for testing whether or not a form belongs to the image of the Davenport–Heilbronn map, which slightly improves on Algorithm 8.4.7 using the reducedness condition.

Algorithm 8.7.1 (Cubic Form Test). Given $F = (a, b, c, d)$ a *reduced* cubic form, $H_F = (P, Q, R)$ its Hessian, and $D = Q^2 - 4PR = -3\,\mathrm{disc}(F)$ the discriminant of H_F, this algorithm outputs `true` or `false` according to whether or not F corresponds to the image of a cubic field by the Davenport–Heilbronn map.

1. [Special case] If there exists a prime p such that $p^2 \mid a$ and $p \mid b$, return `false`.
2. [Check condition at 2] If $16 \mid D$ or if $D \equiv 4 \pmod{16}$ and P or R is odd, return `false`.
3. [Compute f_H] Set $f_H \leftarrow \gcd(P, Q, R)$.
4. [Check conditions at 3] If $27 \mid D$ and if either $3 \nmid f_H$ or the conditions of Proposition 8.3.3 (5) are not satisfied, return `false`.
5. [Check f_H almost squarefree] If $p^2 \mid f_H$ for some $p > 3$, return `false`.
6. [Check f_H and D/f_H^2 almost coprime] Set $t \leftarrow D/f_H^2$, $t \leftarrow t/\gcd(t, 72)$ (so that t is now prime to 6). If $\gcd(t, f_H) > 1$, return `false`.
7. [Check D/f_H^2 almost squarefree] If t is squarefree, return `true`; otherwise, return `false`.

Proof. First, let us show the validity of step 1. If $4 \mid a$ and $2 \mid b$, then $\operatorname{disc}(F) \equiv b^2c^2 \pmod{16}$. Hence either $b/2$ or c is even and then $16 \mid \operatorname{disc}(F)$, or $b/2$ and c are odd and then $\operatorname{disc}(F) \equiv 4 \pmod{16}$ and $R = c^2 - 3bd$ is odd, so in any case $F \notin U_2$. If $9 \mid a$ and $3 \mid b$, then $9 \mid \operatorname{disc}(F)$, so we have $F = (1^3)$ and $9 \mid a$ implies that $F \notin U_3$ by Proposition 8.3.3. Finally, if $p^2 \mid a$ and $p \mid b$ with $p > 3$, then $p^2 \mid \operatorname{disc}(F)$, hence $F \notin V_p$ so $p \mid f_H$, and in particular $p \mid c$; hence $p^4 \mid \operatorname{disc}(F)$, so $F \notin U_p$ by Proposition 8.3.3, proving our claim.

Since F is reduced, Propositions 8.5.4 and 8.6.3 show that F is irreducible, hence step 1 of Algorithm 8.4.7 is satisfied.

I claim that if the above algorithm returns `true`, then F is also primitive. Indeed, let p be a divisor of all the coefficients of F. Then $p^2 \mid f_H$. If $p > 3$, then step 5 will have returned `false`. If $p = 2$, we will have $16 \mid \operatorname{disc}(F)$ so step 2 will have returned `false`. Finally, if $p = 3$, then $27 \mid D = -3\operatorname{disc}(F)$ and $3 \mid a$, $3 \mid d$, hence step 4 will have returned `false`.

As the other steps are the same as in Algorithm 8.4.7, this proves the algorithm's validity. □

Remark. In this algorithm, we must check that f_H and t are squarefree. Since this may be done on billions of forms, it becomes too lengthy to do this squarefreeness test in a naive way. To avoid this, one should use *sieve* methods. The details are left to the reader (Exercise 10), who can also refer to [Bel1] and [Bel3].

From the above algorithm and the definition of reduced forms, it is easy to write algorithms for making tables of cubic fields.

Algorithm 8.7.2 (Real Cubic Field Table). Given a positive number X, this algorithm outputs the reduced defining polynomial of all real cubic fields of discriminant less than or equal to X, as well as their total number.

1. [Initialize loop on a] Set $x \leftarrow \sqrt{X}$, $n \leftarrow 0$, $U_a \leftarrow \lfloor 2\sqrt{x/27} \rfloor$, $M_a \leftarrow \lfloor \sqrt{x}/3 \rfloor$, and $a \leftarrow 0$.
2. [Loop on a, terminate?] Set $a \leftarrow a+1$. If $a > U_a$, output the total number n and terminate the algorithm. Otherwise, let f_a be the product of the primes whose square divides a.
3. [Initialize loop on b] Set $U_b \leftarrow \lfloor 3a/2 + \sqrt{x - 27a^2/4} \rfloor$. If $a \le M_a$, set $L_b \leftarrow 0$ and $M_b \leftarrow U_b - 3a$; otherwise, set $L_b \leftarrow 3a - U_b$ (we do not need M_b in this case). Finally, set $b \leftarrow L_b - 1$.
4. [Loop on b] Set $b \leftarrow b + 1$. If $b > U_b$, go to step 2. If $\gcd(f_a, b) > 1$, go to step 4.
5. [Initialize loop on c] Set $U_c \leftarrow b - 3a$. If $a > M_a$ or $a \le M_a$ and $b > M_b$, set $L_c \leftarrow \lceil (b^2 - x)/(3a) \rceil$. Otherwise, compute an upper bound y for the unique positive solution of the equation $4P_2^3 - (3a + 2b)^2 P_2^2 - 27a^2X = 0$ and set $L_c \leftarrow \lceil (b^2 - y)/(3a) \rceil$. Finally, set $c \leftarrow L_c - 1$.
6. [Loop on c] Set $c \leftarrow c + 1$. If $c > U_c$, go to step 4.

7. [Initialize loop on d] Set $L_d \leftarrow \lceil (bc+3ac-b^2)/(9a) \rceil$. If $b = 0$, set $U_d \leftarrow -1$; otherwise, set $U_d \leftarrow \min(\lfloor (c^2+3ac-b^2)/(3b) \rfloor, \lfloor (b^2-3ac+bc)/(9a) \rfloor)$.
8. [Initialize loop on d (continued)] Set $S \leftarrow 4(b^2-3ac)^3 - 27a^2X$. If $S \leq 0$, set $E_d \leftarrow [L_d, U_d]$. Otherwise, set $s \leftarrow \sqrt{S}$, $d_1 \leftarrow \lfloor (9abc-2b^3-s)/(27a^2) \rfloor + 1$, $d_2 \leftarrow \lceil (9abc-2b^3+s)/(27a^2) \rceil - 1$, and $E_d \leftarrow [L_d, U_d] \cap \mathbb{C}[d_1, d_2]$ (now E_d is either empty, an interval, or a union of two intervals). Finally, set $d \leftarrow \min(E_d) - 1$.
9. [Loop on d] If no elements of E_d are strictly larger than d, go to step 6. Otherwise, replace d by the smallest element of E_d strictly larger than d.
10. [Compute Hessian] Set $F \leftarrow (a,b,c,d)$, $P \leftarrow b^2 - 3ac$, $Q \leftarrow bc - 9ad$, $R \leftarrow c^2 - 3bd$, $H_F \leftarrow (P,Q,R)$, and $D \leftarrow Q^2 - 4PR$. If $|D| > 3X$, go to step 9.
11. [Check special cases] (Here, we know that $|Q| \leq P \leq R$ and $0 < \mathrm{disc}(F) \leq X$.) If $P = Q$ and $b \geq |3a-b|$, go to step 9. If $P = R$ and $a > |d|$ or $a = |d|$ and $b \geq |c|$, go to step 9.
12. [Check form belongs to U] (Here $F = (a,b,c,d)$ is reduced.) Using Algorithm 8.7.1 on the form F (omitting step 1, which has already been performed), check whether F belongs to the image of the Davenport–Heilbronn map. If the answer is `true`, output the reduced polynomial $F(x,1) = ax^3 + bx^2 + cx + d$ and set $n \leftarrow n+1$. In any case, go to step 9.

Proof. This algorithm's validity is an immediate consequence of Proposition 8.5.5 together with the inequalities for d coming from $|Q| \leq P \leq R$ and $\mathrm{disc}(F) \leq X$. Note that in step 8, we compute d_1 and d_2 to avoid missing any field, but at the price of having a few extra fields that are then rejected by the test $|D| \leq 3X$ made in step 10. □

The algorithm in the complex case, which follows, is a little more complicated because of the inequalities for d.

Algorithm 8.7.3 (Complex Cubic Field Table). Given a positive number X, this algorithm outputs the reduced defining polynomial of all complex cubic fields of discriminant less than or equal to X in absolute value, as well as their total number.

1. [Initialize loop on a] Set $x \leftarrow \sqrt{X/3}$, $n \leftarrow 0$, $U_a \leftarrow \lfloor 2\sqrt{x/3} \rfloor$, and $a \leftarrow 0$.
2. [Loop on a, terminate?] Set $a \leftarrow a+1$. If $a > U_a$, output the total number n and terminate the algorithm. Otherwise, let f_a be the product of the primes whose square divides a.
3. [Initialize loop on b] Set $U_b \leftarrow \lfloor 3a/2 + \sqrt{x - 3a^2/4} \rfloor$. If $f_a = 1$, set $b \leftarrow -1$; otherwise, set $b \leftarrow 0$.
4. [Loop on b] Set $b \leftarrow b+1$. If $b > U_b$, go to step 2. If $\gcd(f_a, b) > 1$, go to step 4.

5. [Initialize loop on c] Set $L_c \leftarrow 1-b$. If $3a \geq 2b$, set $U_c \leftarrow \lfloor b^2/(3a) + \sqrt[3]{X/(4a)} \rfloor$; otherwise, set $U_c \leftarrow \lfloor b - 3a/4 + \sqrt[3]{X/(4a)} \rfloor$. Finally, set $c \leftarrow L_c - 1$.
6. [Loop on c] Set $c \leftarrow c+1$. If $c > U_c$, go to step 4.
7. [Initialize loop on d] Set $s \leftarrow \sqrt{4(b^2-3ac)^3 + 27a^2X}$ (this will be a real number),

$$L_d \leftarrow \max\left(1 + \left\lfloor \frac{(b-a)(a-b+c)}{a} \right\rfloor, \left\lceil \frac{9abc - 2b^3 - s}{27a^2} \right\rceil\right),$$

$$U_d \leftarrow \min\left(-1 + \left\lceil \frac{(a+b)(a+b+c)}{a} \right\rceil, \left\lfloor \frac{9abc - 2b^3 + s}{27a^2} \right\rfloor\right),$$

and if $b = 0$, set $L_d \leftarrow \max(L_d, 1)$.
8. [Initialize loop on d (continued)] Let $P \leftarrow b^2 - 3ac$. If $P < 0$, set $E_d \leftarrow [L_d, U_d]$. Otherwise, set $s \leftarrow \sqrt{4P^3}$, $d_1 \leftarrow \lceil (9abc - 2b^3 - s)/(27a^2) \rceil$, $d_2 \leftarrow \lfloor (9abc - 2b^3 + s)/(27a^2) \rfloor$, and $E_d \leftarrow [L_d, U_d] \cap \mathbb{C}[d_1, d_2]$. (Now E_d is either empty, an interval, or a union of two intervals.)
9. [Initialize loop on d (continued again)] Let $S \leftarrow b^2 + 4a^2 - 4ac$. If $S < 0$, go to step 10. Otherwise, set $s \leftarrow \sqrt{S}$, $d_3 \leftarrow \lceil (b-s)/2 \rceil - 1$, $d_4 \leftarrow \lfloor (b+s)/2 \rfloor + 1$, and $E_d \leftarrow E_d \cap \mathbb{C}[d_3, d_4]$. (Now E_d is a union of at most four intervals.) Finally, set $d \leftarrow \min(E_d) - 1$.
10. [Loop on d] If no elements of E_d are strictly larger than d, go to step 6. Otherwise, replace d by the smallest element of E_d strictly larger than d.
11. [Compute Hessian] Set $F \leftarrow (a,b,c,d)$, $P \leftarrow b^2 - 3ac$, $Q \leftarrow bc - 9ad$, $R \leftarrow c^2 - 3bd$, $H_F \leftarrow (P,Q,R)$, and $D \leftarrow Q^2 - 4PR$. If $D > 3X$, go to step 10.
12. [Check form belongs to U] (Here $F = (a,b,c,d)$ is reduced and $-X \leq \mathrm{disc}(F) \leq 0$.) Using Algorithm 8.7.1 on the form F (omitting step 1, which has already been performed), check whether F belongs to the image of the Davenport–Heilbronn map. If the answer is `true`, output the reduced polynomial $F(x,1) = ax^3 + bx^2 + cx + d$ and set $n \leftarrow n+1$. In any case, go to step 10.

8.7.2 Results

The following tables give for $1 \leq n \leq 11$ the total number $N_3^+(X)$ and $N_3^-(X)$ of real and complex cubic fields of discriminant up to $X = 10^n$ in absolute value, as well as the maximal value a for the first coefficient of the cubic form in the interval considered. This value is always within 1 of the largest possible value given by Propositions 8.5.5 and 8.6.4.

X	Number of fields	a
10^1	0	0
10^2	2	1
10^3	27	2
10^4	382	3
10^5	4,804	6
10^6	54,600	12
10^7	592,922	21
10^8	6,248,290	38
10^9	64,659,361	68
10^{10}	661,448,081	121
10^{11}	6,715,824,025	216

Real cubic fields

X	Number of fields	a
10^1	0	0
10^2	7	1
10^3	127	3
10^4	1520	7
10^5	17,041	14
10^6	182,417	26
10^7	1,905,514	49
10^8	19,609,185	86
10^9	199,884,780	155
10^{10}	2,024,660,098	276
10^{11}	20,422,230,540	492

Complex cubic fields

It should be remarked that, thanks to the notion of reducedness, our algorithm gives for every cubic field a canonical defining polynomial (which we can call *reduced*) and which, in addition, has all the nice properties described in Section 8.4. In particular, the integral basis and decomposition of primes is immediate. One consequence is that when the cubic number field does not have a power basis, the polynomial we will find will not be monic. If a power basis exists, however, the reduced polynomial produced by our algorithm is not necessarily monic. See Exercises 11, 12, and 13 for some examples.

8.8 Exercises for Chapter 8

1. Let F be a binary form of degree n with coefficients (a_i) and roots $(\alpha_i : \beta_i)$ normalized as explained in the text. Show that changing a_i into λa_i is equivalent to changing α_i into $\lambda^{1/n}\alpha_i$ and β_i into $\lambda^{1/n}\beta_i$, and deduce from this that the discriminant is a homogeneous polynomial of degree $2n-2$ in the variables of the form F.
2. Let f be a covariant of degree d, weight w, with image in Φ_m. Using differential equations, show that, as stated in the text, we have $w = (nd-m)/2$.
3. Let K be a quartic number field with Galois group isomorphic to the dihedral group D_4, and let $\mathcal{B} = (1, \alpha, \beta, \gamma)$ be an integral basis of K.
   a) Define a ternary form $F_{\mathcal{B}}(x, y, z)$ of degree 6 in a manner analogous to Proposition 8.2.1.
   b) Show that there exists a ternary quadratic form $F_{\mathcal{B},2}(x, y, z)$ and a ternary quartic form $F_{\mathcal{B},4}(x, y, z)$ such that $F_{\mathcal{B}}(x, y, z) = F_{\mathcal{B},2}(x, y, z)F_{\mathcal{B},4}(x, y, z)$.
   c) Prove some results analogous to those of Proposition 8.2.1 for the forms $F_{\mathcal{B},2}$ and $F_{\mathcal{B},4}$.
4. Let K be a cubic field such that δ is an inessential discriminantal divisor. Show that $\delta \leq 2$ and that $\delta = 2$ is possible if and only if 2 is totally split in K.
5. More generally, let L/K be an extension of number fields, let $\mathfrak{p}$ be a prime ideal of K, let $q = \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{p})$, and for any integer f denote by $r(f)$ the number of prime ideals $\mathfrak{P}$ of L above $\mathfrak{p}$ such that $f(\mathfrak{P}/\mathfrak{p}) = f$. Show that $\mathfrak{p}$ is an

inessential discriminantal divisor (in other words, that $\mathfrak{p}$ divides $d(\theta)/\mathfrak{d}(L/K)$ for any $\theta \in \mathbb{Z}_L$ such that $L = K(\theta)$) if and only if there exists f such that

$$r(f) > \frac{1}{f}\sum_{d|f} \mu(d) q^{f/d} \ .$$

Deduce from this that if $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{p}) \geq [L : K]$, then $\mathfrak{p}$ is not an inessential discriminantal divisor (see [Has2] for hints on this exercise).

6. With the notation of the proof of Proposition 8.3.3, show that $p^3 \mid \mathrm{disc}(G)$ implies that $G \notin V_p$ when $p \neq 3$.
7. Prove Corollary 8.3.4.
8. Fill in the details of the proof of Proposition 8.5.7.
9. (J. Cremona) Let $F = (a, b, c, d)$ be an irreducible binary cubic form corresponding to a complex cubic field, in other words, such that $\mathrm{disc}(F) < 0$. Let θ be the unique real root of $F(x, 1) = 0$, and let β and $\overline{\beta}$ be the two other complex conjugate roots of $F(x, 1) = 0$. Consider the binary quadratic form

$$J_F(X, Y) = a^2 \left(2(\theta - \beta)(\theta - \overline{\beta})(X - \beta Y)(X - \overline{\beta} Y) - (\beta - \overline{\beta})^2 (X - \theta Y)^2\right) \ .$$

   a) Compute the coefficients h_i of the form $J_F(X, Y) = h_0 X^2 + h_1 XY + h_2 Y^2$ in terms of a, b, c, d, and θ.
   b) If $M \in \mathrm{GL}_2(\mathbb{Z})$, compute $J_{F \circ M}$ in terms of $J_F \circ M$.
   c) Compute the discriminant of $J_F(X, Y)$ in terms of $\mathrm{disc}(F)$, and deduce that $J_F(X, Y)$ is a positive definite quadratic form.
   d) Define F to be *Julia-reduced* if $a > 0$ and J_F is a reduced, positive definite, binary quadratic form. Give an algorithm for Julia-reducing a cubic form (using θ explicitly), and show that in any class of cubic forms there exists a unique Julia-reduced cubic form.
   e) Assume that $F = (a, b, c, d)$ is Julia-reduced, and let X be such that $-X \leq \mathrm{disc}(F) < 0$. Show that $a \leq (2/3)^{3/2} X^{1/4}$ and that $|P| = |b^2 - 3ac| \leq 2^{1/3} X^{1/2}$.

   Deduce from these results an algorithm analogous to Algorithm 8.7.3 to enumerate complex cubic fields (because of the better bound on a, this algorithm will be faster, but probably only slightly).
10. Modify Algorithm 8.7.1 by precomputing tables and using a sieve, so that the squarefreeness tests of f_H and t become as fast as possible (see [Bel1]).
11. Let $K = \mathbb{Q}(\alpha)$ be the cyclic cubic field of discriminant 961 given by a root α of the monic polynomial $X^3 + X^2 - 10X - 8$.

    a) Show that $(1, \alpha, (\alpha^2 + \alpha)/2)$ is an integral basis of K.
    b) Prove that $\mathbb{Z}_K$ does not have a power basis $(1.\theta, \theta^2)$.
    c) Using the algorithm given in the text, show that the reduced polynomial for the field K is the nonmonic polynomial $2X^3 + X^2 - 5X - 2$.

12. Let $K = \mathbb{Q}(\alpha)$ be the noncyclic cubic field of discriminant 1304 given by a root α of the monic polynomial $X^3 - X^2 - 11X - 1$.

    a) Show that $(1, \alpha, (\alpha^2 + 1)/2)$ is an integral basis of K.
    b) Prove that $\mathbb{Z}_K$ does not have a power basis $(1.\theta, \theta^2)$.
    c) Using the algorithm given in the text, show that the reduced polynomial for the field K is the nonmonic polynomial $2X^3 + 3X^2 - 4X - 2$.

13. Let $K = \mathbb{Q}(\alpha)$ be the noncyclic cubic field of discriminant 2228 given by a root α of the monic polynomial $X^3 - 14X - 18$.
    a) Show that $(1, \alpha, \alpha^2)$ is an integral basis of K, so that K has a power basis.
    b) Using the algorithm given in the text, show that the reduced polynomial for the field K is, however, the nonmonic polynomial $2X^3 + 2X^2 - 6X - 1$.

9. Number Field Table Constructions

9.1 Introduction

In this chapter, we will describe in detail the known methods for computing tables of number fields with small discriminant. We can try to pursue two different goals.

One goal is to find a systematic list of all number fields up to isomorphism, for a given degree, signature, and possibly also Galois group or other properties, up to some bound. Even then, the *ordering* of the fields is not completely canonical. We usually choose the fields ordered by increasing absolute value of their discriminant. It is, however, quite possible to consider completely different orderings, for instance, by specifying a very small number of ramified primes. We do not consider these types of orderings in this book.

A complementary goal, which is the one we try to reach if the first one is too difficult, is to find *some* number fields with the desired properties and not too large discriminant.

We will also consider the corresponding problems in the *relative* case, at least for small relative degrees.

The methods used for computing these tables or these fields are quite diverse and differ enormously in their efficiency. Historically, the first methods used were based on the geometry of numbers, essentially on Hunter's theorem (see [Coh0, Theorem 6.4.2]), and more recently on Martinet's theorem (see Theorem 9.3.2 below), which is a relative version of Hunter's theorem. These methods have the advantage of being completely systematic. Their enormous disadvantage is that one literally has to pick needles out of a haystack: to find one or two fields, it may be necessary to look at billions (10^9) of polynomials. For example, the number fields of smallest discriminant are not known in degree 10 or above, even assuming the GRH, which usually helps, and in degree 9 only the totally real case is known, assuming the GRH (see [Let]).

The other methods used are usually less systematic but substantially faster. One such method is the class field method, which we have already described in detail in Chapter 5. In this chapter, we give several applications of this method. Another method, used by D. Simon in [Sim2], for example,

is to look for *polynomials* having small discriminant by quite elementary but clever methods. We will study several of these methods.

Finally, recall that constructing a table of (absolute) quadratic fields is essentially trivial since it amounts to checking the squarefreeness of an integer, but note also that thanks to the work of K. Belabas, constructing a table of (absolute) cubic fields is almost as easy. We refer to Chapter 8 for complete details on this.

9.2 Using Class Field Theory

9.2.1 Finding Small Discriminants

If at first we are not too ambitious and only want to find examples of number fields with small discriminant, but not necessarily the smallest or a complete table, using class field theory, we can compute Abelian extensions corresponding to congruence subgroups $(\mathfrak{m}, C)$ and hence make *tables* of (hopefully interesting) number fields.

To begin with, we need tables of base fields K. For this, the only two large databases are available either from M. Pohst's Kant group in Berlin or from the author's Pari group in Bordeaux (see Appendix B).

These tables contain nearly a million number fields of degree up to 7, together with their invariants such as discriminant, integral basis, class group, unit group, and regulator.

For a given base field K in this list, we would like to compute a list of moduli $\mathfrak{m}$, the ray class group $Cl_{\mathfrak{m}}(K)$ using Algorithm 4.3.1, a list of congruence subgroups $(\mathfrak{m}, C)$ using the algorithms explained in Section 4.1.10, the discriminant and signature of the corresponding number fields L using the results of Section 3.5.2, and finally a defining polynomial for the fields L that we find interesting using one of the methods presented in Chapter 5 or Chapter 6.

We will almost always want number fields having relatively small root discriminant, and perhaps satisfying some additional conditions. Recall that by definition, the *root discriminant* of a number field L is

$$|d(L)|^{1/[L:\mathbb{Q}]} ,$$

where as usual $d(L)$ is the absolute discriminant of L. A well-known theorem of Minkowski implies that there exists a constant $c(R_1, R_2)$ (with $c(R_1, R_2) > 1$ if $(R_1, R_2) \neq (1, 0)$) depending only on the signature (R_1, R_2) of the number field, such that the root discriminant of any number field of signature (R_1, R_2) is at least equal to $c(R_1, R_2)$. The computation of the best value of this constant has been considerably improved by the use of analytic techniques introduced by H. Stark, A. Odlyzko, J.-P. Serre, G. Poitou, F. Diaz y Diaz, and others. The best known bounds (which we will simply

call the Odlyzko bounds) are obtained by assuming the generalized Riemann hypothesis (GRH), and so we will do this.

The GRH bounds have been carefully recomputed for all signatures and degrees up to 100 and are available by anonymous `ftp` (see Appendix B). An excerpt from this table is given in Appendix C.

Assume that we want to make a table of Abelian extensions of our base fields with degree at most equal (or exactly equal) to some integer N, with given signature (R_1, R_2), with root discriminant at most equal to some bound $b(R_1, R_2)$, and possibly with some extra conditions, such as a specific Galois group. We can proceed as follows.

(1) If the degree of L is specified to be a given number N, we may of course restrict to base fields K whose degree is a divisor of N. If we only ask that the degree of L is at most equal to N, we restrict to base fields whose degree is at most equal to $N/2$ (otherwise, we would have at most the trivial extension $L = K$).

(2) We can always restrict to congruence subgroups $(\mathfrak{m}, C)$ such that $\mathfrak{m}$ is the conductor of the congruence subgroup, hence of the corresponding extension L/K. In particular, if we want a given signature (R_1, R_2), Proposition 3.5.8 tells us exactly how many real places must divide the modulus $\mathfrak{m}$; in other words, it gives us $k = |\mathfrak{m}_\infty|$. This leaves us with $\binom{r_1}{k}$ choices for $\mathfrak{m}_\infty$ (where r_1 is the number of real places of K), which must all be looked at, unless $K/\mathbb{Q}$ is Galois, since in that case all the infinite places of K play the same role, so only the *number* of such places counts, leaving only one possible choice.

(3) One of the most important restrictions is that we want the root discriminant to be at most $b(R_1, R_2)$. This allows us to bound both the base field K and the norm of the modulus $\mathfrak{m}$ (as always, assumed to be the conductor), thanks to the following lemma.

Lemma 9.2.1. *Let L/K be an Abelian extension of number fields of conductor $\mathfrak{m}$, let $n_L = [L : \mathbb{Q}]$ and $n_K = [K : \mathbb{Q}]$, and assume that the root discriminant satisfies $|d(L)|^{1/n_L} \le B$ for some number B. Then*

$$\mathcal{N}(\mathfrak{m}) d(K)^2 \le B^{2n_K} \ .$$

In particular, the number of possible base fields K and moduli $\mathfrak{m}$ is finite.

Proof. By Theorem 3.5.11, we have

$$|d(L)|^{1/h_{\mathfrak{m},C}} = |d(K)| \frac{\mathcal{N}(\mathfrak{m})}{\prod_{\mathfrak{p}|\mathfrak{m}} \mathcal{N}(\mathfrak{p})^{\sum_{1 \le k \le v_{\mathfrak{p}}(\mathfrak{m})} h_{\mathfrak{m}/\mathfrak{p}^k,C}/h_{\mathfrak{m},C}}} \ .$$

Since $\mathfrak{m}$ is assumed to be the conductor, we have $h_{\mathfrak{m}/\mathfrak{p}^k,C} < h_{\mathfrak{m},C}$ for $k \ge 1$, and in particular $h_{\mathfrak{m}/\mathfrak{p}^k,C} \le h_{\mathfrak{m},C}/2$ since it is a divisor. From this and the above formula giving $d(L)$, we obtain

$$\mathcal{N}(\mathfrak{m})^{1/2}\,|d(K)| \le |d(L)|^{1/h_{\mathfrak{m},C}} = |d(L)|^{n_K/n_L} \le B^{n_K} \ ,$$

from which the result follows. □

The result of this lemma is usually very pessimistic, but the simple fact of having a bound on possible pairs $(K, \mathfrak{m})$ is important. Furthermore, the result does not depend on $h_{\mathfrak{m},C}$ so, in particular, not on the group C. For a given degree of $[L : K]$, the bound can often be considerably improved (see Exercise 1).

(4) Assume that we have now chosen the base field K and the modulus $\mathfrak{m}$ subject to the above restrictions. There now remains the task of enumerating congruence subgroups C modulo $\mathfrak{m}$ or, equivalently, subgroups $\overline{C}$ of $Cl_{\mathfrak{m}}$. In principle, this can be done using Birkhoff's Theorem 4.1.18 (or a more naive method if $Cl_{\mathfrak{m}}$ does not have too many cyclic components, or if we are looking only for subgroups of prime index), but in general the number of subgroups is exponentially large (recall, for example, that if p is prime, the number of subgroups of index p of the elementary p-group C_p^r is equal to $(p^r-1)/(p-1)$). Thus, we need some methods to weed out undesirable subgroups. First, we must consider only subgroups of sufficiently small index. More precisely, if we want number fields L of absolute degree at most equal to N, then we must have $[Cl_{\mathfrak{m}} : \overline{C}] = h_{\mathfrak{m},C} \le N/n_K$. Second, we must consider only subgroups C such that $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$. In particular, in view of Proposition 3.3.12, we may restrict to moduli $\mathfrak{m}$ such that the conductor of $(\mathfrak{m}, P_{\mathfrak{m}})$ is equal to $\mathfrak{m}$ — in other words, such that $h_{\mathfrak{m}/v} < h_{\mathfrak{m}}$ for all v dividing $\mathfrak{m}$. We may of course use the full force of that proposition, however, and assert that if the conductor of $(\mathfrak{m}, C_2)$ is not equal to $\mathfrak{m}$, then the conductor of $(\mathfrak{m}, C_1)$ will not be equal to $\mathfrak{m}$ for $C_2 \subset C_1$; hence it is not necessary to consider such subgroups.

(5) There is a final easy restriction useful for weeding out unnecessary congruence subgroups. Assume that $(\mathfrak{m}, C_1)$ and $(\mathfrak{m}, C_2)$ are two congruence subgroups modulo the same modulus $\mathfrak{m}$ and that $C_2 \subset C_1$. For $i = 1$ and $i = 2$, denote by L_i the field extension corresponding to these congruence subgroups, and set $n_i = [L_i : \mathbb{Q}]$. Then L_2 is an extension of L_1, hence

$$|d(L_2)| = |d(L_1)|^{[L_2:L_1]}\,\mathcal{N}_{L_1/\mathbb{Q}}(\mathfrak{d}(L_2/L_1)) \ge |d(L_1)|^{[L_2:L_1]} \ ;$$

in other words,

$$|d(L_2)|^{1/n_2} \ge |d(L_1)|^{1/n_1} \ .$$

This gives a nontrivial lower bound for the root discriminant of the field L_2, which may eliminate a priori the subgroup C_2 of C_1 if this lower bound is larger than the desired bound for the root discriminant of L_2.

We can then ask if it is plausible to find all the Abelian extensions of the number fields of degree less than or equal to 7 which are in our tables, satisfying some limitations on the degree and discriminant (for example, degree

up to 100 and root discriminant up to 1.2 times the GRH bound). While not absolutely impossible, this seems like a huge amount of computation. Thus, instead of doing a complete search, we can limit the size of $\mathcal{N}(\mathfrak{m})$ to a much lower bound than the one given by Lemma 9.2.1, and we can also limit the number of congruence subgroups, for instance by always choosing $C = P_{\mathfrak{m}}$.

In any case, this produces large tables of number fields, and we also obtain in this way number fields with root discriminant very close to the GRH bound (see Appendix C and in particular Section 12.2). Note that this method was already used by J. Martinet in 1980 (see [Mart2]) using a pocket calculator, and it is remarkable that many of his records still hold (of course, it may be that they are best possible).

9.2.2 Relative Quadratic Extensions

Let K be a fixed base field of signature (r_1, r_2) and degree $r_1 + 2r_2$. We would like to make a table of quadratic extensions L/K of signature (R_1, R_2) such that the norm of the relative discriminant $\mathfrak{d}(L/K)$ is less than or equal to a given bound B, up to K-isomorphism. (Note that even though two such quadratic extensions will not be K-isomorphic, they may be $\mathbb{Q}$-isomorphic; see Exercise 2.) Using the formula relating absolute and relative discriminants (Theorem 2.5.1), this is equivalent to asking that $|d(L)| \leq d(K)^2 B$. Note also that (R_1, R_2) must satisfy the necessary and sufficient conditions $R_1 + 2R_2 = 2(r_1 + 2r_2)$, $R_1 \leq 2r_1$ given by Corollary 2.2.7 (the condition $2 \mid R_1$ is superfluous here).

One method is to imitate what is done in the absolute case $K = \mathbb{Q}$. In that case $L = \mathbb{Q}(\sqrt{D})$ and D can be chosen either squarefree or, perhaps more canonically, equal to the discriminant of the number field we are looking for. Making a table of such number fields then essentially amounts to making a table of squarefree integers, which is easily done.

In the general case, we may still write $L = K(\sqrt{D})$ with $D \in K^* \setminus K^{*2}$, but this time the choice of D is not so clear. By Kummer theory (which in this case is trivial), we know that $L_1 = K(\sqrt{D_1})$ and $L_2 = K(\sqrt{D_2})$ will be K-isomorphic if and only if $D_2/D_1 \in K^{*2}$. We must then solve two problems. First, find a reasonable representative of the classes of K^*/K^{*2} (for $K = \mathbb{Q}$, we took either the squarefree numbers or the fundamental discriminants); and second, compute the discriminant ideal of $K(\sqrt{D})$.

For the first task, we may use the following lemma.

Lemma 9.2.2. *As usual, denote by $V_2(K)$ the group of 2-virtual units of K (see Definition 5.2.4), and let I_s be the group of squarefree integral ideals of K whose ideal class is a square. Then*

(1) *The map ϕ, which sends the class modulo $V_2(K)$ of an element $\alpha \in K^*$ to the squarefree part of the ideal $\alpha\mathbb{Z}_K$, is an isomorphism from $K^*/V_2(K)$ to I_s.*

(2) *There is a natural exact sequence*

$$1 \longrightarrow \frac{V_2(K)}{K^{*2}} \longrightarrow \frac{K^*}{K^{*2}} \longrightarrow I_s \longrightarrow 1 \ .$$

Note that the group operation on I_s is defined by taking the squarefree part of the ideal product.

Proof. (1). We first show that the map ϕ is well-defined. Indeed, we can write in a unique way $\alpha\mathbb{Z}_K = \mathfrak{a}\mathfrak{q}^2$ with $\mathfrak{a}$ integral and squarefree, and the ideal class of $\mathfrak{a}$ is equal to that of $\mathfrak{q}^{-2}$, so is a square, so $\phi(\alpha) \in I_s$. In addition, if $\alpha' = \alpha\beta$ with $\beta \in V_2(K)$ is equivalent to α modulo $V_2(K)$, then since β is a virtual unit there exists an ideal $\mathfrak{b}$ such that $\beta\mathbb{Z}_K = \mathfrak{b}^2$, so that $\alpha'\mathbb{Z}_K = \mathfrak{a}(\mathfrak{q}\mathfrak{b})^2$, hence $\mathfrak{a}$ is also equal to the image of α'.

In addition, ϕ is injective since $\phi(\alpha) = \phi(\alpha')$ implies that $\alpha/\alpha' = (\mathfrak{q}/\mathfrak{q}')^2$ for some ideals $\mathfrak{q}$ and $\mathfrak{q}'$, which means exactly that $\alpha/\alpha' \in V_2(K)$.

Finally, ϕ is surjective since if $\mathfrak{a} \in I_s$ then $\mathfrak{a} = \alpha\mathfrak{b}^2$ for some ideal $\mathfrak{b}$, so $\alpha\mathbb{Z}_K = \mathfrak{a}\mathfrak{q}^2$ with $\mathfrak{q} = \mathfrak{b}^{-1}$, finishing the proof of (1).

(2) follows from (1) and the natural exact sequence

$$1 \longrightarrow \frac{V_2(K)}{K^{*2}} \longrightarrow \frac{K^*}{K^{*2}} \longrightarrow \frac{K^*}{V_2(K)} \longrightarrow 1 \ .$$

□

It follows from this lemma that we can specify elements of K^*/K^{*2} as follows. Let $(v_1, \ldots, v_{r_c+r_u+1})$ be a fundamental system of 2-virtual units as in Definition 5.2.7. Then we choose an ideal $\mathfrak{a} \in I_s$, we take any ideal $\mathfrak{q}$ such that $\mathfrak{a}\mathfrak{q}^2 = \alpha_0\mathbb{Z}_K$ is a principal ideal, and the $2^{r_c+r_u+1}$ elements of K^*/K^{*2} corresponding to $\mathfrak{a}$ are the classes of elements of the form $\alpha_0 \prod_i v_i^{x_i}$ with $x_i = 0$ or 1.

In addition, if $\mathfrak{q}$ is chosen to be an integral ideal and if $\alpha\mathbb{Z}_K = \mathfrak{a}\mathfrak{q}^2$ and $L = K(\sqrt{\alpha})$, it follows from Hecke's Theorem 10.2.9 (and it is easy to prove directly) that we have $\mathfrak{d}(L/K) = \mathfrak{a}\mathfrak{b}^2$, where $\mathfrak{b}$ is an integral ideal whose prime factors are only prime ideals above 2. In particular, $\mathcal{N}(\mathfrak{a}) \leq \mathcal{N}(\mathfrak{d}(L/K))$. This suggests the following algorithm for computing quadratic extensions.

Algorithm 9.2.3 (List of Relative Quadratic Extensions Using Squarefree Ideals). Given a number field K of signature (r_1, r_2), a signature (R_1, R_2) such that $R_1 + 2R_2 = 2(r_1 + 2r_2)$ and $R_1 \leq 2r_1$, and a bound B, this algorithm determines all relative quadratic extensions L/K up to K-isomorphism such that $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \leq B$ (or, equivalently, $|d(L)| \leq d(K)^2 B$) and such that the signature of L is (R_1, R_2). We assume that $Cl(K) = (A, D_A)$ and $U(K)$ have been computed as well as the necessary data for using the principal ideal algorithm in K.

1. [Initial computations in K] Let $D_A = \mathrm{diag}(a_1, \ldots, a_s)$ and $A = (\overline{\mathfrak{a}_1}, \ldots, \overline{\mathfrak{a}_s})$ with $a_{i+1} \mid a_i$ for $i \leq s-1$, let $(\varepsilon_1, \ldots, \varepsilon_{r_u})$ be a system of fundamental units

and let ε_0 be a generator of the roots of unity in K. Furthermore, let r_c be the 2-rank of the class group, in other words, the largest index i (possibly 0) such that a_i is even. Using the principal ideal algorithm in K, compute v_i such that $\mathfrak{a}_i^{a_i} = v_i\mathbb{Z}_K$ for $1 \le i \le r_c$, and set $r_v \leftarrow r_c + r_u + 1$ and $v_i \leftarrow \varepsilon_{i-r_c-1}$ for $r_c + 1 \le i \le r_v$. Finally, for all $i \le r_v$, let S_i be the r_1-component vector of signatures of v_i expressed as elements of $\mathbb{F}_2$.

2. [Compute Selmer group] Construct two lists $\mathcal{V}$ and $\mathcal{S}$ as follows. For $j = 0, \ldots, 2^{r_v} - 1$, let $j = \sum_{1\le i\le r_v} x_i 2^{i-1}$ be the binary decomposition of j and set $\mathcal{V}_j \leftarrow \prod_{x_i=1} v_i$ and $\mathcal{S}_j \leftarrow \sum_{x_i=1} S_i$ ($\mathcal{V}$ will contain representatives of all elements of the Selmer group $V_2(K)/K^{*2}$, and $\mathcal{S}$ the corresponding signatures).

3. [Compute ideal list] Using Algorithm 2.3.24, compute the list $\mathcal{L}$ of all square-free ideals $\mathfrak{a}$ of norm less than or equal to B, and let $k \leftarrow 0$ (k will be a pointer on the list $\mathcal{L}$).

4. [Next $\mathfrak{a}$] Set $k \leftarrow k + 1$. If $k > |\mathcal{L}|$, terminate the algorithm. Otherwise, let $\mathfrak{a}$ be the kth element of the list $\mathcal{L}$. If $\mathfrak{a} = \mathbb{Z}_K$, set $\alpha_0 = 1$ and $j \leftarrow 0$, let S be the r_1-component vector with entries equal to $0 \in \mathbb{F}_2$, and go to step 8.

5. [Is $\overline{\mathfrak{a}}$ a square?] Using the principal ideal algorithm in K, compute $\beta \in K$ and exponents x_i such that $0 \le x_i < a_i$ and

$$\mathfrak{a} = \beta \prod_{1\le i\le s} \mathfrak{a}_i^{x_i} .$$

If for some $i \le r_c$ the exponent x_i is odd, go to step 4. Otherwise, for each i such that $r_c < i \le s$ and x_i is odd, set $x_i \leftarrow x_i - a_i$.

6. [Compute $\mathfrak{q}$] Compute

$$\mathfrak{q}_0 \leftarrow \prod_{1\le i\le s,\ x_i\neq 0} \mathfrak{a}_i^{(a_i-x_i)/2} .$$

Using Algorithm 4.3.4 with $\mathfrak{m} = \mathbb{Z}_K$, let $\mathfrak{q}$ be an almost-reduced integral ideal in the same ideal class as $\mathfrak{q}_0$.

7. [Compute α_0] Using the principal ideal algorithm, compute an α_0 such that $\mathfrak{a}\mathfrak{q}^2 = \alpha_0\mathbb{Z}_K$, let S be the r_1-component vector with entries in $\mathbb{F}_2$ of the signatures of α_0, and set $j \leftarrow -1$.

8. [Loop through virtual units] Set $j \leftarrow j+1$. If $j \ge 2^{r_v}$, go to step 4. Otherwise, set $T \leftarrow S + \mathcal{S}_j$. If the number of 0s among the entries of T is not equal to $R_1/2$, go to step 8.

9. [Extension suitable?] Set $\alpha \leftarrow \alpha_0\mathcal{V}_j$. Using Algorithm 2.4.9 or Hecke's Theorem 10.2.9, compute the ideal discriminant $\mathfrak{d}$ of the relative quadratic extension of K defined by the polynomial $X^2 - \alpha$. If $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}) \le B$, output the relative extension $K(\sqrt{\alpha})/K$. Go to step 8.

The proof of this algorithm's validity is left to the reader (Exercise 3).

Remarks

(1) It is not difficult to generalize this algorithm to the construction of cyclic extensions of K of degree ℓ for a prime number ℓ, when one assumes that $\zeta_\ell \in K$ (in other words, for the case of a Kummer extension); see Exercise 4.

(2) As already mentioned, we have $\mathcal{N}(\mathfrak{a}) \leq \mathcal{N}(\mathfrak{d}(L/K)) \leq 4^n \mathcal{N}(\mathfrak{a})$. To be sure to have all possible quadratic extensions, we must therefore use all squarefree ideals of norm up to B. Since 4^n can become quite large, there is therefore a considerable amount of waste in the above algorithm, since ideals $\mathfrak{a}$ such that $\mathcal{N}(\mathfrak{a})$ is much larger than $B/4^n$ will contribute a very small number of extensions. This is the price to pay for simplicity. The algorithm presented below does not have this disadvantage, but the computations are considerably longer.

A different way of computing quadratic extensions of a base field K is to use class field theory. Since quadratic extensions of number fields are necessarily Abelian, we can simply make a table of possible conductors $\mathfrak{m}$ and congruence subgroups C. Recall that in this very simple case, if $\mathfrak{m}$ is the conductor of the quadratic extension L/K, then its finite part is equal to the relative discriminant ideal (see Corollary 3.5.12). In addition, the congruence subgroup C defines a quadratic extension of K if and only if it is of index 2 in $I_{\mathfrak{m}}(K)$. Thus, we can use the following algorithm.

Algorithm 9.2.4 (List of Relative Quadratic Extensions). Given a number field K of signature (r_1, r_2), a signature (R_1, R_2) such that $R_1 + 2R_2 = 2(r_1 + 2r_2)$ and $R_1 \leq 2r_1$, and a bound B, this algorithm determines all relative quadratic extensions L/K up to K-isomorphism such that $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \leq B$ (or, equivalently, $|d(L)| \leq d(K)^2 B$) and such that the signature of L is (R_1, R_2).

1. [Compute ideal list] Using Algorithm 2.3.25 with $\ell = 2$, compute the list $\mathcal{L}_0$ of all ideals $\mathfrak{m}_0$ of norm less than or equal to B which are conductors at 2 — in other words, such that for all prime ideals $\mathfrak{p}$ dividing $\mathfrak{m}_0$, $v_{\mathfrak{p}}(\mathfrak{m}_0) = 1$ if $\mathfrak{p} \nmid 2$ while $2 \leq v_{\mathfrak{p}}(\mathfrak{m}_0) \leq 2e(\mathfrak{p}/2)+1$ if $\mathfrak{p} \mid 2$ — and set $k \leftarrow R_2 - 2r_2 = r_1 - R_1/2$.

2. [Compute moduli] Let $\mathcal{L}$ be the list of all moduli of the form $\mathfrak{m} \leftarrow \mathfrak{m}_0 \mathfrak{m}_\infty$, where $\mathfrak{m}_0 \in \mathcal{L}_0$ and $\mathfrak{m}_\infty$ ranges through either all the $\binom{r_1}{k}$ subsets of the real places of K of cardinality equal to k if $K/\mathbb{Q}$ is not Galois, or a single fixed such subset if $K/\mathbb{Q}$ is Galois, and set $i \leftarrow 0$ (i will be a pointer to the list $\mathcal{L}$).

3. [Check if $\mathfrak{m}$ is a conductor] Set $i \leftarrow i + 1$. If $i > |\mathcal{L}|$, terminate the algorithm. Otherwise, let $\mathfrak{m}$ be the ith element of the list $\mathcal{L}$. Using Algorithm 4.4.2, check whether $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, P_{\mathfrak{m}})$. If it is not, or if we find in that algorithm that $h_{\mathfrak{m}}$ is odd, go to step 3.

4. [Compute congruence subgroups] Using Algorithm 4.3.1, compute the SNF (A, D_A) of $Cl_{\mathfrak{m}}(K)$. Then using Algorithm 4.1.20 with $\ell = 2$, compute the list $\mathcal{C}$ of the left HNF divisors of D_A corresponding to all congruence subgroups C of $I_{\mathfrak{m}}(K)$ of index 2, and set $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{C}$).

5. [Check if $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$] Set $j \leftarrow j+1$. If $j > |\mathcal{C}|$, go to step 3. Otherwise, let C be the congruence subgroup corresponding to the jth element of $\mathcal{C}$. Using once again Algorithm 4.4.2, check whether $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$. If it is not, go to step 5.
6. [Compute defining polynomial] (We now know that $(\mathfrak{m}, C)$ is the conductor of a suitable quadratic extension of K.) Using Algorithm 5.4.8 for $n = 2$, output the defining polynomial for the quadratic extension L/K corresponding to $(\mathfrak{m}, C)$ and go to step 5.

Proof. The proof of this algorithm's validity is straightforward once one notices that the conditions on $\mathfrak{m}_0$ in step 1 are consequences of Corollary 3.5.12 and of Hecke's theorem (more precisely, of Theorem 5.2.2). □

Remark. Steps 4, 5, and 6 may be replaced by the following single step.

4'. [Find list of polynomials] Using Algorithm 5.2.14 with $\ell = 2$, output the relative defining polynomials of all the quadratic extensions of K of conductor $\mathfrak{m}$, and go to step 3.

The reason for this is that, contrary to Algorithm 5.4.8, which specifically finds the defining polynomial corresponding to a given congruence subgroup $(\mathfrak{m}, C)$, Algorithm 5.2.14 essentially uses only the information about $\mathfrak{m}$ and the fact that $h_{\mathfrak{m},C} = 2$, and so it gives all the polynomials at once. Both methods should be implemented in order to compare their relative efficiency, but it is quite clear that the modification explained here will be superior if there are many quadratic extensions having a given conductor.

The two algorithms presented above are reasonable methods to compute a list of relative quadratic extensions. In the author's and collaborators' implementations, the first algorithm using squarefree ideals is usually faster. In addition, it enables us to give an asymptotic estimate for the *number* of relative quadratic extensions (see Exercise 6). On the other hand, the class field method does lead to a formula for the number of such extensions (see Exercise 5), but since ray class numbers are involved in this formula, it does not seem easy to deduce from it the result of Exercise 6.

9.2.3 Relative Cubic Extensions

Let K be a fixed base field of signature (r_1, r_2) and degree $r_1 + 2r_2$. Similarly to the quadratic case, we would like to make a table of cubic extensions L/K of signature (R_1, R_2) such that the norm of the relative discriminant $\mathfrak{d}(L/K)$ is less than or equal to a given bound B, up to K-isomorphism. Using the formula relating absolute and relative discriminants (Theorem 2.5.1), this is equivalent to asking that $|d(L)| \leq |d(K)|^3 B$. Note also that (R_1, R_2) must satisfy the necessary and sufficient conditions $R_1 + 2R_2 = 3(r_1 + 2r_2)$ and $r_1 \leq R_1 \leq 3r_1$ given by Corollary 2.2.7.

We explored in detail in Chapter 8 how to compute extensive tables of cubic extensions of $\mathbb{Q}$. Considering [Dat-Wri], it is quite plausible that these methods can be generalized to the relative case as well, but to the author's knowledge this has not been done.

We would first like to make a table of *cyclic* cubic extensions. Since 3 is odd, Corollary 2.2.6 tells us that the real places are unramified — in other words, that we have $R_1 = 3r_1$ and $R_2 = 3r_2$ — and Corollary 3.5.12 tells us that $\mathfrak{d}(L/K) = \mathfrak{m}^2$, where $\mathfrak{m}$ is the conductor the extension. As in the quadratic case, we have two different methods. The first method is the analog of the use of squarefree ideals and is left as an excellent exercise for the reader (Exercise 7). The second method uses class field theory and gives the following algorithm, which is very close to Algorithm 9.2.4.

Algorithm 9.2.5 (List of Relative Cyclic Cubic Extensions). Given a number field K of signature (r_1, r_2) and a bound B, this algorithm determines all relative cyclic cubic extensions L/K up to K-isomorphism such that $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \leq B$ (or, equivalently, $|d(L)| \leq |d(K)|^3 B$). The signature of L will necessarily be equal to $(3r_1, 3r_2)$.

1. [Compute ideal list] Using Algorithm 2.3.25 with $\ell = 3$, compute the list $\mathcal{L}$ of all ideals $\mathfrak{m}$ of norm less than or equal to $B^{1/2}$ which are conductors at 3 — in other words, such that for all prime ideals $\mathfrak{p}$ dividing $\mathfrak{m}$, $v_{\mathfrak{p}}(\mathfrak{m}) = 1$ if $\mathfrak{p} \nmid 3$ while $2 \leq v_{\mathfrak{p}}(\mathfrak{m}) \leq \lfloor 3e(\mathfrak{p}/3)/2 \rfloor + 1$ if $\mathfrak{p} \mid 3$ — and set $i \leftarrow 0$ (i will be a pointer to the list $\mathcal{L}$).
2. [Check if $\mathfrak{m}$ is a conductor] Set $i \leftarrow i + 1$. If $i > |\mathcal{L}|$, terminate the algorithm. Otherwise, let $\mathfrak{m}$ be the ith element of the list $\mathcal{L}$. Using Algorithm 4.4.2, check whether $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, P_{\mathfrak{m}})$. If it is not, or if we find in that algorithm that $h_{\mathfrak{m}}$ is not divisible by 3, go to step 2.
3. [Compute congruence subgroups] Using Algorithm 4.3.1, compute the SNF (B, D_B) of $Cl_{\mathfrak{m}}(K)$. Then using Algorithm 4.1.20 with $\ell = 3$, compute the list $\mathcal{C}$ of the left HNF divisors of D_B corresponding to all congruence subgroups C of $I_{\mathfrak{m}}(K)$ of index 3, and set $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{C}$).
4. [Check if $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$] Set $j \leftarrow j + 1$. If $j > |\mathcal{C}|$, go to step 2. Otherwise, let C be the congruence subgroup corresponding to the jth element of $\mathcal{C}$. Once again using Algorithm 4.4.2, check whether $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$. If it is not, go to step 4.
5. [Compute defining polynomial] (We now know that $(\mathfrak{m}, C)$ is the conductor of a suitable cyclic cubic extension of K.) Using Algorithm 5.5.5 for $n = 3$ or Algorithm 5.3.17, output the defining polynomial for the cubic extension L/K corresponding to $(\mathfrak{m}, C)$ and go to step 4.

Remarks

(1) In the relative quadratic case, we used the upper bound $v_{\mathfrak{p}}(\mathfrak{m}) \leq 2e(\mathfrak{p}/2) + 1$ coming from Hecke's Theorem 10.2.9, which assumed $\zeta_\ell \in K$. Here in step 1, we use the more general bound

$$v_{\mathfrak{p}}(\mathfrak{m}) \leq \left\lfloor \frac{\ell e(\mathfrak{p}/\ell)}{\ell - 1} \right\rfloor + 1$$

already stated in Proposition 3.3.22 (see Corollary 10.1.24).

(2) Since L/K is unramified at infinity, contrary to Algorithm 9.2.4 we do not have to worry about infinite places.

(3) Since Algorithm 5.3.17 requires the specific knowledge of the congruence subgroup C to compute the exact conductor of the extension L_z/K_z, contrary to Algorithm 9.2.4 it is in general not possible to replace steps 3, 4, and 5 by a unique step. When $\mathfrak{m}$ is prime to 3, however, and also in many other cases, this is possible. We leave the details to the reader.

(4) If $\zeta_3 \notin K$, the computation of the defining polynomials in step 5 will be more expensive since they will involve first extending the extension L/K to the extension L_z/K_z with $K_z = K(\zeta_3)$. In this case, it is evidently important to precompute all the necessary information about the field K_z so as not to repeat it each time. On the other hand, if the defining polynomial is not desired but only the relative discriminant (which is equal to $\mathfrak{m}^2$) is desired, this is, of course, not necessary.

Once the cyclic cubic extensions have been listed, we want to list the noncyclic (or, equivalently, the non-Galois) extensions. Even in this case, we can use class field theory thanks to the following theorem kindly communicated to us by J. Martinet. The proof involves a refinement of the Hasse–Arf theorem (see [Ser]). I refer to Section 10.1.5 for the case of prime degree and to an unpublished manuscript of J. Martinet for the general case.

Theorem 9.2.6. *Let K be a number field, L an extension of K of degree n, and let L_2 be the Galois closure of L/K in some algebraic closure of K. Assume that $\mathrm{Gal}(L_2/K)$ is isomorphic to the dihedral group D_n, and that $n \geq 3$ with n odd. Finally, let K_2 be the unique quadratic subextension of L_2/K, and let $(\mathfrak{m}, C)$ be the conductor of the Abelian extension L_2/K_2 (see the diagram below). Then for each $d \mid n$ there exists an integral ideal $\mathfrak{a}_d$ of K (not only of K_2) such that the following holds.*

(1) *The conductor $\mathfrak{m}$ of L_2/K_2 is obtained by extending $\mathfrak{a}_n$ to K_2, in other words, $\mathfrak{m}_0 = \mathfrak{a}_n \mathbb{Z}_{K_2}$ and $\mathfrak{m}_\infty = \varnothing$.*

(2) *More generally, the conductor of the unique subextension $L_{2,d}/K_2$ of degree d of L_2/K_2 is obtained by extending $\mathfrak{a}_d$ to K_2. In particular, $\mathfrak{a}_1 = \mathbb{Z}_K$ and $d_1 \mid d_2 \mid n$ implies that $\mathfrak{a}_{d_1} \mid \mathfrak{a}_{d_2}$.*

(3) *If τ is a generator of $\mathrm{Gal}(K_2/K)$ and if we set*

$$Cl_{\mathfrak{m}}(K_2)/\overline{C} = \langle \overline{I} \rangle$$

for some ideal I of K_2, then $\tau(\overline{I}) = \overline{I^{-1}}$.

(4) *The relative discriminant ideal $\mathfrak{d}(L/K)$ is given by*

$$\mathfrak{d}(L/K) = \mathfrak{d}(K_2/K)^{(n-1)/2} \prod_{d \mid n} \mathfrak{a}_d^{\phi(d)} \ .$$

(5) *Let* (r_1, r_2) *be the signature of* K *and* (r_1', r_2') *the signature of* K_2. *Then the signature of* L_2 *is equal to* (nr_1', nr_2') *and the signature of* L *is equal to* $(r_1 + ((n-1)/2)r_1', r_2 + ((n-1)/2)r_2')$.

(6) *Let* $\mathfrak{p}$ *be a prime ideal of* K. *Then we have the following properties:*
   a) $\mathfrak{p} \mid \mathfrak{a}_n$ *if and only if* $\mathfrak{p}$ *is totally ramified in* L/K;
   b) $\mathfrak{p} \mid (\mathfrak{a}_d, \mathfrak{d}(K_2/K))$ *implies* $\mathfrak{p} \mid d$;
   c) $\mathfrak{p}^2 \mid \mathfrak{a}_d$ *implies* $\mathfrak{p} \mid d$.

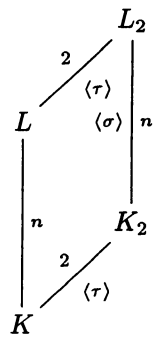

See Section 10.1.5 for a proof of this theorem in the case n prime.

Remarks

(1) The point of this theorem is not only that it gives exactly the relative discriminant ideal of L/K, but also that one can obtain L as a subfield of an Abelian extension of K_2 of conductor coming from K and not only from K_2. Ordinary results of Galois theory would tell us only that the conductor of L_2/K_2 must be stable under $\mathrm{Gal}(K_2/K)$, but not necessarily that it comes from a modulus of K.

(2) A similar but slightly more complicated statement is true when n is even. In particular, in that case, if K_2' denotes the unique quadratic subextension of L/K, we have

$$\mathfrak{d}(L/K) = \mathfrak{d}(K_2/K)^{(n-2)/2}\mathfrak{d}(K_2'/K) \prod_{d \mid n,\ d>2} \mathfrak{a}_d^{\phi(d)} .$$

(3) If we assume in addition that n is an odd prime number, the formula for the discriminant simplifies to

$$\mathfrak{d}(L/K) = \mathfrak{d}(K_2/K)^{(n-1)/2}\mathfrak{a}^{n-1} ,$$

where $\mathfrak{a}$ is the ideal of K such that $\mathfrak{a}\mathbb{Z}_{K_2}$ is the conductor of L_2/K_2.

(4) If $n = 3$, statement (6) is the exact generalization to the relative case of Proposition 8.4.1.

(5) The corresponding statement in the more general case of *metacyclic* groups (where the condition $\tau^2 = 1$ is replaced by $\tau^k = 1$ for some $k > 2$) is false.

Using this theorem with $n = 3$ and the fact that a noncyclic cubic extension has Galois group isomorphic to D_3, it is easy to write an algorithm for computing relative noncyclic cubic extensions as follows.

Algorithm 9.2.7 (List of Relative Noncyclic Cubic Extensions). Given a number field K of signature (r_1, r_2), a signature (R_1, R_2) such that $R_1 + 2R_2 = 3(r_1 + 2r_2)$ and $r_1 \leq R_1 \leq 3r_1$, and a bound B, this algorithm determines all non-Galois relative cubic extensions L/K up to K-isomorphism such that $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K)) \leq B$ (or, equivalently, $|d(L)| \leq |d(K)|^3 B$) and such that the signature of L is (R_1, R_2).

1. [Compute ideal list] Using a simple modification of Algorithm 2.3.25, compute the list $\mathcal{L}_0$ of all ideals $\mathfrak{a}$ of K of norm less than or equal to $B^{1/2}$ which are conductors at 3, except that we allow $v_{\mathfrak{p}}(\mathfrak{a}) = 1$ for $\mathfrak{p} \mid 3$ (see Exercise 8).
2. [Compute list of quadratic extensions] Using either Algorithm 9.2.3 or 9.2.4, compute the list $\mathcal{Q}$ of quadratic extensions K_2/K up to K-isomorphism of signature $(R_1 - r_1, R_2 - r_2)$ such that $\mathcal{N}(\mathfrak{d}(K_2/K)) \leq B$, and set $i \leftarrow 0$ (i will be a pointer to the list $\mathcal{Q}$).
3. [Compute quadratic extension] Set $i \leftarrow i + 1$. If $i > |\mathcal{Q}|$, terminate the algorithm. Otherwise, let K_2 be the ith element of $\mathcal{Q}$. Compute the necessary data to work in K_2, such as integral basis, class and unit group, and compute the action of the generator τ of $\mathrm{Gal}(K_2/K)$ on a K-basis of K_2/K. Finally, let $\mathfrak{d} \leftarrow \mathfrak{d}(K_2/K)$ be the relative discriminant of K_2/K.
4. [Compute ideal sublist] Let $\mathcal{L}_1$ be the sublist of the ideals $\mathfrak{a} \in \mathcal{L}_0$ such that

$$\mathcal{N}(\mathfrak{a}) \leq (B/\mathcal{N}(\mathfrak{d}(K_2/K)))^{1/2} = |d(K)|\,(B/|d(K_2)|)^{1/2}$$

   and such that if $\mathfrak{p} \nmid 3$ and $\mathfrak{p} \mid \mathfrak{a}$, then $\mathfrak{p} \nmid \mathfrak{d}$; if $\mathfrak{p} \mid 3$ and $v_{\mathfrak{p}}(\mathfrak{a}) = 1$, then $\mathfrak{p} \mid \mathfrak{d}$; and finally if $\mathfrak{p} \mid 3$ and $v_{\mathfrak{p}}(\mathfrak{a}) = 3e(\mathfrak{p}/3)/2 + 1$, then $\mathfrak{p} \nmid \mathfrak{d}$. Using Algorithm 2.5.4, compute the list $\mathcal{L}_2$ of ideals $\mathfrak{m}_0$ of K_2 of the form $\mathfrak{m}_0 = \mathfrak{a}\mathbb{Z}_{K_2}$ with $\mathfrak{a} \in \mathcal{L}_1$, and set $j \leftarrow 0$ (j will be a pointer to the list $\mathcal{L}_2$).
5. [Check if $\mathfrak{m}$ is a conductor] Set $j \leftarrow j+1$. If $j > |\mathcal{L}_2|$, terminate the algorithm. Otherwise, let $\mathfrak{m}$ be the modulus whose finite part $\mathfrak{m}_0$ is the jth element of the list $\mathcal{L}_2$ and with $\mathfrak{m}_\infty = \varnothing$. Using Algorithm 4.4.2, check whether $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, P_{\mathfrak{m}})$. If it is not, or if we find in that algorithm that $h_{\mathfrak{m}}$ is not divisible by 3, go to step 5.
6. [Compute congruence subgroups] Using Algorithm 4.3.1, compute the SNF (A, D_A) of $Cl_{\mathfrak{m}}(K_2)$. Then using Algorithm 4.1.20 with $\ell = 3$, compute the list $\mathcal{C}$ of the left HNF divisors of D_A corresponding to all congruence subgroups C of $I_{\mathfrak{m}}(K_2)$ of index 3, and set $c \leftarrow 0$ (c will be a pointer to the list $\mathcal{C}$).
7. [Check if $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$] Set $c \leftarrow c + 1$. If $c > |\mathcal{C}|$, go to step 5. Otherwise, let C be the congruence subgroup corresponding to the cth element of $\mathcal{C}$. Once again using Algorithm 4.4.2, check whether $\mathfrak{m}$ is the conductor of $(\mathfrak{m}, C)$. If it is not, go to step 7.

8. [Check condition on τ] Let H_C be the left HNF divisor of D_A corresponding to C, let k be the index of the unique row whose diagonal entry is equal to 3, and let $\overline{I}$ be the kth generator of $Cl_{\mathfrak{m}}(K_2)$ — in other words, the kth element of A. Using the explicit action of τ, compute $J \leftarrow \tau(I)$, and compute the discrete logarithm Z of $\overline{J}$ in $Cl_{\mathfrak{m}}(K_2)$. If the kth component Z_k of Z is not congruent to 2 modulo 3, go to step 7.
9. [Compute defining polynomial for L_2/K_2] (We now know that $(\mathfrak{m}, C)$ is the conductor of a suitable cyclic cubic extension of K_2.) Using Algorithm 5.4.8 for $n = 3$, or Algorithm 5.2.14 when $\zeta_3 \in K_2$, or Algorithm 5.5.5 for $n = 3$, or Algorithm 5.3.17 when $\zeta_3 \notin K_2$, compute a defining polynomial $P(X) \in K_2[X]$ for the cubic extension L_2/K_2 corresponding to $(\mathfrak{m}, C)$.
10. [Compute defining polynomial for L/K] Let $P_c \leftarrow P^{\tau}(X)$ be the polynomial obtained by applying τ to all the coefficients of $P \in K_2[X]$, and set $Q(X) \leftarrow \mathcal{R}_Y(P_c(Y), P(X - Y))$, where as usual $\mathcal{R}_Y$ denotes the resultant in the variable Y. Then $Q(X) \in K[X]$. Factor $Q(X)$ in $K[X]$, output one of the irreducible factors of $Q(X)$ of degree 3 in $K[X]$ (it will have one) as a defining polynomial for L/K, and go to step 7.

Proof. Thanks to Theorem 9.2.6, we know that the conductor of the cyclic cubic extension L_2/K_2 is of the form $\mathfrak{a}\mathbb{Z}_{K_2}$ with $\mathfrak{a}$ an ideal of K and that $\mathfrak{d}(L/K) = \mathfrak{d}(K_2/K)\mathfrak{a}^2$, so that $\mathcal{N}(\mathfrak{d}(L/K)) \le B$ is equivalent to $\mathcal{N}(\mathfrak{a}) \le (B/\mathcal{N}(\mathfrak{d}(K_2/K)))^{1/2}$. In addition, by Propositions 3.3.21 and 3.3.22, for every prime ideal $\mathfrak{P}$ of K_2 dividing $\mathfrak{a}\mathbb{Z}_{K_2}$ we must have $v_{\mathfrak{P}}(\mathfrak{a}\mathbb{Z}_{K_2}) = 1$ if $\mathfrak{P} \nmid 3$, while $2 \le v_{\mathfrak{P}}(\mathfrak{a}\mathbb{Z}_{K_2}) \le \lfloor 3e(\mathfrak{P}/3)/2 \rfloor + 1$ if $\mathfrak{P} \mid 3$. In terms of the ideal $\mathfrak{p}$ below $\mathfrak{P}$ in K, this implies the following when $\mathfrak{p} \mid \mathfrak{a}$.

(1) If $\mathfrak{p} \nmid 3$ and $\mathfrak{p}$ is ramified in K_2/K as $\mathfrak{p}\mathbb{Z}_{K_2} = \mathfrak{P}^2$, then $v_{\mathfrak{P}}(\mathfrak{a}\mathbb{Z}_{K_2}) = 2v_{\mathfrak{p}}(\mathfrak{a}) \ge 2$, which is impossible.
(2) If $\mathfrak{p} \nmid 3$ and $\mathfrak{p}$ is unramified in K_2/K, then we must have $v_{\mathfrak{p}}(\mathfrak{a}) = 1$.
(3) If $\mathfrak{p} \mid 3$ and $\mathfrak{p}$ is ramified in K_2/K as $\mathfrak{p}\mathbb{Z}_{K_2} = \mathfrak{P}^2$, then the condition on $v_{\mathfrak{P}}(\mathfrak{a})$ given above is equivalent to the condition $1 \le v_{\mathfrak{p}}(\mathfrak{a}) \le 3e(\mathfrak{p}/3)/2 + 1/2$.
(4) Finally, if $\mathfrak{p} \mid 3$ and $\mathfrak{p}$ is unramified in K_2/K, then we must have $2 \le v_{\mathfrak{p}}(\mathfrak{a}) \le \lfloor 3e(\mathfrak{p}/3)/2 \rfloor + 1$.

These conditions explain the list computed in step 1 as well as the additional conditions given in step 4 once K_2 has been chosen.

In step 8, the ideal class $\overline{I}$ is clearly a generator of the group $Cl_{\mathfrak{m}}(K_2)/\overline{C}$ of order 3. We will have $Z_k \equiv \pm 1 \pmod 3$, and the condition $Z_k \equiv -1 \equiv 2 \pmod 3$ is equivalent to the condition $\tau(\overline{I}) = \overline{I^{-1}}$ of Theorem 9.2.6.

Finally, in step 10, the roots of $Q(X)$ are of the form $\theta + \tau(\theta')$, where θ and θ' denote roots of $P(X)$ in L_2. Denote by θ_1 some fixed root of $P(X)$ in L_2, and set $\theta_2 = \sigma(\theta_1)$ and $\theta_3 = \sigma(\theta_2)$. If we set

$$Q_1(X) = (X - (\theta_1 + \tau(\theta_1)))(X - (\theta_2 + \tau(\theta_3)))(X - (\theta_3 + \tau(\theta_2))) \ ,$$
$$Q_2(X) = (X - (\theta_2 + \tau(\theta_2)))(X - (\theta_1 + \tau(\theta_3)))(X - (\theta_3 + \tau(\theta_1))) \ , \text{ and}$$
$$Q_3(X) = (X - (\theta_3 + \tau(\theta_3)))(X - (\theta_1 + \tau(\theta_2)))(X - (\theta_2 + \tau(\theta_1))) \ ,$$

then $Q(X) = Q_1(X)Q_2(X)Q_3(X)$, the $Q_i(X)$ are clearly stable by τ, and they are also stable by σ by definition of the θ_i, since $\tau\sigma\tau^{-1} = \sigma^{-1}$. It follows by Galois theory that $Q_i(X) \in K[X]$ for $i = 1$, 2, and 3, so $Q(X) = Q_1(X)Q_2(X)Q_3(X)$ is a (possibly partial) factorization of $Q(X)$ in $K[X]$.

I claim that at least two of the $Q_i(X)$ are irreducible in $K[X]$. Indeed, assume by contradiction that two of them are reducible, say $Q_1(X)$ and $Q_2(X)$. Then $Q_1(X)$ will have a factor of degree 1 in $K[X]$. But since the three roots of $Q_1(X)$ differ only by application of σ, if one of them is in K, then they will all be equal. In other words, if $Q_1(X)$ is reducible, then in fact $Q_1(X) = (X - \alpha_1)^3$, where

$$\alpha_1 = \theta_1 + \tau(\theta_1) = \theta_2 + \tau(\theta_3) = \theta_3 + \tau(\theta_2) \in K \ .$$

Similarly, if $Q_2(X)$ is reducible, then in fact $Q_2(X) = (X - \alpha_2)^3$ with

$$\alpha_2 = \theta_2 + \tau(\theta_2) = \theta_1 + \tau(\theta_3) = \theta_3 + \tau(\theta_1) \in K \ .$$

Subtracting, we deduce that

$$\alpha_1 - \alpha_2 = \beta = \theta_1 - \theta_3 = \theta_2 - \theta_1 = \theta_3 - \theta_2 \in K \ ,$$

and adding these three expressions for β we obtain $3\beta = 0$, so $\theta_3 = \theta_2$, which is absurd. This shows that at least two of the $Q_i(X)$ are irreducible in $K[X]$.

Assume, for example, that $Q_1(X)$ is irreducible in $K[X]$. Then $\alpha_1 = \theta_1 + \tau(\theta_1)$ is a root of $Q_1(X)$, it does not belong to K, and it is stable by τ, hence it belongs to L. Since $[L : K] = 3$ is prime, we deduce that $L = K(\alpha_1)$ and that $Q_1(X)$ is a defining polynomial for L/K, finishing the proof of the algorithm. □

Remarks

(1) I thank F. Diaz y Diaz for the resultant method explained in step 10.

(2) If $\zeta_3 \in K_2$, the computation of the Kummer extension in step 9 will be cheap. On the other hand, if $\zeta_3 \notin K_2$, we will need to extend K_2 by adjoining ζ_3, work on the extension field, and come back down. There seems to be no way to avoid working in $K_2(\zeta_3)$ so as to be able to apply Kummer theory. On the other hand, to obtain a defining polynomial for L/K, it is not necessary first to obtain a defining polynomial for L_2/K_2 and then go down to L/K, since it is possible to perform both steps at once (see Exercise 9).

(3) As usual, if the explicit polynomials are not needed but only the relative discriminant ideals are, we can replace the computations performed in steps 9 and 10 by the simple computation of $\mathfrak{d}(L/K) \leftarrow \mathfrak{d}\mathfrak{a}^2$.

(4) As in the quadratic case, we can modify the preceding algorithm so that it uses squarefree ideals (more precisely, Exercise 7) instead of class field theory in the construction of the extensions L_2/K_2 (see Exercise 10).

9.2.4 Finding the Smallest Discriminants Using Class Field Theory

Finding the smallest possible discriminants for a given signature using class field theory is not possible unless one adds some extra conditions, such as conditions on the Galois group. Indeed, if the extra conditions do not exclude *primitive* fields, then class field theory can only find those as Abelian extensions of $\mathbb{Q}$, hence as subfields of cyclotomic fields by the Kronecker–Weber theorem, and these can all easily be found if the signature and bounds are given and are usually not the smallest possible, except in very small degrees.

Consider the first few degrees.

In degree 4, the imprimitive fields are the fields with Galois group isomorphic to subgroups of D_4, in other words to C_4, $C_2 \times C_2$, and D_4, and of course they all contain a quadratic subfield. Thus, although the Abelian number fields with Galois group C_4 and $C_2 \times C_2$ can be obtained easily as extensions of $\mathbb{Q}$, we may also deal with the three Galois groups at the same time by considering relative quadratic extensions. Since a quadratic extension is necessarily Abelian, the desired number fields can be obtained as class fields from quadratic base fields (see also Section 9.4.5).

Similar ideas apply in degree 6. We use the permutation group notation given, for example, in [Coh0]. A sextic field contains a cubic subfield (of which it will necessarily an Abelian extension) if and only if its Galois group is isomorphic to a transitive subgroup of $S_4 \times C_2$; in other words, to C_6, S_3, D_6, A_4, S_4^+, S_4^-, $A_4 \times C_2$, and $S_4 \times C_2$. Thus, complete tables of number fields having those specific Galois groups can be made by using class field theory together with tables of cubic fields, easily made by using the algorithms of Chapter 8.

We can also compute tables of sextic fields with Galois group isomorphic to G_{18} since these (in addition to C_6 and S_3) are the sextic fields that are Abelian extensions of a quadratic subfield.

Using Algorithm 9.2.7, we can also compute tables of sextic fields with Galois group isomorphic to the transitive subgroups of G_{72} not already considered; in other words, to G_{18}^+, G_{18}^-, G_{36}, and G_{72}, since these are the sextic fields that are non-Galois extensions of a quadratic subfield.

In degree 8, we can treat in the same way octic fields that contain a quartic subfield, which correspond to a large list of Galois groups. We refer to [Co-Di-Ol3] and [Co-Di-Ol6] for details. On the other hand, imprimitive octic fields that contain a quadratic subfield but no quartic subfield cannot be treated by class field theory. (The only nontrivial case would be when the relative Galois group is is isomorphic to D_4, for which there exists a theorem

analogous to Theorem 9.2.6, but unfortunately in this case the octic field contains a quartic subfield since a quartic D_4 extension contains a quadratic subextension.)

Apart from the general methods coming from the geometry of numbers (in this case Martinet's Theorem 9.3.2 for the relative case below), the only known method, which is in fact a vast generalization of class field theory to GL_2, is based on the use of Galois representations of GL_2 (see, for example, [Cas-Jeh]).

These ideas can of course be pushed to larger degrees if desired.

9.3 Using the Geometry of Numbers

The books [Con-Slo] and [Mart5] are the essential modern references for the study of this subject for its own sake. On the other hand, for the construction of complete tables of number fields using the geometry of numbers, we refer to [Mart3] and [Poh] for a modern description of the general methods, to [For1] for totally complex quartic fields, to [Buc-For] for totally real quartic fields, to [Bu-Fo-Po] for quartic fields of mixed signature, to [Diaz] for totally real quintic fields, to [Sc-Po-Di] for nontotally real quintic fields, to [Be-Ma-Ol] for sextic fields with a quadratic subfield, to [Oli1] for sextic fields with a cubic subfield, to [Oli2] for primitive sextic fields, and to [Let] for septic fields.

9.3.1 The General Procedure

For the reader's convenience, we recall here Hunter's theorem (see, for example, [Coh0, Theorem 6.4.2]). Let γ_n denote Hermite's constant in dimension n, whose first few values are given by $\gamma_1 = 1$, $\gamma_2^2 = 4/3$, $\gamma_3^3 = 2$, $\gamma_4^4 = 4$, $\gamma_5^5 = 8$, $\gamma_6^6 = 64/3$, $\gamma_7^7 = 64$, and $\gamma_8^8 = 256$.

Theorem 9.3.1. *Let K be a number field of degree n over $\mathbb{Q}$ and discriminant $d(K)$. There exists $\alpha \in \mathbb{Z}_K \setminus \mathbb{Z}$ that satisfies the following additional properties.*

(1) *If $\alpha^{(j)}$ denotes the conjugates of α in $\mathbb{C}$, then*

$$\sum_{1 \le j \le n} |\alpha^{(j)}|^2 \le \frac{(\mathrm{Tr}_{K/\mathbb{Q}}(\alpha))^2}{n} + \gamma_{n-1} \left(\frac{|d(K)|}{n} \right)^{1/(n-1)} ,$$

where γ_{n-1} is Hermite's constant in dimension $n-1$.

(2) $0 \le \mathrm{Tr}_{K/\mathbb{Q}}(\alpha) \le n/2$.

There also exists a relative version of this theorem, due to J. Martinet, as follows.

Theorem 9.3.2. *Let K be a number field of degree m and let L/K be a relative extension of number fields of relative degree $n = [L : K]$, so that L is of absolute degree nm. There exists $\alpha \in \mathbb{Z}_L \setminus \mathbb{Z}_K$ that satisfies the following additional properties.*

(1) *If $\alpha^{(j)}$ denotes the conjugates of α in $\mathbb{C}$ and $t^{(h)}$ denotes the conjugates of some $t \in K$ in $\mathbb{C}$, then*

$$\sum_{1 \le j \le nm} |\alpha^{(j)}|^2 \le \frac{\sum_{1 \le h \le m} \left|\mathrm{Tr}_{L/K}(\alpha)^{(h)}\right|^2}{n} + \gamma_{m(n-1)} \left(\frac{|d(L)|}{n^m |d(K)|}\right)^{1/(m(n-1))}.$$

(2) *Such an α can be chosen arbitrarily modulo addition of an arbitrary element of $\mathbb{Z}_K$ and also modulo multiplication by an arbitrary root of unity belonging to K.*

We include this relative version not only because it is in the spirit of this book, for the sake of it, but also because it is essential even in the absolute case. Indeed, the element guaranteed to exist by Hunter's theorem is not necessarily primitive (unless n is a prime number), hence the number field search is split into two distinct parts. Given a discriminant bound and a signature, we first search for polynomials of degree n whose roots satisfy Hunter's conditions and keep only those that are irreducible and satisfy the discriminant and signature conditions. In doing so, we may miss defining polynomials for number fields K such that the element α guaranteed to exist by Hunter's theorem does not generate K but rather a subfield of K. In particular, all primitive fields (that is, number fields having no nontrivial subfields) will be obtained, but also many imprimitive fields as well (those for which $K = \mathbb{Q}(\alpha)$).

In a second part, we search for imprimitive fields using Martinet's theorem. We are helped by the first part because Hunter's inequality gives upper bounds for the discriminants of the possible subfields (see the degree 4 examples below).

Another important result, also due to J. Martinet, is the following. (I thank F. Diaz y Diaz for having brought this theorem to my attention.)

Theorem 9.3.3. *Let K be a number field of degree n, let $\alpha \in K$ be given by Hunter's Theorem 9.3.1, let $k = \mathbb{Q}(\alpha)$ and $m = [k : \mathbb{Q}]$, so that $m \mid n$ and $m > 1$. Then*

$$|d(k)| \le \left(\frac{2\gamma_{n-1}(|d(K)|/n)^{1/(n-1)}}{n - n/m}\right)^{m(m-1)/2}.$$

Remarks

(1) Since we also have the bound $|d(k)| \le |d(K)|^{m/n}$ coming from Theorem 2.5.1, Martinet's bound is useful only for $m = 2$ and also for $m = 3$ if $|d(K)|$ is not too large (see Exercise 11).

(2) It is essential to note that in this theorem α is an element given by Hunter's theorem, and not any element of a subfield of K; otherwise, the result is trivially false (see Exercise 13).

We proceed as follows. Assume that we want to find all number fields of signature (r_1, r_2) and absolute value of discriminant less than or equal to some bound B. Let α be the element guaranteed to exist by Hunter's theorem, and let

$$X^n - a_1 X^{n-1} + \cdots + (-1)^n a_n = \prod_{j=1}^{n} (X - \alpha^{(j)})$$

be its characteristic polynomial. Hunter's theorem implies that we can take $0 \le a_1 \le n/2$ and $\sum_{1 \le j \le n} |\alpha^{(j)}|^2 \le t_2$ for some bound

$$t_2 = \frac{a_1^2}{n} + \gamma_{n-1} \left(\frac{B}{n}\right)^{1/(n-1)} \le \frac{n}{4} + \gamma_{n-1} \left(\frac{B}{n}\right)^{1/(n-1)}$$

depending only on n and B.

The coefficient a_k is the kth elementary symmetric function of the $\alpha^{(j)}$. Using this inequality very crudely, we can say that $|\alpha^{(j)}| \le t_2^{1/2}$ for all j, hence that

$$|a_k| \le \binom{n}{k} t_2^{k/2} .$$

If we use the arithmetic–geometric mean inequality on the numbers $|\alpha^{(j)}|^2$, we obtain

$$|a_n|^{2/n} = \left(\prod |\alpha^{(j)}|^2\right)^{1/n} \le \frac{1}{n} \sum |\alpha^{(j)}|^2 \le \frac{t_2}{n} ,$$

so that $|a_n| \le t_2^{n/2}/n^{n/2}$, which is much better than the bound $|a_n| \le t_2^{n/2}$ obtained above.

It will also be extremely useful to use the power sums $s_k = \sum_{1 \le j \le n} \alpha^{(j)^k}$, which are linked to the coefficients a_k through Newton's formulas

$$k a_k = \sum_{j=1}^{k} (-1)^{j-1} a_{k-j} s_j ,$$

where we have set $a_0 = 1$ (we will see this in more detail in Section 9.3.2 below).

The s_k are integers, and $s_1, \ldots, s_k$ determine the a_j uniquely for $j \le k$. We will see in Lemma 9.3.6 that $|s_k| \le t_2^{k/2}$ (the bound on $|\alpha^{(j)}|$ only gives $|s_k| \le n t_2^{k/2}$). It follows that there are at most $2t_2^{k/2} + 1$ possible values for s_k. In fact, Newton's formulas imply that $s_k \equiv \sum_{j=1}^{k-1} (-1)^{k-j+1} a_{k-j} s_j \pmod{k}$,

so that given $s_1, \dots, s_{k-1}$, the congruence class of s_k modulo k is determined. Thus, given s_j for $j < k$, there are approximately $(2/k)t_2^{k/2}$ possible values for s_k, hence for a_k. Since $0 \leq a_1 = s_1 \leq n/2$ and $|a_n| \leq t_2^{n/2}/n^{n/2}$, it follows that the total number of polynomials to be considered will be at most approximately equal to

$$\frac{2^{n-2}}{n^{(n-2)/2}(n-1)!} t_2^{(n-1)(n+2)/4} .$$

Most of the inequalities above can and should be improved in the specific cases considered. We must also use the information coming from the signature (r_1, r_2) of the desired number fields. This gives precious additional inequalities that considerably restrict the domain in which we must look for suitable polynomials.

A remark should be made, however. Whatever clever inequalities are used (some very specific to small degrees), as far as the author is aware, nobody knows how to decrease the dependence on t_2: the search region still stays of the order of $t_2^{(n-1)(n+2)/4}$ or, equivalently, of the order of $B^{(n+2)/4}$, which is the reason for which it becomes prohibitively expensive even for small degrees such as 10. Recall also that according to the Odlyzko bounds, or simply by Minkowski's theorem, the smallest discriminant in degree n independently of the signature grows at least like C^n for some C; hence the search region is always at least of the order of $C^{n(n+2)/4}$.

The above reasoning applied also to the imprimitive case gives, in fact, an upper bound for the number of number fields of bounded discriminant.

Proposition 9.3.4. *The number of nonisomorphic number fields of fixed degree n and discriminant in absolute value bounded by B is at most equal to $c \cdot B^{(n+2)/4}$ for some constant c depending only on n.*

This proposition is sharp for $n = 2$, but not for $n = 3$ since we know from the Davenport–Heilbronn theory that the number of cubic fields is bounded by $c \cdot B$ (see Theorems 8.5.6 and 8.6.5). For $n = 4$, it gives the bound $c \cdot B^{3/2}$, which is almost certainly not sharp.

In fact, J. Martinet, the author (and certainly other people) have formulated the following strong conjectures.

Conjecture 9.3.5. *Let $n \geq 2$ be an integer.*

(1) *The number of nonisomorphic number fields of degree n and absolute value of discriminant less than or equal to B is asymptotic to $c_n \cdot B$ when $B \to \infty$ for some positive constant c_n depending on n.*
(2) *The number of nonisomorphic number fields of all possible degrees n and absolute value of discriminant less than or equal to B is asymptotic to $c \cdot B$ when $B \to \infty$ for some absolute positive constant c.*

(3) *Conjecture* (1) *should also be true if one fixes a signature* (r_1, r_2) *such that* $r_1 + 2r_2 = n$, *or more generally if one fixes a specific decomposition type of a finite number of prime numbers.*

Very little is known about these conjectures. If G is a transitive subgroup of S_n, denote by $N_n(G, X)$ the number of isomorphism classes of number fields whose discriminant is in absolute value less than or equal to X and whose Galois closure has Galois group isomorphic to G, all signatures together (similar results exist when the signatures are also given). Then it is known that

$$N_2(C_2, X) \sim \frac{6}{\pi^2} X \ ,$$

$$N_3(C_3, X) \sim \frac{11\sqrt{3}}{36\pi} \prod_{p \equiv 1 \pmod 6} \left(1 - \frac{2}{p(p+1)}\right) X^{1/2} \ ,$$

$$N_3(S_3, X) \sim \frac{1}{3\zeta(3)} X \ ,$$

$$N_4(C_4, X) \sim \frac{3}{\pi^2} \left(\left(1 + \frac{\sqrt{2}}{24}\right) \prod_{p \equiv 1 \pmod 4} \left(1 + \frac{2}{p^{3/2} + p^{1/2}}\right) - 1\right) X^{1/2} \ ,$$

$$N_4(C_2 \times C_2, X) \sim \frac{23}{960} \prod_p \left(\left(1 + \frac{3}{p}\right)\left(1 - \frac{1}{p}\right)^3\right) X^{1/2} \log^2 X \ ,$$

$$N_4(D_4, X) \sim \frac{3}{\pi^2} \left(\sum_D \frac{2^{(\mathrm{sign}(D)-1)/2}}{D^2} \frac{L\left(\left(\frac{D}{\cdot}\right), 1\right)}{L\left(\left(\frac{D}{\cdot}\right), 2\right)}\right) X \ ,$$

$$c_1 X^{1/2} \leq N_4(A_4, X) \leq c_2 X \log^4 X \ ,$$

$$c_3 X \leq N_4(S_4, X) \leq c_4 X^{3/2} \ ,$$

for positive constants c_i.

In the above, the sum in the expression for $N_4(D_4, X)$ is over all fundamental discriminants (not including 1), and $L\left(\left(\frac{D}{\cdot}\right), s\right)$ is the usual Dirichlet L-function for the quadratic character $\left(\frac{D}{n}\right)$.

The result for C_2 is elementary (see, for example, [Coh0, Exercise 1 of Chapter 5]). For C_3 it follows from the explicit description of cyclic cubic fields given, for example, in [Coh0, Theorem 6.4.6] (see Exercise 12). For S_3 it is exactly the result of Davenport and Heilbronn (Theorems 8.5.6 and 8.6.5). The results for C_4, $C_2 \times C_2$, and A_4 and the lower bound for S_4 follow from [Bai] with several errors corrected. This paper also gives the weaker result $c \cdot X \leq N_4(D_4, X) \leq c' \cdot X$. The stronger result given for D_4 follows from forthcoming work of the author and collaborators (see Exercise 6 for an indication). Finally, the upper bound for S_4 is a special case of Proposition 9.3.4. Note that [Bai], which uses class field theory, gives a slightly weaker upper bound.

Thus, in degree 4 it is known that number fields with Galois groups $C_2 \times C_2$ and C_4 have density 0, which is probably also the case for A_4, although it has not been proved. On the other hand, assuming the above conjecture, number fields with Galois group D_4 have positive density. The great difficulty comes from the fields with Galois group S_4 for which the upper bound $c_4 X^{3/2}$ stated above is certainly far from the truth. Work is in progress on this subject by the author and collaborators, and by D. Wright and A. Yukie (see [Wri1] and [Wri-Yuk] for advances in this direction), and it seems that a proof of asymptotic equalities for A_4 and S_4 is within reach.

Resuming our discussion of the use of the geometry of numbers, the method, when reasonably applicable, works as follows.

Using Hunter's theorem, first look for number fields of signature (r_1, r_2) and absolute discriminant bounded by B which are of the form $K = \mathbb{Q}(\alpha)$ for Hunter's element α, by making a careful analysis of the best inequalities that can be obtained from the a priori knowledge of the signature, of B, of the coefficient a_1, and usually also of a_n. This can consist of very complicated sets of inequalities depending perhaps on auxiliary conditions, but every bit is good to take since the time spent in analyzing these inequalities is completely negligible compared to the time spent in looking for the polynomials. This is a case-by-case study for each of the possible small signatures.

Second, explore as cleverly as possible the complete range of polynomials whose coefficients satisfy the inequalities. For each of these polynomials, check whether it is irreducible and if the discriminant of the corresponding number field is less than B in absolute value (this will rarely be the case). For each polynomial obtained in this way, apply a strong polynomial reduction algorithm such as [Coh0, Algorithm 4.4.12]. If the polynomial thus obtained is new, keep it. As a last step, check whether or not the number fields defined by the polynomials obtained are isomorphic. Thanks to the strong polynomial reduction that has been performed, this will very rarely be the case.

Finally, do the same for imprimitive fields that are not of the form $K = \mathbb{Q}(\alpha)$ by using Martinet's Theorems 9.3.2 and 9.3.3 (this will be a much faster search for a given absolute degree), and perform an isomorphism check with the fields already obtained by using Hunter's theorem.

Note that it is not necessary to check whether the fields obtained by using Hunter's theorem are indeed primitive. It can easily be done by computing the Galois group, but in any case if a nonprimitive field is encountered, then either it is rejected because the characteristic polynomial of α is not irreducible (it will even be a power of a polynomial) or it is kept because α is a primitive element, and then the isomorphism check done at the very end will ensure that we do not keep it twice.

In the next three subsections, we give useful inequalities of the sort stated above in the use of Hunter's theorem. We give the results and proofs for degree 4, and refer to the literature for the results in larger degree. Note that

Lagrange multipliers (which we explain below) can be used in degree 4 but are really useful only in degrees 5 or more.

9.3.2 General Inequalities

Let K be a number field of degree n and signature (r_1, r_2), let α be the element given by Hunter's theorem, and let

$$P(X) = X^n - a_1 X^{n-1} + a_2 X^{n-2} - \cdots + (-1)^n a_n$$

be the characteristic polynomial of α, so that a_k is the kth elementary symmetric function of α in K. Note that since we do not necessarily assume that $K = \mathbb{Q}(\alpha)$, the polynomial is not necessarily the minimal polynomial of α, but only a power of it in general; hence $P(X)$ is equal to the power of a nonlinear, irreducible polynomial (since $\alpha \notin \mathbb{Z}$) and, in particular, $a_n \neq 0$.

We set $a_0 = 1$, we denote for simplicity by α_j the conjugates of α, ordered in the usual way so that $\alpha_j \in \mathbb{R}$ for $1 \leq j \leq r_1$ and $\alpha_{j+r_2} = \overline{\alpha_j}$ for $r_1+1 \leq j \leq r_1 + r_2$, and finally we set $s_k = \sum_{1 \leq j \leq n} \alpha_j^k$. Recall that Newton's formulas give the induction

$$k a_k = \sum_{j=1}^{k} (-1)^{j-1} a_{k-j} s_j \ .$$

Let us see what general inequalities can be obtained from Hunter's theorem without making any assumption on the signature of K.

Hunter's theorem tells us that $0 \leq a_1 \leq n/2$ and that

$$T_2 = \sum |\alpha_j|^2 \leq t_2 = \frac{a_1^2}{n} + \gamma_{n-1} \left(\frac{B}{n}\right)^{1/(n-1)} \ .$$

Using the arithmetic–geometric mean inequality as above, we deduce that $|a_n| \leq T_2^{n/2}/n^{n/2} \leq t_2^{n/2}/n^{n/2}$.

For the other coefficients, using Newton's formulas we see that we must give bounds for s_k. We already have $s_1 = a_1$, so $0 \leq s_1 \leq n/2$.

For s_2, we can write

$$|s_2| = \left|\sum_j \alpha_j^2\right| \leq \sum_j |\alpha_j|^2 \leq t_2 \ ,$$

so that $-t_2 \leq s_2 \leq t_2$.

This can be slightly improved as follows. Write $\alpha_j = x_j + i y_j$ (real and imaginary parts). Then since a_1, a_2, and s_2 are real, we have $a_1 = \sum_j x_j$, $t_2 = \sum_j x_j^2 + \sum_j y_j^2$, and $s_2 = \sum_j x_j^2 - \sum_j y_j^2$. It follows by Schwartz's inequality that

$$s_2 + t_2 = 2 \sum_j x_j^2 \geq \frac{2}{n} \Big(\sum_j x_j\Big)^2 = \frac{2}{n} a_1^2 \ ,$$

so we obtain the more refined inequality $(2/n)a_1^2 - t_2 \le s_2 \le t_2$, which is equivalent to the inequality

$$\frac{a_1^2 - t_2}{2} \le a_2 \le \frac{((n-2)/n)a_1^2 + t_2}{2} .$$

For $k \ge 3$, we write

$$|s_k| = \left|\sum_j \alpha_j^k\right| \le \sum_j |\alpha_j|^k ,$$

so if we set $T_k = \sum_j |\alpha_j|^k$, we have $-T_k \le s_k \le T_k$, giving the bounds

$$\frac{\sum_{1\le j\le k-1}(-1)^{j-1}a_{k-j}s_j - T_k}{k} \le a_k \le \frac{\sum_{1\le j\le k-1}(-1)^{j-1}a_{k-j}s_j + T_k}{k} .$$

We thus compute inductively bounds for $a_3, \ldots, a_{n-1}$ (note that once the a_i have been fixed for $i \le j$, the value of s_j is also fixed). It remains to compute bounds for T_k.

A simple method consists in using the following lemma.

Lemma 9.3.6. (1) *If the x_j for $1 \le j \le n$ are nonnegative real numbers and k is a real number such that $k \ge 1$, then*

$$\sum_{1\le j\le n} x_j^k \le \left(\sum_{1\le j\le n} x_j\right)^k .$$

(2) *If the x_j for $1 \le j \le n$ are nonnegative real numbers and k is a real number such that $k \ge 2$, then*

$$\sum_{1\le j\le n} x_j^k \le \left(\sum_{1\le j\le n} x_j^2\right)^{k/2} .$$

Proof. For (1), we could say that for all $p \ge 1$, the L^p norm is less than or equal to the L^1 norm. For a simpler proof, let us show the statement by induction on n. It is trivially true for $n \le 1$. Set

$$f(x_n) = \left(\sum_{1\le j\le n} x_j\right)^k - \sum_{1\le j\le n} x_j^k .$$

Then

$$f'(x_n) = k\left(\left(\sum_{1\le j\le n} x_j\right)^{k-1} - x_n^{k-1}\right) ,$$

and since $k \ge 1$ this is nonnegative. It follows that $f(x_n)$ is a nondecreasing function of x_n, hence that $f(x_n) \ge f(0) \ge 0$ by our induction hypothesis, proving (1).

(2) follows trivially by applying (1) to x_j^2 and $k/2$. □

This lemma implies the inequalities

$$|s_k| \le T_k \le T_2^{k/2} \le t_2^{k/2} \ .$$

For example, with $k = 3$ we obtain

$$\frac{-a_1^3 + 3a_1a_2 - t_2^{3/2}}{3} \le a_3 \le \frac{-a_1^3 + 3a_1a_2 + t_2^{3/2}}{3} \ .$$

In addition to this, there are two special cases where we can significantly reduce the domain of exploration.

(1) If $a_1 = 0$, changing x into $-x$ changes the sign of all the coefficients of degree having a different parity than n, so we may assume, for example, that $a_3 \ge 0$.
(2) Slightly less trivial is that if n is even and $a_1 = n/2$, changing x into $1 - x$ does not change a_1 and a_2, and a_3 is changed into

$$a_3' = (n-2)a_2 - \frac{n(n-1)(n-2)}{12} - a_3 \ .$$

Therefore, we may assume that a_3 is greater than or equal to, or less than or equal to $(n-2)a_2/2 - n(n-1)(n-2)/24$, whichever is preferable. It is not difficult to show that we should choose $a_3 \ge (n-2)a_2/2 - n(n-1)(n-2)/24$ if $a_2 \le n(3n-2)/24$, and $a_3 \le (n-2)a_2/2 - n(n-1)(n-2)/24$ if $a_2 \ge n(3n-2)/24$, so that the difference between the upper and lower bounds for a_3 is as small as possible (see Exercise 14).

Finally, see Exercise 15 for still other inequalities.

9.3.3 The Totally Real Case

In the totally real case ($r_2 = 0$), we have stronger results. First, note that $s_2 = T_2$ by definition; hence in all the bounds involving T_2, where in the general case we replace T_2 by an upper bound t_2, here we should replace T_2 by s_2. For example, we use the inequality $|s_k| \le s_2^{k/2}$ instead of the much weaker inequality $|s_k| \le t_2^{k/2}$, where t_2 is the bound given by Hunter's theorem.

Second, we have the following easy result.

Proposition 9.3.7. *If $r_2 = 0$, then we have the following inequalities:*

(1) $|a_n| < s_2^{n/2}/n^{n/2}$;
(2) $s_2 \ge n+1$ *or, equivalently,* $a_2 \le (a_1^2 - 1 - n)/2$;
(3) *for k even,* $s_k > n\,|a_n|^{k/n}$.

Proof. Since $T_2 = s_2$, the above remark and the arithmetic–geometric mean inequality used above gives the bound $|a_n| < s_2^{n/2}/n^{n/2}$, the inequality being strict since the arithmetic–geometric mean inequality becomes an equality only if all components are equal, which is not possible in the totally real case. Since $|a_n| \geq 1$, this same inequality can also be used backwards to say that $s_2 \geq n+1$ or, equivalently, that $a_2 \leq (a_1^2 - 1 - n)/2$. Similarly for k even, the arithmetic–geometric mean inequality gives the stated inequality. □

Remark. It follows from deeper work of C.-L. Siegel (see [Sie]) that in fact $s_2 \geq 3n/2$ for $n \geq 2$ or, equivalently, $a_2 \leq a_1^2/2 - 3n/4$ for $n \geq 2$. Later work by C. Smyth (see [Smy2]) has improved this result to $s_2 \geq 1.7719n$ with a small number of exceptions. In addition, in [Smy1] it is shown that $s_2 \geq 2n - 1$ for $n \leq 7$, while this is not true for $n \geq 8$. It has in fact been proved by C. Smyth (see [Smy3]) that there exists a constant $c < 2$ such that $s_2 \geq c \cdot n$ cannot be true for sufficiently large n.

A consequence of the above results is the inequality $a_2 \leq a_1^2/2 - n + 1/2$ valid for $n \leq 7$.

In fact, it is easy to prove that $s_2 \geq ((n-1)/2)\,\mathrm{disc}(P)^{2/(n(n-1))}$ (see Exercise 18), so that $s_2 \to \infty$ when $|\mathrm{disc}(P)| \to \infty$ for fixed degree n. Thus, to check that $s_2 \geq 2n - 1$ for a given small n, it is enough to look at the finite number of totally real number fields K of degree n up to isomorphism, such that $d(K) \leq ((4n-2)/(n-1))^{n(n-1)/2}$, and all the possible polynomials P. Unfortunately, this is totally impractical for $n \geq 5$ and even quite difficult for $n = 4$ (see Exercise 19). A much better and more realistic method to obtain optimal lower bounds for s_2 is given in [Smy1].

Third, in the totally real case we also have slightly subtler inequalities coming from Newton's formulas.

Proposition 9.3.8. *If $r_2 = 0$, then for $1 \leq k \leq n-1$ we have the inequality*

$$a_{k-1}a_{k+1} \leq \frac{k(n-k)}{(k+1)(n-k+1)}a_k^2 \ .$$

Proof. Write as usual $P(X) = \prod_j (X - x_j)$ with the x_j real by assumption, and let x be a real number different from the x_j. Then

$$\frac{P'(x)}{P(x)} = \sum_j \frac{1}{x - x_j} \ ;$$

hence

$$\frac{P(x)P''(x) - P'(x)^2}{P(x)^2} = -\sum_j \frac{1}{(x-x_j)^2} \ .$$

By Schwartz's inequality we have

$$\left(\sum_j \frac{1}{x-x_j}\right)^2 \le n \sum_j \frac{1}{(x-x_j)^2} \ ;$$

hence

$$\frac{P(x)P''(x)-P'(x)^2}{P(x)^2} \le -\frac{1}{n}\left(\frac{P'(x)}{P(x)}\right)^2 \ ,$$

from which it follows that

$$P(x)P''(x) \le \frac{n-1}{n}P'(x)^2 \ .$$

Now since $P(X)$ is totally real, it follows that all the derivatives of $P(X)$ are also totally real (there must always be a root of $P'(X)$ between two roots of $P(X)$), hence we may apply this result to the $(n-k+1)$st derivative of $P(x)$. Using the formula $P^{(n-k)}(0) = (-1)^{n-k}(n-k)!a_k$ and simplifying, we obtain the inequality of the proposition. □

See Exercise 16 for a way to obtain slightly stronger inequalities of this type.

Finally, in the totally real case we can also obtain other useful inequalities by using the positive definiteness of a canonical quadratic form (see Exercise 17).

9.3.4 The Use of Lagrange Multipliers

It is possible to improve the above inequalities for s_k with $k \ge 3$ to a considerable extent by using an idea due to M. Pohst (see [Poh]). There is, however, a price to pay, in that the computations of the bounds obtained by Pohst's method are not as easy as the ones above (we must usually solve several polynomial equations). This is why we have insisted in giving many simple tricks, although most are superseded by the inequalities obtained by Pohst's method. In actual practice, it is essential to use first the simple inequalities of the preceding sections to restrict as much as possible the number of polynomials to be studied before using the more sophisticated bounds obtained in this section.

We have given bounds for a_1, a_2, and a_n, and these bounds can be considered as quite reasonable. Pohst's idea consists of computing a bound for s_k by considering a_1, $s_2 = a_1^2 - 2a_2$, t_2, and a_n as given and using the theory of Lagrange multipliers as follows. Recall that the α_j are ordered in a way compatible with the signature, so we set $x_j = \alpha_j$ for $1 \le j \le r_1$, $x_j = \mathrm{Re}(\alpha_j) = \mathrm{Re}(\alpha_{j+r_2})$ for $r_1 < j \le r_1 + r_2$, and $x_j = \mathrm{Im}(\alpha_{j-r_2}) = -\mathrm{Im}(\alpha_j)$ for $r_1 + r_2 < j \le r_1 + 2r_2 = n$. Then, setting $\mathbf{x} = (x_1, \ldots, x_n)$, we clearly have

$$g_1(\mathbf{x}) = \sum_{1\le j\le r_1} x_j + 2 \sum_{r_1<j\le r_1+r_2} x_j - a_1 = 0 \ ,$$
$$g_2(\mathbf{x}) = \sum_{1\le j\le r_1} x_j^2 + 2 \sum_{r_1<j\le r_1+r_2} (x_j^2 - x_{j+r_2}^2) - s_2 = 0 \ ,$$
$$g_3(\mathbf{x}) = \prod_{1\le j\le r_1} x_j \prod_{r_1<j\le r_1+r_2} (x_j^2 + x_{j+r_2}^2) - a_n = 0 \ ,$$
$$g_4(\mathbf{x}) = \sum_{1\le j\le r_1} x_j^2 + 2 \sum_{r_1<j\le r_1+r_2} (x_j^2 + x_{j+r_2}^2) - t_2 \le 0 \ .$$

On the other hand, we want to find bounds for s_k, hence for the functions

$$f_k(\mathbf{x}) = \sum_{1\le j\le r_1} x_j^k + 2 \sum_{r_1<j\le r_1+r_2} \mathrm{Re}\left((x_j + ix_{j+r_2})^k\right) \ .$$

Here we assume k integral, but not necessarily positive.

Following Pohst, let G be the set of vectors $\mathbf{x} \in \mathbb{R}^n$ satisfying $g_k(\mathbf{x}) = 0$ for $1 \le k \le 3$ and $g_4(\mathbf{x}) \le 0$, let G_3 be the set of vectors $\mathbf{x} \in \mathbb{R}^n$ satisfying only $g_k(\mathbf{x}) = 0$ for $1 \le k \le 3$, and finally let G_4 be the set of vectors $\mathbf{x} \in \mathbb{R}^n$ satisfying $g_k(\mathbf{x}) = 0$ for $1 \le k \le 4$. We thus have $G_4 \subset G \subset G_3$. In addition, because of the function g_4, it is clear that G and G_4 are closed and bounded and hence are compact subsets of $\mathbb{R}^n$. This is not true for G_3 in general, unless $r_2 = 0$ (in other words, in the totally real case), in which case everything simplifies anyway since the equality $g_2(\mathbf{x}) = 0$ implies $g_4(\mathbf{x}) \le 0$; hence $G = G_3$ in that case.

Since $a_n \ne 0$, we have $x_i \ne 0$ for $i \le r_1$ and $(x_i, x_{i+r_2}) \ne (0,0)$ for $r_1 < i \le r_1 + r_2$. Since G is compact and f_k is continuous on G (even for $k < 0$ by what we have just said), it follows that f_k has a global maximum B_k and a global minimum b_k on G, and hence the desired inequality for s_k will be $b_k \le s_k \le B_k$.

Let $\mathbf{x} \in \mathbb{R}^n$ be a global extremum of f_k. Then either $\mathbf{x} \in G_4$, and then $\mathbf{x}$ is a local extremum of f_k in G_4, or $\mathbf{x} \in G_3$ but $\mathbf{x} \notin G_4$ — in other words, $\mathbf{x} \in G_3$ and $g_4(\mathbf{x}) < 0$, so that g_4 does not enter into the local conditions for x — hence $\mathbf{x}$ is a local extremum of f_k in G_3. To summarize, the procedure for finding the global extrema of f_k in G is as follows.

(1) Find the local extrema $\mathbf{x}$ of f_k in G_3, and keep those such that $g_4(\mathbf{x}) < 0$.
(2) Find the local extrema $\mathbf{x}$ of f_k in G_4.
(3) Find the global minimum and maximum of f_k in G by respectively finding the minimum and maximum values of f_k among the finite number of vectors $\mathbf{x}$ found in (1) and (2).

Finding local extrema is easily done in principle using Lagrange multipliers. For any C^1 function f from $\mathbb{R}^n$ to $\mathbb{R}$ and $\mathbf{x} \in \mathbb{R}^n$, define $f'(\mathbf{x})$ to be the vector of $\mathbb{R}^n$ formed by all the partial derivatives of f at $\mathbf{x}$; in other words,

$$f'(\mathbf{x}) = \left(\frac{\partial f}{\partial x_j}(\mathbf{x})\right)_{1 \le j \le n} .$$

The result is as follows (see any multivariable calculus textbook).

Proposition 9.3.9. *Let $g_1, \ldots, g_m$ and f be C^1 functions in $\mathbb{R}^n$, let A be the subset of $\mathbb{R}^n$ defined by the equations $g_k(\mathbf{x}) = 0$ for $1 \le k \le m$, and let $\mathbf{x}$ be a local extremum of the function f on A. Then the vectors $f'(\mathbf{x})$ and $g_k'(\mathbf{x})$ for $1 \le k \le m$ are linearly dependent; in other words, there exist $\lambda_k \in \mathbb{R}$ and $\lambda_0 \in \mathbb{R}$ not all equal to zero such that if we set $g(\mathbf{x}) = \lambda_0 f(\mathbf{x}) + \sum_{1 \le k \le m} \lambda_k g_k(\mathbf{x})$, then $\dfrac{\partial g}{\partial x_j}(\mathbf{x}) = 0$ for all j with $1 \le j \le n$.*

Thus, we may apply this proposition to $A = G_3$ and $A = G_4$, but *not* to G itself since there is an inequality in the definition of G, and this is the reason for which we have had to introduce the auxiliary sets G_3 and G_4. Note also that, as in the one variable case, the condition is necessary but not sufficient for $\mathbf{x}$ to be an extremum (consider $f(x) = x^3$ at $x = 0$).

Let us study the consequences of this proposition for our specific problem. Since we want to obtain inequalities for s_k with $3 \le k \le n-1$ (and also for $k = -1$), we may assume that $n \ge 4$.

Consider first the extrema of f_k in G_3. Let $J(\mathbf{x})$ be the Jacobian matrix of the g_k at $\mathbf{x}$ for $1 \le k \le 3$; in other words, the $n \times 3$ matrix $J(\mathbf{x}) = (J_{j,k}(\mathbf{x}))_{1 \le j \le n,\ 1 \le k \le 3}$ with

$$J_{j,k}(\mathbf{x}) = \frac{\partial g_k}{\partial x_j}(\mathbf{x}) .$$

If $\mathbf{x}$ is a local extremum of f_k, then the above proposition says that if $M(\mathbf{x})$ is the $n \times 4$ matrix obtained by concatenating $J(\mathbf{x})$ with $f_k'(\mathbf{x})$, then $M(\mathbf{x})$ must have rank at most 3, that is, all the 4×4 subdeterminants of $M(\mathbf{x})$ are equal to 0. This gives one or several equations for the x_j, which one then solves numerically. Note that the condition that $M(\mathbf{x})$ has rank at most equal to 3 is necessary for $\mathbf{x}$ to be a local extremum, but it is not sufficient since the matrix $J(\mathbf{x})$ may have rank at most equal to 2. This is of no importance, however, since it will only add an unnecessary finite number of points at which to compute the value of f_k.

Finding the extrema of f_k in G_4 is similar. The matrix $M(\mathbf{x})$ is now an $n \times 5$ matrix that must have rank at most 4. If $n \ge 5$, this leads to one or several equations for the x_j, which one solves numerically. For $n = 4$, however, the situation is different since for $r_2 > 0$, G_4 is reduced to a finite number of points, so we simply evaluate f_k on these points to find the extrema in this case (we give the explicit results below).

We may also use Pohst's method in a weaker form, which already gives good results as follows.

Proposition 9.3.10. *Let t_2 as above be any upper bound for T_2 (in particular, we may take $t_2 = s_2$ in the totally real case), and let $r = t_2/|a_n|^{2/n} \geq n$ by the arithmetic–geometric mean inequality.*

(1) *For $1 \leq m \leq n-1$, the equation*

$$mX^{m-n} + (n-m)X^m = r$$

has either one or two positive roots. Call z_m the smallest such root.

(2) *For any $k \in \mathbb{Z}$, set*

$$t_k = |a_n|^{k/n} \max_{1 \leq m \leq n-1} \left(m z_m^{k(m-n)/2} + (n-m) z_m^{km/2} \right) .$$

Then we have the inequality $|s_k| \leq T_k \leq t_k$.

Proof. Let $x_j = |\alpha_j|$, so that $|s_k| \leq T_k = \sum_j x_j^k$. In Proposition 9.3.9, we take $f(\mathbf{x}) = \sum_j x_j^k$, $g_1(\mathbf{x}) = \sum_j x_j^2 - T_2$, and $g_2(\mathbf{x}) = \prod_j x_j - |a_n|$, restricted to $x_j \geq 0$ for all j. First note that the x_j cannot all be equal. Indeed, the arithmetic–geometric mean inequality applied to the x_j^k shows that $f(\mathbf{x}) \geq n\left(\prod_j x_j\right)^{k/n} = n\,|a_n|^{k/n}$ with equality if and only if all the x_j are equal, hence equality of the x_j corresponds to a *minimum* of $f(\mathbf{x})$ and so can be excluded.

Note further that the set A defined by $g_1(\mathbf{x}) = g_2(\mathbf{x}) = 0$ and $x_j \geq 0$ is compact. It follows from Proposition 9.3.9 that if $\mathbf{x}$ is a local maximum of f in A, then either all the x_j are equal, which as we have seen is excluded, or there exist real numbers λ and μ such that $x_j^{k-1} = \lambda x_j + \mu/x_j$ for all j or, equivalently, $R(x_j) = 0$ with $R(X) = X^k - \lambda X^2 - \mu$ for $k \geq 0$, or $R(1/x_j) = 0$ with $R(X) = X^{2-k} - \mu X^2 - \lambda$ for $k < 0$. Assume $k \geq 0$, the case $k < 0$ being similar. Then all the x_j are positive real roots of $R(X) = 0$. By Exercise 20, we know that $R(X)$ has at most two such roots. On the other hand, $R(X) = 0$ has at least one such root. Indeed, f has at least one extremum $\mathbf{x}$ in A, and for such an extremum we have $x_j \neq 0$ for all j, hence $\mathbf{x}$ is in the interior of A, and in particular is a local extremum, whose components must satisfy $R(X) = 0$.

It follows that the x_j can take at most two distinct values, hence exactly two since the x_j are not all equal. So assume $x_1 = x_2 = \cdots = x_m = x$ and $x_{m+1} = \cdots = x_n = y$, where we may assume that $1 \leq m \leq n/2$. Thus, $0 = g_1(\mathbf{x}) = mx^2 + (n-m)y^2 - T_2$ and $0 = g_2(\mathbf{x}) = x^m y^{n-m} - |a_n|$.

Set

$$R_{m,k}(X) = mX^{k(m-n)/2} + (n-m)X^{km/2} .$$

Since $m \geq 1$, we can solve for y and deduce that $R_{m,2}(z) - T_2/|a_n|^{2/n} = 0$, with $z = y\,|a_n|^{-1/n}$.

We have the following lemma.

Lemma 9.3.11. *Let $R_{m,k}(x)$ be the above function with $1 \leq m \leq n-1$.*

(1) *For $x > 0$, the function $R_{m,k}(x)$ has a unique minimum at 1, it decreases for $x < 1$, and it increases for $x > 1$.*
(2) *For any $t_2 \geq T_2$, the function $R_{m,2}(x) - t_2/|a_n|^{2/n}$ has exactly two (possibly equal) roots z_m and z'_m, which satisfy $z_m \leq 1 \leq z'_m$.*
(3) *We have $z'_m = 1/z_{n-m}$ and $R_{m,k}(z'_m) = R_{n-m,k}(z_{n-m})$.*
(4) *The root $z_m = z_m(t_2)$ is a decreasing function of t_2, and $R_{m,k}(z_m)$ is an increasing function of t_2.*

Assuming this lemma for the moment, we see from Proposition 9.3.9 that

$$T_k = f(\mathbf{x}) \leq |a_n|^{k/2} \max_{1 \leq m \leq n/2} (\max(R_{m,k}(z_m(T_2)), R_{m,k}(z'_m(T_2))))$$
$$= |a_n|^{k/2} \max_{1 \leq m \leq n-1} R_{m,k}(z_m(T_2)) \leq |a_n|^{k/2} \max_{1 \leq m \leq n-1} R_{m,k}(z_m(t_2)) \ ,$$

using the symmetry relation and the fact that $R_{m,k}(z_m(t_2))$ is a nondecreasing function of t_2, proving the proposition.

The proof of the lemma is straightforward. We have

$$R'_{m,k}(x) = \frac{km(n-m)}{2} \left(x^{km/2-1} - x^{k(m-n)/2-1} \right) \ ,$$

so $R'_{m,k}(x) = 0$ if and only if $x = 1$, and it is negative for $x < 1$ and positive for $x > 1$, proving (1).

Since $r = t_2/|a_n|^{2/n} \geq n$ by the inequality for $|a_n|$, it follows that the minimum of $R_{m,2}(x) - r$ for $x > 0$, attained at $x = 1$, is equal to $n - r \leq 0$, hence there exist exactly two (possibly equal) roots z_m and z'_m of $R_{m,2}(x) - r = 0$ such that $z_m \leq 1 \leq z'_m$, proving (2).

For (3), we immediately check the symmetry relation $R_{n-m,k}(x) = R_{m,k}(1/x)$ (which must exist since we may change m into $n-m$ if we exchange x and y).

As the roots of $R_{n-m}(x)$ are z_{n-m} and z'_{n-m} and the function $1/x$ is decreasing, it follows that we have $z_{n-m} = 1/z'_m$. Using the symmetry relation we get $R_{m,k}(z'_m) = R_{n-m,k}(z_{n-m})$, proving (3).

For (4) we use the implicit function theorem, which tells us that

$$\frac{dz_m(t_2)}{dt_2} = \frac{1}{R'_{m,2}(z_m(t_2))} \ .$$

Now $R'_{m,2}(x) = m(n-m)x^{m-1}(1-x^{-n})$, and since $z_m(t_2) \leq 1$, we have $R'_{m,2}(z_m(t_2)) < 0$, so $dz_m(t_2)/dt_2 < 0$, hence $z_m(t_2)$ is a decreasing function of t_2.

Finally, since $R_{m,k}(x)$ is a decreasing function of x for $x < 1$, it follows that $R_{m,k}(z_m(t_2))$ is an increasing function of t_2, finishing the proof of the lemma and of the proposition. □

Remarks

(1) This proposition has the advantage of being simpler than the general method of Lagrange multipliers. The results obtained in this way are weaker, since we have omitted the condition $\sum_j x_j = a_1$, but they are not that much weaker in general (see, for example, Exercise 21).
(2) The equation of the proposition can also be written as the polynomial equation $(n-m)X^n - rX^{n-m} + m = 0$.
(3) See Exercise 22 for an efficient practical method to compute z_m.

9.4 Construction of Tables of Quartic Fields

In the next subsections, we first consider quartic fields K for which the element α given by Hunter's theorem is such that $K = \mathbb{Q}(\alpha)$, which includes in particular all primitive quartic fields, but not only those. Note that a quartic number field K is primitive if and only if the Galois group of its Galois closure is isomorphic to A_4 or to S_4, the cases being distinguished by whether or not the discriminant is a square (see Exercise 23).

We will consider the imprimitive case in Section 9.4.5.

9.4.1 Easy Inequalities for All Signatures

Let $P(X) = X^4 - a_1X^3 + a_2X^2 - a_3X + a_4$ be the characteristic polynomial of the element α given by Hunter's theorem. Since $\gamma_3 = 2^{1/3}$, we have

$$\sum |\alpha^{(j)}|^2 \le t_2 = \frac{a_1^2}{4} + \left(\frac{B}{2}\right)^{1/3} .$$

Hence the general inequalities in the quartic case are

$$\begin{aligned} 0 \le a_1 &\le 2 , \\ \frac{a_1^2 - t_2}{2} \le a_2 &\le \frac{a_1^2/2 + t_2}{2} , \\ \frac{-a_1^3 + 3a_1a_2 - t_2^{3/2}}{3} \le a_3 &\le \frac{-a_1^3 + 3a_1a_2 + t_2^{3/2}}{3} , \\ -\frac{t_2^2}{16} \le a_4 &\le \frac{t_2^2}{16} . \end{aligned}$$

In addition, if $a_1 = 0$, we may assume $a_3 \ge 0$; hence in that case we have the simple inequality $0 \le a_3 \le t_2^{3/2}/3$.

Furthermore, if $a_1 = 2$, then if $a_2 \leq 1$ we may assume $a_3 \geq a_2 - 1$, while if $a_2 \geq 2$ we may assume $a_3 \leq a_2 - 1$.

Remark. If we take into account only the above inequalities, we find that the number of polynomials to consider is asymptotic to $(B/2)^{3/2}/6$ when $B \to \infty$. In each specific signature, however, this bound is improved by at least a factor of 2 (see Exercise 24).

In the quartic case, we also have inequalities coming from the cubic resolvent polynomial (see [Coh0, Exercise 1 of Chapter 4], where the result is stated incorrectly; see the errata sheets). The (corrected) result is the following.

Proposition 9.4.1. *Let* $P(X) = X^4 - a_1X^3 + a_2X^2 - a_3X + a_4$ *be a squarefree monic polynomial with real coefficients. Set* $A(P) = 3a_1^2 - 8a_2$ *and*

$$B(P) = a_2^2 - a_1^2a_2 + \frac{3}{16}a_1^4 + a_1a_3 - 4a_4 \ .$$

(1) P *has signature* $(0,2)$ *if and only if* $\mathrm{disc}(P) > 0$ *and* $A(P) \leq 0$ *or* $B(P) \leq 0$.
(2) P *has signature* $(2,1)$ *if and only if* $\mathrm{disc}(P) < 0$.
(3) P *has signature* $(4,0)$ *if and only if* $\mathrm{disc}(P) > 0$, $A(P) > 0$, *and* $B(P) > 0$.

We can thus use three different tools: first, the general inequalities stated above; second, the inequalities specific to a given signature coming from the above proposition; third, the inequalities (also specific to a given signature) coming from Pohst's idea of using Lagrange multipliers. It should be noted that solving the equations involved in the Lagrange multiplier method is relatively costly, although much less than the gain obtained on the bounds for s_k. This shows, however, that these refined inequalities should be used only for large searches. In particular, in degree 4 they become really useful only if the search bound B is over 10^6 or so. On the other hand, for larger degrees, they become useful (and even essential) much more rapidly.

Note that in Pohst's method we have mentioned the possibility of finding extrema of $s_{-1} = a_{n-1}/a_n$. This is indeed useful in larger degrees, but it does not bring any improvement in the quartic case since s_{-1} is in that case a function of a_1, a_2, s_3, and a_4.

We now consider the different signatures separately.

9.4.2 Signature (0, 2): The Totally Complex Case

If K has signature $(0,2)$ — in other words, if K is totally complex — then $a_4 = \alpha_1\overline{\alpha_1}\alpha_2\overline{\alpha_2} > 0$, so $a_4 \geq 1$. In addition, thanks to Proposition 9.4.1, we can say that $\mathrm{disc}(P) > 0$ and that either $a_2 \geq a_1$ (since $a_2 \geq 3a_1^2/8$ is equivalent to $a_2 \geq a_1$ for $0 \leq a_1 \leq 2$) or that $B(P) \leq 0$.

Pohst's method gives the following result for bounding s_3.

Proposition 9.4.2. *Assume that $P(X)$ has signature $(0,2)$, keep the notation of Section 9.3.4, and let $\mathbf{x} = (x_1, x_2, x_3, x_4)$ be a local extremum on G_3 or G_4 of the function*

$$s_3(\mathbf{x}) = 2(x_1^3 + x_2^3) - 6(x_1x_3^2 + x_2x_4^2) \ .$$

(1) *If $\mathbf{x} \in G_3 \setminus G_4$, then either $x_1 = x_2$, $x_3 = 0$, or $x_4 = 0$.*

a) The case $x_1 = x_2$ is possible only if the relation $a_1^2(8a_2 - 3a_1^2) = 64a_4$ holds, in which case we have either $a_1 = 1$ and $a_2 \geq 9$ or $a_1 = 2$ and $a_2 \geq 4$, and $s_3(\mathbf{x}) = 5a_1^3/8 - 3a_1a_2/2$.

b) If $x_3 = 0$, then $x_4^2 = \left(-s_2 + (s_2^2 + 16a_4)^{1/2}\right)/2$, x_1 and x_2 are the two real roots of the equation

$$8X^2 - 4a_1X + 2a_2 - (s_2^2 + 16a_4)^{1/2} = 0 \ ,$$

and $s_3(\mathbf{x}) = 2(x_1^3 + x_2^3) - 6x_2x_4^2$.

c) The case $x_4 = 0$ is identical to the case $x_3 = 0$ with the indices 1 and 2 exchanged and 3 and 4 exchanged. In particular, for $\mathbf{x} \in G_3 \setminus G_4$ only two values of $s_3(\mathbf{x})$ have to be computed in addition to the given value in a).

(2) *If $\mathbf{x} \in G_4$, then for some choice of signs ε_1 and ε_2, if we set $d_1 = a_1^2 + 2t_2 - 4a_2$ and $d_2 = t_2^2 - 16a_4$ (which are both nonnegative), then $x_1 = (a_1 + \varepsilon_1\sqrt{d_1})/4$, $x_2 = (a_1 - \varepsilon_1\sqrt{d_1})/4$,*

$$x_3^2 = \frac{t_2 - s_2 + 2\varepsilon_2\sqrt{d_2} - \varepsilon_1a_1\sqrt{d_1}}{8} \ ,$$

$$x_4^2 = \frac{t_2 - s_2 - 2\varepsilon_2\sqrt{d_2} + \varepsilon_1a_1\sqrt{d_1}}{8} \ ,$$

and the two extremal values of $s_3(\mathbf{x})$ on G_4 are

$$s_3(\mathbf{x}) = \frac{a_1(3t_2 + 4a_1^2 - 12a_2) + 3\varepsilon\sqrt{d_1d_2}}{4}$$

with $\varepsilon = -\varepsilon_1\varepsilon_2 = \pm 1$.

Proof. With the notation of Section 9.3.4, we have $g_1(\mathbf{x}) = 2(x_1+x_2) - a_1$, $g_2(\mathbf{x}) = 2(x_1^2 + x_2^2 - x_3^2 - x_4^2) - s_2$, $g_3(\mathbf{x}) = (x_1^2 + x_2^2)(x_3^2 + x_4^2) - a_4$, $g_4(\mathbf{x}) = 2(x_1^2 + x_2^2 + x_3^2 + x_4^2) - t_2$, and $s_3(\mathbf{x}) = 2(x_1^3 + x_2^3) - 6(x_1x_3^2 + x_2x_4^2)$.

For (1), assume that $\mathbf{x} \in G_3 \setminus G_4$ is a local extremum of $s_3(\mathbf{x})$. Then by Proposition 9.3.9, the 4×4 matrix whose columns are $g_1'(\mathbf{x})$, $g_2'(\mathbf{x})$, $g_3'(\mathbf{x})$, and $s_3'(\mathbf{x})$ must have rank at most equal to 3; in other words, its determinant D must be equal to 0. A computation shows that

$$D = -192x_3x_4(x_1 - x_2)^2(x_1^2 + x_2^2 + x_3^2 + x_4^2) \ .$$

Thus $D = 0$ if and only if $x_1 = x_2$ or $x_3 = 0$ or $x_4 = 0$.

Assume first that $x_1 = x_2$. Replacing this in the equations for G_3 and in $s_3(\mathbf{x})$, a computation shows that we obtain the relation $a_1^2(8a_2 - 3a_1^2) = 64a_4$, and $s_3(\mathbf{x}) = 5a_1^3/8 - 3a_1a_2/2$, proving a). In the other two cases we can solve for x_1, x_2, x_3, and x_4, replace in $s_3(\mathbf{x})$ and obtain the desired result.

For (2), assume that $\mathbf{x} \in G_4$ (we do not even need to assume that it is a local extremum of $s_3(\mathbf{x})$). This gives a system of four equations in four unknowns, which we can easily solve and once again obtain the desired result. □

9.4.3 Signature (2, 1): The Mixed Case

If K has signature $(2,1)$, Proposition 9.4.1 gives the additional inequality $\mathrm{disc}(P) < 0$, but no other. In particular, note that there is a misprint in [Bu-Fo-Po] ($a_4 = n \geq 0$ with their notation is incorrect), as can be seen, for example, from the field defined by the polynomial $X^4 - X^3 - X^2 + 4X - 1$ (see Exercise 25). The table given in the paper is correct, however.

Pohst's method gives the following result for bounding s_3.

Proposition 9.4.3. *Assume that $P(X)$ has signature $(2,1)$, keep the notation of Section 9.3.4, set*

$$Q(X) = 3X^4 - 2a_1X^3 + a_2X^2 - a_4 = XP'(X) - P(X) \ ,$$

$$S(X) = 12X^3 - 9a_1X^2 + 6a_2X + a_1^3 - 3a_1a_2 \ ,$$

and let $\mathbf{x} = (x_1, x_2, x_3, x_4)$ be a local extremum on G_3 or G_4 of the function

$$s_3(\mathbf{x}) = x_1^3 + x_2^3 + 2x_3^3 - 6x_3x_4^2 \ .$$

(1) *If $\mathbf{x} \in G_3 \setminus G_4$, then either $x_1 = x_2$ or $x_4 = 0$.*
   a) *If $x_1 = x_2$, then x_1 is a real root of $Q(X) = 0$, $x_3 = a_1/2 - x_1$, $x_4 = (x_1^2 + x_3^2 - s_2/2)^{1/2}$, and $s_3(\mathbf{x}) = S(x_1)$. In addition, the condition $g_4(\mathbf{x}) < 0$ is equivalent to $4x_1^2 - 2a_1x_1 + a_2 = x_1^2 + a_4/x_1^2 < t_2/2$.*
   b) *If $x_4 = 0$, then x_3 is a real root of $Q(X) = 0$, x_1 and x_2 are the two real roots of*

$$X^2 - (a_1 - 2x_3)X + 3x_3^2 - 2a_1x_3 + a_2 = 0 \ ,$$

   and we have $s_3(\mathbf{x}) = S(x_3)$ and $g_4(\mathbf{x}) = s_2 < t_2$.
   c) *In both cases $s_3(\mathbf{x})$ is the value of the polynomial $S(X)$ at a real root of $Q(X) = 0$.*

(2) *If $\mathbf{x} \in G_4$, then x_3 must be a real root of*

$$48X^4 - 32a_1X^3 + 8(t_2 - 2s_2 + a_1^2)X^2 - 8a_1(t_2 - s_2)X - (t_2 - s_2)(t_2 + s_2 - 2a_1^2) - 16a_4 = 0 \ ,$$

$x_4 = (t_2 - s_2)^{1/2}/2$, *$x_1$ and x_2 are the two real roots of*

$$4X^2 - 4(a_1 - 2x_3)X + 12x_3^2 - 8a_1x_3 + 2a_1^2 - t_2 - s_2 = 0 \ ,$$

and $s_3(\mathbf{x}) = S(x_3) - 3(t_2 - s_2)(x_3 - a_1/4)$.

Proof. With the notation of Section 9.3.4, we have $g_1(\mathbf{x}) = x_1 + x_2 + 2x_3 - a_1$, $g_2(\mathbf{x}) = x_1^2 + x_2^2 + 2x_3^2 - 2x_4^2 - s_2$, $g_3(\mathbf{x}) = x_1x_2(x_3^2 + x_4^2) - a_4$, $g_4(\mathbf{x}) = x_1^2 + x_2^2 + 2x_3^2 + 2x_4^2 - t_2$, and $s_3(\mathbf{x}) = x_1^3 + x_2^3 + 2x_3^3 - 6x_3x_4^2$.

For (1), assume that $\mathbf{x} \in G_3 \setminus G_4$ is a local extremum of $s_3(\mathbf{x})$. Then by Proposition 9.3.9, the 4×4 matrix whose columns are $g_1'(\mathbf{x})$, $g_2'(\mathbf{x})$, $g_3'(\mathbf{x})$, and $s_3'(\mathbf{x})$ must have rank at most equal to 3; in other words, its determinant D must be equal to 0. A computation (using the fact that $g_3(\mathbf{x}) = 0$) shows that

$$D = 24x_4(x_1 - x_2)((x_1 - x_3)^2 + x_4^2)((x_2 - x_3)^2 + x_4^2) \ .$$

Thus $D = 0$ if and only if $x_1 = x_2$ or $x_4 = 0$ (note that if one of the last two factors is equal to 0, we also have $x_4 = 0$), and in both cases we can solve for x_1, x_2, x_3, and x_4, replace the resulting values in $s_3(\mathbf{x})$, and obtain the desired result.

For (2), assume that $\mathbf{x} \in G_4$ (once again we do not even need to assume that it is a local extremum of $s_3(\mathbf{x})$). This gives a system of four equations in four unknowns, which we can easily solve and once again obtain the desired result. □

The above proposition shows that to compute the extrema of s_3 it is sufficient to compute its value on at most eight points.

9.4.4 Signature (4, 0): The Totally Real Case

If K has signature $(4,0)$, in other words if K is totally real, then thanks to Proposition 9.4.1, we can say that $a_2 \le a_1 - 1$ (since $a_2 < 3a_1^2/8$ is equivalent to $a_2 \le a_1 - 1$); hence the inequality for a_2 is improved to

$$\frac{a_1^2 - t_2}{2} \le a_2 \le a_1 - 1 \ .$$

In addition, we can also say that $\mathrm{disc}(P) > 0$ and $B(P) > 0$.

In this case we can also apply Proposition 9.3.7, which gives the stronger inequalities $a_2 \le (a_1^2 - 5)/2$, equivalent to $a_2 \le a_1 - 3$ for $0 \le a_1 \le 2$, $|a_4| < (a_1^2 - 2a_2)^2/16$ (which is usually considerably better than $|a_4| \le t_2^2/16$). Proposition 9.3.8 gives $a_3^2 \ge (8/3)a_2a_4$ and $a_1a_3 \le (4/9)a_2^2$, and using Exercise 16, this can be improved to $a_3^2 \ge (8/3)a_2a_4 + 4|a_4|$ and $a_1a_3 \le (4/9)a_2^2 - (a_3^2/6)^{1/3}$, which further restrict the range of possible values of a_3, given a_1, a_2, and a_4.

If we apply Smyth's result $s_2 \ge 2n - 1$ for $n \le 7$ mentioned above, we obtain the improvement $a_2 \le a_1 - 4$.

Pohst's method gives the following result for bounding s_3. Recall that in the totally real case we have $G = G_3$ and $G_4 = G_3$ if $t_2 = s_2$, $G_4 = \varnothing$ otherwise, so we want to study extrema on G_3.

Proposition 9.4.4. *Assume that $P(X)$ has signature $(4,0)$, keep the notation of Section 9.3.4, set*

$$Q(X) = 3X^4 - 2a_1X^3 + a_2X^2 - a_4 = XP'(X) - P(X) \ ,$$
$$S(X) = 12X^3 - 9a_1X^2 + 6a_2X + a_1^3 - 3a_1a_2 \ ,$$

and let $\mathbf{x} = (x_1, x_2, x_3, x_4)$ be a local extremum on G_3 of the function

$$s_3(\mathbf{x}) = x_1^3 + x_2^3 + x_3^3 + x_4^3 \ .$$

Then at least two of the x_i are equal. If, for example, $x_3 = x_4$, then $x_3 = x_4$ is a real root of $S(X) = 0$, x_1 and x_2 are the two real roots of

$$X^2 - (a_1 - 2x_3)X + 3x_3^2 - 2a_1x_3 + a_2 = 0 \ ,$$

and $s_3(\mathbf{x}) = S(x_3)$.

Proof. This signature is much simpler than the other two. Indeed, we have $g_1(\mathbf{x}) = x_1 + x_2 + x_3 + x_4 - a_1$, $g_2(\mathbf{x}) = x_1^2 + x_2^2 + x_3^2 + x_4^2 - s_2$, $g_3(\mathbf{x}) = x_1x_2x_3x_4 - a_4$, and $s_3(\mathbf{x}) = x_1^3 + x_2^3 + x_3^3 + x_4^3$. If follows that the 4×4 matrix that we must consider according to Proposition 9.3.9 is, up to trivial factors, a Vandermonde matrix in the x_i, so its determinant vanishes if and only if two of the x_i are equal. Replacing and solving for the x_i and $s_3(\mathbf{x})$ gives the proposition. □

9.4.5 Imprimitive Degree 4 Fields

These are quartic fields L whose Galois group is isomorphic to $C_2 \times C_2$, C_4, or D_4. If K denotes one of the quadratic subfields of L (there is only one in the case C_4 and D_4), then L/K is a quadratic extension, and we know how to construct all such extensions (see Algorithms 9.2.3 and 9.2.4). Furthermore, Theorem 2.5.1 tells us that $|d(K)| \leq |d(L)|^{1/[L:K]}$ so in our case that $|d(K)| \leq |d(L)|^{1/2}$. Finally, if L is of signature (R_1, R_2) and K of signature (r_1, r_2), we must have $r_2 \leq \lfloor R_2/2 \rfloor$. The procedure is thus as follows.

(1) To obtain imprimitive quartic fields of signature $(0,2)$ and discriminant less than or equal to B, we first make the list of all real and imaginary quadratic fields K of discriminant in absolute value less than or equal to $B^{1/2}$. For each such field, we use Algorithm 9.2.3 with $(R_1, R_2) = (0,2)$ to compute all quadratic extensions whose relative discriminant is in norm less than or equal to $B/d(K)^2$. Finally, we remove isomorphic fields, which may only occur when the Galois group is isomorphic to $C_2 \times C_2$.

(2) To obtain imprimitive quartic fields of signature $(2,1)$, which are necessarily with Galois group isomorphic to D_4 (see Exercise 26), we first make a list of all real quadratic fields of discriminant less than or equal to $B^{1/2}$ and then proceed as above, asking for $(R_1, R_2) = (2,1)$.
(3) Finally, to obtain imprimitive quartic fields of signature $(4,0)$, we first make a list of all real quadratic fields of discriminant less than or equal to $B^{1/2}$ and then proceed as above, asking for $(R_1, R_2) = (4,0)$.

The above procedure assumes that we want to make a table of imprimitive quartic fields for their own sake. However, frequently the search for imprimitive fields is done after a search using Hunter's theorem, in which all primitive and some imprimitive fields are discovered. In that case, the discriminant bound $|d(K)| \leq |d(L)|^{1/2}$ can be considerably improved as follows.

Let α be the element given by Hunter's theorem, assumed not to generate K over $\mathbb{Q}$. Since $\alpha \in \mathbb{Z}_K \setminus \mathbb{Z}$, α is an algebraic integer belonging to a quadratic number subfield k of K. If we denote the conjugate of α in k by β, it follows that $(\alpha - \beta)^2/d(k)$ is a perfect integer square. On the other hand, the conjugates of α in K are α, α, β, and β, hence $a_1/2 = \alpha + \beta$ and $s_2/2 = \alpha^2 + \beta^2$, from which it follows that

$$(\alpha - \beta)^2 = -(\alpha + \beta)^2 + 2(\alpha^2 + \beta^2) = s_2 - \frac{a_1^2}{4} .$$

The inequality $\left|s_2 - a_1^2/4\right| \leq t_2 - a_1^2/4$, which we have proved in Section 9.3.2, thus shows that $|d(k)| \leq t_2 - a_1^2/4$. In particular, using Hunter's bound on T_2, we obtain the upper bound $|d(k)| \leq (B/2)^{1/3}$ which is much stronger than the trivial upper bound $|d(k)| \leq B^{1/2}$.

Note that this is exactly what Martinet's Theorem 9.3.3 tells us, but we have given the easy proof in our case.

9.5 Miscellaneous Methods (in Brief)

There are a number of other nonsystematic methods for constructing number fields with small discriminant. As is the case with the class field theory methods (in which case the lists are not systematic unless the Galois group is also specified), these methods cannot find a systematic list of number fields with discriminant up to a given bound, but can try only to find examples of number fields with small discriminant.

The author knows of the following methods, in addition to the ones described in the rest of this chapter:

(1) the search for Euclidean number fields of reasonable degree (less than 12, say) using the notion of cliques of exceptional units invented by H. W. Lenstra. We refer to [Leu] and [Leu-Nik] for details on the most recent results on this subject;

(2) elementary searches for polynomials of small discriminant using various techniques. This is, for example, the subject of recent work by D. Simon (see [Sim2], [Sim3]);
(3) the search for number fields having prescribed ramification at small primes.

I will briefly describe the first two methods. For the third method, I refer to the survey paper [Har] and to the web page of J. Jones mentioned in Appendix B.

9.5.1 Euclidean Number Fields

Let K be a number field. Recall that a unit $\varepsilon \in U(K)$ is said to be *exceptional* if $1-\varepsilon$ is also a unit, and a family $\varepsilon_1, \ldots, \varepsilon_k$ of units is called a *clique* of exceptional units if $\varepsilon_1 = 1$ and $\varepsilon_i - \varepsilon_j$ is a unit for all $i \neq j$. It can be shown that the number of exceptional units is finite, but more importantly for our purposes, conjecturally a field with many exceptional units relative to its degree should have a small discriminant. In addition, if there exists a sufficiently large clique, then K can be shown to be Euclidean for the field norm (see the above-mentioned papers of Leutbecher et al. for precise statements and details).

First H. W. Lenstra and then others such as A. Leutbecher have developed powerful methods for constructing number fields having many exceptional units, in order to construct Euclidean fields. In the process, most of the number fields that they find have a small discriminant, even those which may not be Euclidean.

The method is limited to small degrees, say less than or equal to 16, but is very useful. For example, [Leu-Nik] have found in this way the totally complex number field of degree 10 of smallest known discriminant (it was also found later using the class field method).

A final remark about Euclidean number fields. Proving that a field is Euclidean or not (for the ordinary field norm) is usually not easy, and more and more extensive tables of known Euclidean fields have been made (see [Cav-Lem], [Lem], [Que]). The subject has now lost most of its appeal because, contrary to what many people thought until rather recently, it is highly probable that almost all number fields having a unit rank at least equal to 3 (and perhaps even 2) are norm-Euclidean, and in particular have class number 1. This conjecture was first formulated by H. W. Lenstra, and seems very plausible.

9.5.2 Small Polynomial Discriminants

Other methods for finding number fields of small discriminant are based on the search for small *polynomial* discriminants. Some of these methods are

closely related to the methods used to construct exceptional units, and most are due to D. Simon (see [Sim2], [Sim3]).

A priori, it seems a bad idea to look for polynomials with small discriminant, since in general the discriminant of the corresponding number field is much smaller. For example, the number fields with smallest known discriminants such as those given in Appendix C usually have a defining polynomial whose discriminant is much larger than the corresponding field discriminant (see Exercise 28).

In fact, J.-P. Serre has raised the following interesting question. One knows that there exist families of number fields of arbitrary large degree with bounded root discriminant: this is an immediate consequence of the existence of infinite class field towers due to Golod–Shafarevitch already mentioned in Section 3.1.

Is the same true for polynomials (say irreducible and monic)? In other words, does there exist a family of polynomials P_n with $d_n = \deg(P_n) \to \infty$ and $|\mathrm{disc}(P_n)|^{1/d_n} \leq B$ for some constant B?

It is difficult to determine whether this question should have a positive or negative answer (this is the reason for which Serre calls it a "question" and not a conjecture).

In any case, the big advantage of looking at small polynomial discriminants is that we have much better control on them than on field discriminants, and indeed the work of Simon has produced a large number of improvements on the smallest known *number field* discriminants. I refer to his paper and his thesis for details.

9.6 Exercises for Chapter 9

1. Assume that the relative degree $[L : K]$ is odd. Improve the bound given by Lemma 9.2.1.
2. Give an example of a number field K and two relative extensions L_1/K and L_2/K that are not K-isomorphic but that are $\mathbb{Q}$-isomorphic.
3. Prove that Algorithm 9.2.3 is valid.
4. Let ℓ be an odd prime and let K be a number field such that $\zeta_\ell \in K$.
   a) Let $\mathfrak{a}$ be an integral ideal of K. Show that $\mathfrak{a}$ can be written in a unique way as $\mathfrak{a} = \prod_{1 \leq i \leq \ell} \mathfrak{a}_i^i$ with $\mathfrak{a}_i$ squarefree and pairwise coprime for $1 \leq i \leq \ell - 1$.
   b) Let α and α' in K^*, and set $\alpha\mathbb{Z}_K = \mathfrak{a}$ and $\alpha'\mathbb{Z}_K = \mathfrak{a}'$. Write the ℓ-Kummer-equivalence relation (see Definition 10.2.8) in terms of the ideals $\mathfrak{a}_i$ and $\mathfrak{a}'_i$ defined in (1) and an equivalence relation on ℓ-virtual units.
   c) Deduce from this a modification of Algorithm 9.2.3 which computes the list of cyclic extensions of relative degree ℓ of K assuming that $\zeta_\ell \in K$.
5. Let K be a number field of signature (r_1, r_2), and denote by $\mathrm{rk}_2(G)$ the 2-rank of a finite Abelian group G. Show that the number $N_{2r_1,2r_2}(B)$ of quadratic extensions L/K up to isomorphism with $\mathcal{N}(\mathfrak{d}(L/K)) \leq B$ and signature $(2r_1, 2r_2)$ is given by

$$\sum_{\mathcal{N}(\mathfrak{a})\leq B}\mu(\mathfrak{a})S\left(\frac{B}{\mathcal{N}(\mathfrak{a})}\right)\quad\text{with}\quad S(x)=\sum_{\mathcal{N}(\mathfrak{m})\leq x}\left(2^{\mathrm{rk}_2(Cl_{\mathfrak{m}}(K))}-1\right)\ ,$$

where μ is the usual Möbius function on ideals already used in Section 3.5.2 and $Cl_{\mathfrak{m}}(K)$ is the ray class group of K corresponding to the modulus $\mathfrak{m}$. Give an analogous formula for an arbitrary signature (R_1,R_2) such that $R_1+2R_2=2(r_1+2r_2)$ and $R_1\leq 2r_1$.

6. The following exercise summarizes some results contained in forthcoming work of the author and collaborators. Let K be a number field of signature (r_1,r_2). Denote by κ the residue of $\zeta_K(s)$ at $s=1$, so that

$$\kappa=\frac{2^{r_1}(2\pi)^{r_2}h(K)R(K)}{w(K)\sqrt{|d(K)|}}$$

with usual notation.

   a) Using Lemma 9.2.2 and Hecke's Theorem 10.2.9, show that when $x\to\infty$ the number of K-isomorphism classes of quadratic extensions L/K with $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K))\leq x$ is asymptotic to $Q_K\cdot x$ with

$$Q_K=\frac{\kappa}{2^{r_2}\zeta_K(2)}\ .$$

   (This question is difficult and uses results from the analytic theory of number fields.)

   b) Let $S_2(K)$ be the 2-Selmer group of K (see Definition 5.2.4), and let R_1 denote an even integer such that $0\leq R_1\leq 2r_1$. Show that the number of K-isomorphism classes of quadratic extensions L/K with $\mathcal{N}_{K/\mathbb{Q}}(\mathfrak{d}(L/K))\leq x$ with R_1 real embeddings is asymptotic to $(\left|A_{R_1/2}\right|/|S_2(K)|)Q_K\cdot x$, where Q_K is as above and $A_{R_1/2}$ is the number of elements of $S_2(K)$ having exactly $R_1/2$ positive real conjugates.

   c) Assuming a reasonable regularity hypothesis, deduce from this that the number of quartic fields $L/\mathbb{Q}$ up to isomorphism with Galois group isomorphic to D_4 and absolute discriminant bounded by x is asymptotic to $C\cdot x$ for some explicit constant C (it is possible to remove the regularity hypothesis).

   d) Express C as an explicit infinite sum if we restrict to totally complex quartic fields with Galois group isomorphic to D_4 and which contain an imaginary quadratic subfield.

7. Let K be a number field, $\zeta=\zeta_3$ a primitive cube root of unity, $K_z=K(\zeta)$, and τ the generator of $\mathrm{Gal}(K_z/K)$, and assume that $\zeta\notin K$ (otherwise, see Exercise 4). Let $V_3(K_z)$ be the group of 3-virtual units of K_z (see Section 5.3.4), and set $K_{z,\tau}=\{\gamma\in K_z^*/\ \gamma^2\tau(\gamma)\in K_z^3\}$ and $V_{z,\tau}=\{\gamma\in K_z^*/\ \gamma^2\tau(\gamma)\in V_3(K_z)\}$.

   a) Show that $\beta^2\tau(\beta)$ is 3-Kummer-equivalent to $\beta'^2\tau(\beta')$ (see Definition 10.2.8) if and only if $\beta/\beta'\in K_{z,\tau}$ or $\beta\beta'\in K_{z,\tau}$.

   b) Let $\beta\in K_z^*$. Show that there exist unique ideals $\mathfrak{a}=\mathfrak{a}(\beta)$ and $\mathfrak{c}=\mathfrak{c}(\beta)$ of K_z such that

$$\beta\mathbb{Z}_{K_z}=\mathfrak{a}\frac{\mathfrak{c}^2}{\tau(\mathfrak{c})}$$

   with $\mathfrak{a}$ a primitive squarefree integral ideal not divisible by inert or ramified primes of K_z/K.

   c) Show that $\mathfrak{a}(1/\beta)=\tau(\mathfrak{a}(\beta))$ and $\mathfrak{c}(1/\beta)=(\mathfrak{c}(\beta)\mathfrak{a}(\beta)\tau(\mathfrak{a}(\beta)))^{-1}$.

d) Show that the map $\beta \mapsto \mathfrak{a}(\beta)$ induces a bijection from $K_z^*/V_{z,\tau}$ to the set I_τ of integral ideals $\mathfrak{a}$ that are primitive, squarefree, not divisible by inert or ramified primes, and such that there exists an ideal $\mathfrak{c}$ with $\mathfrak{a}\mathfrak{c}^2/\tau(\mathfrak{c})$ a principal ideal.

e) Let $S_3(K_z) = V_3(K_z)/K_z^{*3}$ be the 3-Selmer group of K_z (see Definition 5.2.4) considered as an $\mathbb{F}_3$-vector space, and denote by $S_3(K_z)[\tau - 2]$ the kernel of the map $x \mapsto \tau(x)/x^2$ from $S_3(K_z)$ to itself. Show that the map $\gamma \mapsto \gamma^2\tau(\gamma)$ induces a bijection from $V_{z,\tau}/K_{z,\tau}$ to $S_3(K_z)[\tau - 2]$.

f) Prove an analog of Lemma 9.2.2 for this case, and deduce from this and the example $\ell = 3$ of Proposition 5.3.9 an analog of Algorithm 9.2.3 for computing the list of relative cyclic cubic extensions of K.

8. Write the simple modification of Algorithm 2.3.25 required in step 1 of Algorithm 9.2.7 giving the list of ideals $\mathfrak{a}$ of K of norm less than or equal to B, such that if $\mathfrak{p} \mid \mathfrak{a}$, then $v_\mathfrak{p}(\mathfrak{a}) = 1$ if $\mathfrak{p} \nmid \ell$, while $1 \le v_\mathfrak{p}(\mathfrak{a}) \le \lfloor \ell e(\mathfrak{p}/\ell)/(\ell - 1) \rfloor + 1$ if $\mathfrak{p} \mid \ell$.

9. In the situation of steps 9 and 10 of Algorithm 9.2.7, assuming that $\zeta_3 \notin K_2$, give an algorithm that directly computes a defining polynomial for L/K from a defining polynomial for $N(\zeta_3)/K_2(\zeta_3)$, without explicitly using a defining polynomial for N/K_2.

10. Write an algorithm analogous to Algorithm 9.2.7 for computing noncyclic cubic extensions of K which uses Exercise 7 instead of class field theory for the construction of the extension L_2/K_2 (with the notation of that exercise, show that it suffices to keep the ideals $\mathfrak{a}$ of I_τ such that $\mathfrak{a}\tau(\mathfrak{a})$ comes from an ideal of K).

11. Let K be a number field of degree 6, and let $\alpha \in \mathbb{Z}_K \setminus \mathbb{Z}$ satisfy the conditions of Hunter's theorem.

a) Assume that $k = \mathbb{Q}(\alpha)$ is a quadratic subfield of K. Show directly that

$$|d(K)| \le \frac{2}{3}\left(\frac{4\,|d(L)|}{3}\right)^{1/5},$$

which is much better than the trivial bound $|d(K)| \le |d(L)|^{1/3}$.

b) Assume that $k = \mathbb{Q}(\alpha)$ is a cubic subfield of K. Show directly that

$$|d(K)| \le \frac{13}{16}\left(\frac{4\,|d(L)|}{3}\right)^{3/5}.$$

(Martinet's Theorem 9.3.3 shows that the constant 13/16 can be improved to 1/8.) What is the upper bound on $|d(L)|$ for which this is better than the bound $|d(K)| \le |d(L)|^{1/2}$? What if the constant 13/16 is replaced by 1/8?.

12. Knowing that for each integer e of the form $p_1 \dots p_t$ or $9p_1 \dots p_{t-1}$ with the p_i distinct prime numbers congruent to 1 modulo 6 there exist up to isomorphism 2^{t-1} cyclic cubic fields of discriminant e^2 (see [Coh0, Theorem 6.4.6]), prove the formula for $N_3(C_3, X)$ given in the text.

13. Give an example of a quartic field K and a quadratic subfield k of K such that the inequality of Theorem 9.3.3 is not satisfied.

14. With the notation of Section 9.3.2, set $A_\pm = a_2 - n(3n-2)/24 \pm t_2^{3/2}/3$ and $X = a_3 - (n-2)a_2/2 + n(n-1)(n-2)/24$, and assume that $a_1 = n/2$.

a) Using the bounds for a_2, show that we always have $A_+ \ge 0$.

b) Show that $A_+ - \max(A, 0) \le \min(A_+, 0) - A = -A$ if and only if $A_+ + A_- \le 0$.

c) Show that $A_- \le X \le A_+$.
d) Deduce from this that, as stated in the text, when $a_1 = n/2$ we should choose $X \ge 0$ if $a_2 \le n(3n-2)/24$, and $X \le 0$ if $a_2 \ge n(3n-2)/24$.

15. As usual, let $P(x) = x^n - a_1 x^{n-1} + \cdots + (-1)^n a_n$, let α_i be the complex roots of P, and let $T_2 = \sum_i |\alpha_i|^2$. Generalizing the inequality $|a_n| \le (T_2/n)^{n/2}$ proved in the text, show that for any real number k we have

$$P(k) \le \left(\frac{T_2 - 2ka_1 + k^2 n}{n}\right)^{n/2} .$$

16.

a) Let $(y_i)_{1 \le i \le n}$ be a family of n real numbers. Refining Schwartz's inequality, show that

$$n \sum_{1 \le i \le n} y_i^2 - \left(\sum_{1 \le i \le n} y_i\right)^2 = \sum_{1 \le i < j \le n} (y_i - y_j)^2 .$$

b) Let $P(X) = \prod_{1 \le i \le n}(X - x_i)$ be a monic polynomial with real nonzero roots x_i. Using a) and the arithmetic–geometric mean inequality, show that

$$n \sum_{1 \le i \le n} \frac{1}{x_i^2} - \left(\sum_{1 \le i \le n} \frac{1}{x_i}\right)^2 \ge \frac{n(n-1)}{2} \frac{|\mathrm{disc}(P)|^{2/(n(n-1))}}{|P(0)|^{4/n}} .$$

c) Deduce from b) the following strengthening of Proposition 9.3.8. For $1 \le k \le n-1$,

$$a_{k-1}a_{k+1} \le \frac{k(n-k)}{(k+1)(n-k+1)} a_k^2 - \frac{k(n-k)(n-k+1)}{2(n-k+1)!^{\frac{4}{k+1}}} |a_{k+1}|^{2-\frac{4}{k+1}} \left|\mathrm{disc}(P^{(n-k-1)})\right|^{\frac{2}{k(k+1)}} .$$

d) In the special case $n = 4$ and $P \in \mathbf{Z}[X]$ monic, squarefree with real roots, deduce the inequalities

$$a_1 a_3 \le \frac{4}{9} a_2^2 - \left(\frac{a_3^2}{6}\right)^{1/3} \quad \text{and}$$

$$a_2 a_4 \le \frac{3}{8} a_3^2 - \frac{3}{2} |a_4| \, |\mathrm{disc}(P)|^{1/6} \le \frac{3}{8} a_3^2 - \frac{3}{2} |a_4| .$$

17. Let P be a monic polynomial of degree n and let $M = (m_{i,j})$ be the $n \times n$ symmetric matrix such that $m_{i,j} = s_{i+j-2}$, the sum of the $(i+j-2)$th powers of the roots of P. Show the following.

a) The determinant of M is equal to $\mathrm{disc}(P)$.
b) The polynomial P is totally real if and only if the quadratic form defined by M is positive definite.
c) The polynomial P is totally real if and only if for every $k \le n$ the $k \times k$ upper-left minor extracted from M is positive.

d) The polynomial P is totally real if and only if any principal $k \times k$ minor (in other words, using the same rows and columns) extracted from M is positive.

This gives additional inequalities in the totally real case which may sometimes be useful.

18. Let $P(X) = X^n - a_1X^{n-1} + a_2X^{n-2} + \cdots + (-1)^n a_n$ be a monic polynomial with only real roots. Using the arithmetic–geometric mean inequality, show that

$$0 \le \operatorname{disc}(P) \le \left(\frac{(n-1)a_1^2 - 2na_2}{n(n-1)/2}\right)^{n(n-1)/2};$$

hence that $s_2 \ge a_1^2/n + ((n-1)/2)\operatorname{disc}(P)^{2/(n(n-1))}$. Deduce from this that if $s_2 < 2n-1$ (which is possible only when $n \ge 8$), then $\operatorname{disc}(P) \le 4^{n(n-1)/2}$.

19. Continuing the above exercise, assume that $n = 3$.
    a) Show that $s_2 \ge 5$, except perhaps if $\operatorname{disc}(P) = 49$.
    b) If $\operatorname{disc}(P) = 49$ and $s_2 \le 4$, show that $a_1 = 0$ and $a_2 = -2$ and deduce a contradiction.
    c) Find all the monic irreducible polynomials $P(X) \in \mathbb{Z}[X]$ of degree 3 such that $s_2 = 5$.
    d) Try to do the same exercise in degree 4, where the optimal inequality is $s_2 \ge 7$. (The reader is referred to [Smy1] for a better method and stronger results.)

20. Let $R(X) = X^k + aX^m + b$ be a monic trinomial, with k and m real numbers, not necessarily integers. Show that $R(X) = 0$ has at most two nonnegative real roots.

21. Let $P(X) = X^4 - a_1X^3 + a_2X^2 - a_3X + a_4$ be a quartic polynomial with only real roots. The arithmetic–geometric mean inequality shows that $|a_4| \le s_2^2/16$. Using the method of Lagrange multipliers, show that if we fix not only $s_2 = T_2$ but also a_1, we obtain the slightly stronger inequality $|a_4| \le (s_2 - a_1^2/2)^2$ (since we may assume that $0 \le a_1 \le 2$, this improvement is very slight and indicates that Proposition 9.3.10, which is a weak form of the general method of Lagrange multipliers, is not that much weaker after all).

22. As in the text, let $R_{m,2}(X) = mX^{m-n} + (n-m)X^m$ with $1 \le m \le n-1$, let $r < n$ be given, and let z_m be the unique root of $R_{m,2}(X) - r = 0$ with $0 < z_m < 1$. Show that if we set $x_0 = (m/r)^{1/(n-m)}$ and $x_{i+1} = x_i - (R_{m,2}(x_i) - r)/R'_{m,2}(x_i)$ by the usual Newton iteration, then x_i is an increasing sequence, for all i we have $x_i < z_m$, and x_i converges quadratically to z_m.

23. Show that a quartic field K is primitive if and only if the Galois group G of its Galois closure is isomorphic to A_4 or to D_4. More precisely, show that if G is isomorphic to C_4 or to D_4, then K has a unique quadratic subfield, while if G is isomorphic to $C_2 \times C_2$, then K has three quadratic subfields.

24.
    a) Show that, as claimed in the text, if we take into account only the simple general inequalities in the quartic case, the number of polynomials to consider is asymptotic to $B^{3/2}/(12\sqrt{2})$ when $B \to \infty$.
    b) For each of the three possible signatures (r_1, r_2) in the quartic case, give the corresponding asymptotic number of polynomials to consider if, in addition to the general inequalities, we consider only the specific simple inequalities for that signature (not using Pohst's method), without using the discriminant inequality.

c) Try to do the same if we also use the bounds for s_3 coming from Pohst's method and the discriminant inequality.

25. Let K be the quartic number field of signature $(2,1)$ defined by a root of the irreducible polynomial $X^4-X^3-X^2+4X-1$. Show that there exist exactly six elements $\alpha \in \mathbb{Z}_K \setminus \mathbb{Z}$ such that $\sum_j |\alpha_j|^2 \leq 1+(d/2)^{1/3}$, that among these six elements only three are such that $0 \leq a_1 \leq 2$, and that for these three elements we have $a_4 < 0$, so that contrary to what is stated in [Bu-Fo-Po], we cannot assume $a_4 > 0$.

26. Show that an imprimitive quartic field of signature $(2,1)$ has Galois group necessarily isomorphic to D_4.

27. Using the results of Section 9.3.4, state and prove inequalities analogous to those obtained in the quartic case for degree 5 fields. Look at [Diaz] and [Sc-Po-Di] if you need help.

28. Using your favorite package, for each of the polynomials defining the smallest known discriminant of totally complex fields K of degree up to 36 given in Appendix C, compute the index $[\mathbb{Z}_K : \mathbb{Z}[\theta]]$ for θ a root of the polynomial, in other words, the square root of the ratio of the polynomial discriminant to the field discriminant.

10. Appendix A: Theoretical Results

In this appendix, we regroup and prove a number of results that we need.

10.1 Ramification Groups and Applications

In most of this section we follow Serre ([Ser]) quite closely. Since we have not used any local arguments in this book, only global proofs are given, which make the proofs slightly more cumbersome.

10.1.1 A Variant of Nakayama's Lemma

We first need some technical results about modules over Dedekind domains, linked to Nakayama's lemma.

Let R be a Dedekind domain, M a finitely generated, torsion-free R-module of rank n, and N a submodule of M of finite index (or, equivalently, having the same rank). Let

$$\mathfrak{a} = \mathrm{Ann}(M,N) = \{x \in R/\ xM \subset N\}\ ,$$

and let $\mathfrak{b}$ be the index-ideal of N in M.

By the elementary divisor theorem (Theorem 1.2.30), we can write $M/N = \bigoplus_i (R/\mathfrak{d}_i)g_i$ with $\mathfrak{d}_i \mid \mathfrak{d}_{i-1}$ for $i \geq 2$. By definition, we have $\mathfrak{b} = \prod_i \mathfrak{d}_i$, and clearly we have $\mathfrak{a} = \mathfrak{d}_1$. In particular, it follows that $\mathfrak{a} \mid \mathfrak{b} \mid \mathfrak{a}^n$, so $\mathfrak{a}$ and $\mathfrak{b}$ have the same prime ideal divisors.

The following lemma is a variant of a special case of Nakayama's lemma.

Lemma 10.1.1. *Let R be a Dedekind domain, M a finitely generated, torsion-free R-module, N a submodule of M of finite index, and $\mathfrak{a} = Ann(M,N)$ and $\mathfrak{b} = [M : N]$ be as above. Finally, let $\mathfrak{p}$ be a (nonzero) prime ideal of R. The following conditions are equivalent:*

(1) $\mathfrak{p} \nmid \mathfrak{a}$;
(2) $\mathfrak{p} \nmid \mathfrak{b}$;
(3) $\mathfrak{p}M + N = M$;
(4) *the injection from N to M induces an isomorphism from $N/\mathfrak{p}N$ to $M/\mathfrak{p}M$;*

(5) $\mathfrak{p}M \cap N = \mathfrak{p}N$.

If, in addition, N is free with basis (ω_i), these conditions are also equivalent to:

(6) *the classes modulo $\mathfrak{p}M$ of the ω_i form an $R/\mathfrak{p}$-basis of $M/\mathfrak{p}M$.*

Proof. Since $\mathfrak{a} \mid \mathfrak{b} \mid \mathfrak{a}^n$, as already stated $\mathfrak{a}$ and $\mathfrak{b}$ have the same prime ideal divisors, so the equivalence of (1) and (2) is trivial.

(2) $\Longrightarrow$ (3) By (2), we have $\mathfrak{b}+\mathfrak{p} = R$, hence multiplying by M and using $\mathfrak{b}M \subset \mathfrak{a}M \subset N$, we obtain $M \subset \mathfrak{p}M + N$. The other inclusion being obvious, this proves (3).

(3) $\Longrightarrow$ (4). The natural map i from $N/\mathfrak{p}N$ to $M/\mathfrak{p}M$ induced by the injection from N to M is well-defined. By (3), it is surjective. On the other hand, both $M/\mathfrak{p}M$ and $N/\mathfrak{p}N$ are $R/\mathfrak{p}$-vector spaces of the same dimension n; hence the map i is an isomorphism.

(4) $\Longrightarrow$ (5). This simply expresses the fact that the kernel of i is trivial.

(5) $\Longrightarrow$ (1). Assume by contradiction that $\mathfrak{p} \mid \mathfrak{a}$, in other words that $\mathfrak{a} \subset \mathfrak{p}$. This implies that $\mathfrak{a}M \subset \mathfrak{p}M$. On the other hand, by definition of $\mathfrak{a}$ we have $\mathfrak{a}M \subset N$, so $\mathfrak{a}M \subset \mathfrak{p}M \cap N = \mathfrak{p}N$ by (5). It follows that $\mathfrak{a}\mathfrak{p}^{-1}M \subset N$, hence that $\mathfrak{a}\mathfrak{p}^{-1} \subset \mathfrak{a}$, a contradiction, thus proving (1). We have thus proved that conditions (1) to (5) are equivalent.

Assume now that N is free and that (ω_i) is a basis of N.

(3) $\Longleftrightarrow$ (6). It is clear that $M = N + \mathfrak{p}M$ is equivalent to the statement that the classes modulo $\mathfrak{p}M$ of the ω_i generate $M/\mathfrak{p}M$. Since this is an $R/\mathfrak{p}$-vector space of dimension n, and since there are n generating elements ω_i, this proves the equivalence of (3) and (6) and hence the lemma. □

We will mainly use this lemma with $R = \mathbb{Z}_K$, the ring of integers of a number field K, and $M = \mathbb{Z}_L$ and $N = \mathbb{Z}_K[\theta]$ for a field extension $L = K(\theta)$. Condition (2) of the lemma then simply states that $\mathbb{Z}_K[\theta]$ is a $\mathfrak{p}$-maximal order.

Lemma 10.1.2. *Let L/K be a finite extension of number fields of degree n, and let $\mathfrak{p}$ be a prime ideal of K that is totally ramified in L, so that $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}^n$. Let $\pi \in \mathfrak{P} \setminus \mathfrak{P}^2$. Then $\mathbb{Z}_K[\pi]$ is a $\mathfrak{p}$-maximal order of $\mathbb{Z}_L$.*

Proof. Since $\mathfrak{P}$ is totally ramified, $\mathbb{Z}_L/\mathfrak{P}^k$ is a $\mathbb{Z}_K/\mathfrak{p}$-vector space (of dimension k) for all $k \leq n$. We will show by induction that for all $k \leq n$, the classes modulo $\mathfrak{P}^k$ of the π^i for $0 \leq i < k$ generate $\mathbb{Z}_L/\mathfrak{P}^k$. Applied to $k = n$, this will show that the classes modulo $\mathfrak{p}\mathbb{Z}_L$ of the π^i for $0 \leq i < n$ generate $\mathbb{Z}_L/\mathfrak{p}\mathbb{Z}_L$, hence form a basis, and the lemma follows from the equivalence (2) $\Longleftrightarrow$ (6) of Lemma 10.1.1.

Since $\mathbb{Z}_L/\mathfrak{P}$ is a $\mathbb{Z}_K/\mathfrak{p}$-vector space of dimension 1, it is generated by the class of 1, so our claim is true for $k = 1$. Assume that it is true for some $k < n$, and let $x \in \mathbb{Z}_L$. Thus, there exist $x_i \in \mathbb{Z}_K$ such that $x \equiv \sum_{0 \leq i < k} x_i \pi^i \pmod{\mathfrak{P}^k}$. On the other hand, since $\pi \in \mathfrak{P} \setminus \mathfrak{P}^2$, multiplication by π^k induces

an isomorphism from $\mathbb{Z}_K/\mathfrak{p} \simeq \mathbb{Z}_L/\mathfrak{P}$ to $\mathfrak{P}^k/\mathfrak{P}^{k+1}$. Since $x - \sum_{0 \le i < k} x_i \pi^i \in \mathfrak{P}^k$, this implies that there exists $x_k \in \mathbb{Z}_K$ such that

$$x - \sum_{0 \le i < k} x_i \pi^i \equiv x_k \pi^k \pmod{\mathfrak{P}^{k+1}} ,$$

thus proving our induction hypothesis and hence the lemma. □

10.1.2 The Decomposition and Inertia Groups

From now on, we let L/K be a *normal* extension of number fields with Galois group $G = \mathrm{Gal}(L/K)$ and degree $n = [L : K]$, let $\mathfrak{p}$ be a prime ideal of K, and let

$$\mathfrak{p}\mathbb{Z}_L = \prod_{1 \le i \le g} \mathfrak{P}_i^{e_i} \quad \text{with} \quad f(\mathfrak{P}_i/\mathfrak{p}) = f_i .$$

We recall (with proof) the following basic facts (see [Marc]).

Proposition 10.1.3. *Let L/K be a normal extension of number fields as above.*

(1) *The ideals $\mathfrak{P}_i$ are permuted transitively by the Galois group G: in other words, for every pair (i, j) there exists a (not necessarily unique) $\sigma_{i,j} \in G$ such that $\sigma(\mathfrak{P}_i) = \mathfrak{P}_j$.*

(2) *All the e_i are equal (to e, say), all the f_i are equal (to f, say), and we have the equality $efg = n$.*

Proof. (1) Fix $\mathfrak{P}_i$, and assume that for some j, $\mathfrak{P}_j$ is not of the form $\sigma(\mathfrak{P}_i)$ for $\sigma \in G$. By the weak approximation theorem (Proposition 1.2.3), we can find $x \in \mathfrak{P}_j$ such that $x \notin \sigma(\mathfrak{P}_i)$ for all $\sigma \in G$. If we let $a = \mathcal{N}_{L/K}(x) = \prod_{\sigma \in G} \sigma(x)$, we have $a \in \mathbb{Z}_K \cap \mathfrak{P}_j = \mathfrak{p}$. On the other hand, for all $\sigma \in G$ we have $\sigma(x) \notin \mathfrak{P}_i$ (otherwise, $x \in \sigma^{-1}(\mathfrak{P}_i)$), and since $\mathfrak{P}_i$ is a prime ideal we have $a \notin \mathfrak{P}_i$ and in particular $a \notin \mathfrak{p}$, which is absurd.

(2) follows immediately, since (1) shows that the $\mathfrak{P}_i$ play a symmetrical role. □

We will denote by $\mathfrak{P}$ any of the prime ideals $\mathfrak{P}_i$ above $\mathfrak{p}$.

Definition 10.1.4. (1) *The* decomposition group $D(\mathfrak{P}/\mathfrak{p})$ *is the subgroup of G defined by*

$$D(\mathfrak{P}/\mathfrak{p}) = \{\sigma \in G/\ \sigma(\mathfrak{P}) = \mathfrak{P}\} .$$

(2) *The kth* ramification group $G_k(\mathfrak{P}/\mathfrak{p})$ *is the subgroup of G defined by*

$$G_k(\mathfrak{P}/\mathfrak{p}) = \{\sigma \in G/\ \forall x \in \mathbb{Z}_L,\ \sigma(x) \equiv x \pmod{\mathfrak{P}^{k+1}}\} .$$

(3) *The* inertia group *is the group $I(\mathfrak{P}/p) = G_0(\mathfrak{P}/\mathfrak{p})$, and the decomposition group $D(\mathfrak{P}/\mathfrak{p})$ will also be denoted $G_{-1}(\mathfrak{P}/\mathfrak{p})$.*

If $\sigma_{i,j}$ is any element of G such that $\sigma_{i,j}(\mathfrak{P}_i) = \mathfrak{P}_j$, it is clear that the set of all such σ is the coset $\sigma_{i,j}D(\mathfrak{P}_i/\mathfrak{p})$ and also that

$$D(\mathfrak{P}_j/\mathfrak{p}) = \sigma_{i,j}D(\mathfrak{P}_i/\mathfrak{p})\sigma_{i,j}^{-1} \ .$$

More generally, this is true for all the ramification groups G_k.

If $\sigma \in D(\mathfrak{P}/\mathfrak{p})$, then $\sigma(\mathfrak{P}) = \mathfrak{P}$ (by definition) and σ fixes $\mathbb{Z}_K$ pointwise. Thus σ induces a field isomorphism $s(\sigma)$ from $\mathbb{Z}_L/\mathfrak{P}$ into itself which leaves $\mathbb{Z}_K/\mathfrak{p}$ pointwise fixed, hence $s(\sigma) \in \mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$.

Proposition 10.1.5. (1) *The map s defined above is a surjective homomorphism from $D(\mathfrak{P}/\mathfrak{p})$ to $\mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$ whose kernel is equal to $I(\mathfrak{P}/\mathfrak{p})$. In other words, s induces an isomorphism from $D(\mathfrak{P}/\mathfrak{p})/I(\mathfrak{P}/\mathfrak{p})$ to $\mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$.*

(2) *There exists $\sigma_{\mathfrak{P}} \in D(\mathfrak{P}/\mathfrak{p})$ such that for all $x \in \mathbb{Z}_L$*

$$\sigma_{\mathfrak{P}}(x) \equiv x^{\mathcal{N}(\mathfrak{p})} \pmod{\mathfrak{P}}$$

and $\sigma_{\mathfrak{P}}$ is defined uniquely up to conjugation by an element of $I(\mathfrak{P}/\mathfrak{p})$.

Proof. (1) The finite field extension $(\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p})$ is Galois (that is, normal and separable, and even cyclic). Hence, in particular, the primitive element theorem applies, so we can find $\overline{\theta} \in \mathbb{Z}_L/\mathfrak{P}$ such that $\mathbb{Z}_L/\mathfrak{P} = (\mathbb{Z}_K/\mathfrak{p})(\overline{\theta})$. By the weak approximation theorem, we can find $\theta \in \mathbb{Z}_L$ such that $\theta \equiv \overline{\theta} \pmod{\mathfrak{P}}$ and $\theta \in \sigma(\mathfrak{P})$ for all $\sigma \notin D(\mathfrak{P}/\mathfrak{p})$. This is possible since by definition of $D(\mathfrak{P}/\mathfrak{p})$ all the $\sigma(\mathfrak{P})$ for $\sigma \notin D(\mathfrak{P}/\mathfrak{p})$ are distinct from $\mathfrak{P}$. Set

$$P(X) = \prod_{\sigma \in G} (X - \sigma(\theta)) \ ,$$

and let $\overline{P}(X)$ be the polynomial in $(\mathbb{Z}_L/\mathfrak{P})[X]$ obtained by reducing the coefficients of P modulo $\mathfrak{P}$. Since $\sigma(\theta) \in \mathfrak{P}$ for all $\sigma \notin D(\mathfrak{P}/\mathfrak{p})$, it follows that for $m = |G| - |D(\mathfrak{P}/\mathfrak{p})|$ we have

$$\overline{P}(X) = X^m \overline{P_1}(X) \quad \text{with} \quad \overline{P_1}(X) = \prod_{\sigma \in D(\mathfrak{P}/\mathfrak{p})} (X - \sigma(\theta)) \ .$$

By Galois theory $\overline{P}(X) \in (\mathbb{Z}_K/\mathfrak{p})[X]$, hence $\overline{P_1}(X) \in (\mathbb{Z}_K/\mathfrak{p})[X]$, so $\overline{\theta}$ is a root of $\overline{P_1}(X)$, hence its minimal polynomial is a divisor of $\overline{P_1}$, so the conjugates of $\overline{\theta}$ are among the $\sigma(\overline{\theta})$ for $\sigma \in D(\mathfrak{P}/\mathfrak{p})$. This means exactly that the homomorphism s is surjective.

The kernel of s is the set of $\sigma \in D(\mathfrak{P}/\mathfrak{p})$ such that $s(\sigma)$ is the identity on $\mathbb{Z}_L/\mathfrak{P}$, in other words such that $\sigma(x) \equiv x \pmod{\mathfrak{P}}$ for all $x \in \mathbb{Z}_L$, and by definition this is the inertia group $I(\mathfrak{P}/\mathfrak{p})$, finishing the proof of (1).

The proof of (2) is a simple and well-known property of finite fields. It is clear that the map $x \mapsto x^{\mathcal{N}(\mathfrak{p})}$, called the *Frobenius homomorphism*, is an

element of $\mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$ (here $\mathcal{N}(\mathfrak{p})$ is the absolute norm of $\mathfrak{p}$, in other words the cardinality of $\mathbb{Z}_K/\mathfrak{p}$). It is also not difficult to prove that $\mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$ is cyclic and generated by the Frobenius homomorphism, which is therefore of order exactly equal to $f = f(\mathfrak{P}/\mathfrak{p})$. Statement (2) follows from immediately from this. □

Corollary 10.1.6. *We have*

$$|D(\mathfrak{P}/\mathfrak{p})| = e(\mathfrak{P}/\mathfrak{p})f(\mathfrak{P}/\mathfrak{p}) \quad \textit{and} \quad |I(\mathfrak{P}/\mathfrak{p})| = e(\mathfrak{P}/\mathfrak{p}) \ .$$

Proof. Since there are g ideals $\mathfrak{P}_i$ above $\mathfrak{p}$ and the $D(\mathfrak{P}_i/\mathfrak{p})$ are conjugate groups, and hence have the same cardinality, we have $D(\mathfrak{P}/\mathfrak{p}) = n/g = e(\mathfrak{P}/\mathfrak{p})f(\mathfrak{P}/p)$. The cardinality of $I(\mathfrak{P}/\mathfrak{p})$ then follows from the proposition. □

Corollary 10.1.7. *If L/K is a normal extension of number fields which is not cyclic, no prime ideal of K is inert in L/K.*

Proof. Indeed, if $\mathfrak{p}$ is inert, then $D(\mathfrak{P}/\mathfrak{p}) = \mathrm{Gal}(L/K)$ and $I(\mathfrak{P}/\mathfrak{p}) = \{1\}$, hence Proposition 10.1.5 shows that $\mathrm{Gal}(L/K) = D(\mathfrak{P}/\mathfrak{p})/I(\mathfrak{P}/p)$ is cyclic. Note that the converse is true: if L/K is a cyclic extension, there exist inert primes in L/K, even a positive density. This follows immediately from the Chebotarev density theorem (see [Lan3]). □

From now on, we consider $\mathfrak{P}$ fixed, and we write D, I, G_k, e, f instead of $D(\mathfrak{P}/\mathfrak{p})$, $I(\mathfrak{P}/\mathfrak{p})$, $G_k(\mathfrak{P}/\mathfrak{p})$, $e(\mathfrak{P}/\mathfrak{p})$, $f(\mathfrak{P}/\mathfrak{p})$, respectively.

Let L^D and L^I be the fixed fields of L by D and I, respectively. The inclusions $\{1_G\} \subset I \subset D \subset G$ lead to the following diagram of fields, where we have also written the corresponding prime ideals and some relations, whose easy proofs are left to the reader (Exercise 1).

$$\begin{array}{ccc}
L & & \mathfrak{P} = \mathfrak{P}_I^e \\
{\scriptstyle e}\,\big| & & \\
L^I & & \mathfrak{P}_I = \mathfrak{P}_D\mathbb{Z}_{L^I} \\
{\scriptstyle f}\,\big| & & \\
L^D & & \mathfrak{P}_D \mid \mathfrak{p}\mathbf{Z}_{L^D} \\
{\scriptstyle g}\,\big| & & \\
K & & \mathfrak{p}
\end{array}$$

Remark. We have $f(\mathfrak{P}_D/\mathfrak{p}) = 1$, in other words $\mathfrak{P}_D$ is of degree 1 over $\mathfrak{p}$, but it is *not* true in general that $f(\mathfrak{P}'_D/\mathfrak{p}) = 1$ for the other prime ideals $\mathfrak{P}'_D$ of L^D above $\mathfrak{p}$ (see Exercise 1).

10.1.3 Higher Ramification Groups

In this section, we follow [Ser] essentially verbatim. We start with a trivial lemma.

Lemma 10.1.8. *Recall that $D = G_{-1}$ and $I = G_0$.*

(1) *We have the inclusions*

$$G \supset G_{-1} \supset G_0 \supset G_1 \cdots \ .$$

(2) *There exists n such that $G_k = \{1_G\}$ for all $k \geq n$.*

Proof. (1) is trivial. Furthermore, if $\sigma \in G_k$ for all k, then for all $x \in \mathbb{Z}_L$ and for all k, $\sigma(x) \equiv x \pmod{\mathfrak{P}^{k+1}}$, so $\sigma(x) = x$ and hence $\sigma = 1_G$. Since G is finite, it follows that $G_k = \{1_G\}$ for k sufficiently large. □

Lemma 10.1.9. *For all $k \geq 0$, G_k is a normal subgroup of $D = G_{-1}$ (not of G itself, however).*

Proof. Let $\sigma \in G_k$ and $\tau \in D$. For all $x \in \mathbb{Z}_L$, we have

$$\sigma(\tau^{-1}(x)) \equiv \tau^{-1}(x) \pmod{\mathfrak{P}^{k+1}} \ ;$$

hence

$$\tau(\sigma(\tau^{-1}(x))) \equiv x \pmod{\tau(\mathfrak{P})^{k+1}} \ ,$$

and since $\tau \in D$, we have $\tau(\mathfrak{P}) = \mathfrak{P}$, so the result follows. □

By definition, $\mathrm{Gal}(L/L^I) = I$ and $\mathrm{Gal}(L/L^D) = D$. By this lemma, the extension L^I/L^D is normal, with Galois group isomorphic to D/I and hence to $\mathrm{Gal}((\mathbb{Z}_L/\mathfrak{P})/(\mathbb{Z}_K/\mathfrak{p}))$. Thus it is even a cyclic extension of degree f.

Lemma 10.1.10. *Keep the above notation. Then for all $k \geq 0$ we have $G_k(\mathfrak{P}/\mathfrak{P}_I) = G_k(\mathfrak{P}/\mathfrak{p})$ and we also have $v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(\mathfrak{D}(L/L^I))$, where as usual $\mathfrak{D}(L/K)$ denotes the relative different (see Definition 2.3.16).*

Proof. Since $I = G_0$, it is clear that for $k \geq 0$, $\sigma \in G_k(\mathfrak{P}/\mathfrak{p})$ if and only if $\sigma \in G_k(\mathfrak{P}/\mathfrak{P}_I)$, so the groups are equal. To prove the second statement, we note that the extension L^I/K is *unramified* at $\mathfrak{P}_I$, hence $\mathfrak{P}_I$ and hence also $\mathfrak{P}$ do not divide the relative different $\mathfrak{D}(L^I/K)$. By the transitivity formula for the different (see Proposition 2.3.17), it follows that $v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(\mathfrak{D}(L/L^I))$, as was to be proved. □

Since $\mathfrak{P}_I$ is totally ramified in the extension L/L^I, this lemma shows that to study higher ramification groups (i.e., the G_k for $k \geq 0$), we can restrict to normal extensions L/K and prime ideals $\mathfrak{p}$ such that $\mathfrak{p}$ is totally ramified in L, in other words such that $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}^e$. It also shows that the computation of the relative different can also be reduced to that case. We will use this reduction in several places, particularly in the proof of Theorem 10.1.22 below.

Lemma 10.1.11. *Assume that $\mathfrak{p} = \mathfrak{P}^e$ is totally ramified in L/K and let $\pi \in \mathfrak{P} \setminus \mathfrak{P}^2$ be a uniformizer at $\mathfrak{P}$. If $\sigma \in G = \mathrm{Gal}(L/K)$, then $\sigma \in G_k$ if and only if $\sigma(\pi) \equiv \pi \pmod{\mathfrak{P}^{k+1}}$.*

Proof. The condition is evidently necessary. Conversely, if it is satisfied, then $\sigma(x) \equiv x \pmod{\mathfrak{P}^{k+1}}$ for all $x \in \mathbb{Z}_K[\pi]$. Let y be an arbitrary element of $\mathbb{Z}_L$. By Lemma 10.1.2 we know that $\mathbb{Z}_K[\pi]$ is a $\mathfrak{p}$-maximal order, so it follows that the index-ideal $\mathfrak{b} = [\mathbb{Z}_L : \mathbb{Z}_K[\pi]]$ is not divisible by $\mathfrak{p}$. By the weak approximation theorem we can find $d \in \mathbb{Z}_K$ such that $v_{\mathfrak{p}}(d) = 0$ and $v_{\mathfrak{q}}(d) \geq v_{\mathfrak{q}}(\mathfrak{b})$ for $\mathfrak{q} \mid \mathfrak{b}$. Thus $dy \in \mathfrak{b}y \subset \mathbb{Z}_K[\pi]$, hence $\sigma(dy) \equiv dy \pmod{\mathfrak{P}^{k+1}}$. Since $d \in \mathbb{Z}_K$ we have $\sigma(d) = d$, and since d is coprime to $\mathfrak{p}$ and hence to $\mathfrak{P}$, we deduce that $\sigma(y) \equiv y \pmod{\mathfrak{P}^{k+1}}$, as desired. □

When $\mathfrak{p}$ is totally ramified, thanks to this lemma, to check if $\sigma \in G_k$ it is sufficient to check the congruence on the single element π.

Definition and Proposition 10.1.12. *Keep all the above notation. For all $\sigma \in G = \mathrm{Gal}(L/K)$, we define*

$$i_G(\sigma) = v_{\mathfrak{P}}(\sigma(\pi) - \pi) \ ,$$

where by convention $v_{\mathfrak{P}}(0) = +\infty$.

Then this does not depend on the choice of the uniformizer π, and in fact $\sigma \in G_k$ if and only if $i_G(\sigma) \geq k+1$, in other words, $i_G(\sigma) = k+1$ if and only if $\sigma \in G_k \setminus G_{k+1}$.

Proof. By Lemma 10.1.11, we have $i_G(\sigma) \geq k+1$ if and only if $\sigma \in G_k$, and since this condition is itself independent of π, the result follows. □

Proposition 10.1.13. *Let $\mathfrak{p}$ be a prime ideal of K which is not necessarily totally ramified in L/K, let π be a uniformizer of $\mathfrak{P}$, and let $\sigma \in G_0 = I$.*

(1) *We have $v_{\mathfrak{P}}(\sigma(\pi)/\pi) = 0$.*
(2) *For $k \geq 1$, $\sigma \in G_k \iff \sigma(\pi)/\pi \equiv 1 \pmod{{}^*\mathfrak{P}^k}$.*

Note that since $\sigma(\pi)/\pi$ is not an algebraic integer in general, the congruence is multiplicative, in other words (2) reads $\sigma \in G_k \iff v_{\mathfrak{P}}(\sigma(\pi)/\pi - 1) \geq k$.

Proof. Since π is a uniformizer of $\mathfrak{P}$, we have

$$v_{\mathfrak{P}}(\sigma(\pi)/\pi) = v_{\sigma^{-1}(\mathfrak{P})}(\pi) - v_{\mathfrak{P}}(\pi) = 0$$

since σ is in the inertia group, hence a fortiori in the decomposition group, proving (1).

For (2), by Lemma 10.1.10, we have $G_k(\mathfrak{P}/\mathfrak{P}_I) = G_k(\mathfrak{P}/\mathfrak{p})$. Since we assume that $\sigma \in G_0 = I$, this implies that we can replace the extension L/K by the extension L/L^I without changing the ramification groups, and in this extension $\mathfrak{P}_I$ is totally ramified. We clearly have

$$v_{\mathfrak{P}}(\sigma(\pi)/\pi - 1) = v_{\mathfrak{P}}(\sigma(\pi) - \pi) - 1 = i_G(\sigma) - 1 \ ;$$

hence the definition shows that $\sigma \in G_k$ if and only if $i_G(\sigma) \geq k+1$ if and only if $\sigma(\pi)/\pi \equiv 1 \pmod{{}^*\mathfrak{P}^k}$, as was to be proved. □

Proposition 10.1.14. *As above, let π be a uniformizer of $\mathfrak{P}$.*

(1) *The map that sends $\sigma \in G_0$ to $\sigma(\pi)/\pi$ induces an injection θ_0 from G_0/G_1 to $(\mathbb{Z}_L/\mathfrak{P})^*$.*
(2) *For $k \geq 1$ the map that sends $\sigma \in G_k$ to $\sigma(\pi)/\pi - 1$ induces an injection θ_k from G_k/G_{k+1} to $\mathfrak{P}^k/\mathfrak{P}^{k+1}$.*
(3) *The maps θ_k do not depend on the choice of the uniformizer π.*

Proof. First note that if $v_{\mathfrak{P}}(\alpha) \geq k$, by Lemma 1.2.31 we may write $\alpha = x/d$ with $v_{\mathfrak{P}}(d) = 0$, hence with $v_{\mathfrak{P}}(x) \geq k$. It follows that we can send α to $\mathfrak{P}^k/\mathfrak{P}^{k+1}$ by sending it to the class of xd^{-1}, where d^{-1} is the inverse of d modulo $\mathfrak{P}^{k+1}$, and this clearly does not depend on the chosen representation of α.

On the other hand, proposition 10.1.13 shows that for $k \geq 1$, $\sigma(\pi)/\pi \equiv 1 \pmod{{}^*\mathfrak{P}^k}$ if and only if $\sigma \in G_k$, showing both that for $k \geq 1$ the maps θ_k are well-defined and that they are injective. For $k = 0$, a similar reasoning shows that the map θ_0 is well-defined, and $\sigma(\pi)/\pi \equiv 1 \pmod{{}^*\mathfrak{P}}$ is equivalent to $\sigma \in G_1$ by Proposition 10.1.13, so the map is also injective in this case. We leave the proof of (3) to the reader. □

Corollary 10.1.15. *Let p be the prime number below $\mathfrak{p}$ and $\mathfrak{P}$ or, equivalently, the characteristic of the residue fields $\mathbb{Z}_K/\mathfrak{p}$ and $\mathbb{Z}_L/\mathfrak{P}$. Then we have the following.*

(1) *The group G_0/G_1 is a cyclic group of order prime to p, and in fact $|G_0/G_1| = e/(p^\infty, e)$, the prime to p part of the ramification index $e = e(\mathfrak{P}/\mathfrak{p})$.*
(2) *For all $k \geq 1$, G_k/G_{k+1} is an Abelian group isomorphic to a product of copies of $\mathbb{Z}/p\mathbb{Z}$ (in other words, an* elementary p-group*). In particular, G_1 is a (not necessarily Abelian) p-group, and $|G_1| = (p^\infty, e)$.*

Proof. Since $(\mathbb{Z}_L/\mathfrak{P})^*$ is the multiplicative group of a finite field with $p^{f(\mathfrak{P}/p)}$ elements, it is a cyclic group of order $p^{f(\mathfrak{P}/p)} - 1$. Proposition 10.1.14 thus implies that G_0/G_1 is a subgroup of this group, hence is a cyclic group of order dividing $p^{f(\mathfrak{P}/p)} - 1$, and in particular of order prime to p. Since we know that $|G_0| = e = e(\mathfrak{P}/\mathfrak{p})$ and we will see in (2) that G_1 is a p-group, (1) follows.

Similarly, Proposition 10.1.14 shows that for $k \geq 1$, G_k/G_{k+1} is isomorphic to a subgroup of the additive group $\mathfrak{P}^k/\mathfrak{P}^{k+1}$, which is (noncanonically) isomorphic to $\mathbb{Z}_L/\mathfrak{P}$ as we saw in Section 4.2.3. Since this is the additive group of a finite field, it is an $\mathbb{F}_p$-vector space and hence an elementary p-group, finishing the proof of the corollary. □

In the same way that the extension L^I/K is the maximal subextension of L in which $\mathfrak{p}$ is unramified, the above corollary implies that the extension L^{G_1}/K is the maximal subextension of L in which $\mathfrak{p}$ is *tamely ramified*, in other words, such that $e(\mathfrak{P}_{G_1}/\mathfrak{p})$ is coprime to p.

Proposition 10.1.16. *Let $\sigma \in G_0$ and let $\tau \in G_k$ for some $k \geq 1$. Then in $\mathfrak{P}^k/\mathfrak{P}^{k+1}$ we have*

$$\theta_k(\sigma\tau\sigma^{-1}) = \theta_0(\sigma)^k\theta_k(\tau) \ .$$

Proof. Note first that this formula makes sense, since $\theta_0(\sigma)^k \in (\mathbf{Z}_L/\mathfrak{P})^*$, which operates multiplicatively on the $\mathbf{Z}_L/\mathfrak{P}$-vector space $\mathfrak{P}^k/\mathfrak{P}^{k+1}$.

Set $\pi_1 = \sigma^{-1}(\pi)$. Since σ is in the decomposition group, π_1 is also a uniformizer of $\mathfrak{P}$. Thus by definition, if we set $a = \tau(\pi_1)/\pi_1 - 1$, we have $v_{\mathfrak{P}}(a) \geq k$ and the class of a modulo $\mathfrak{P}^{k+1}$ is equal to $\theta_k(\tau)$. On the other hand, we have

$$\sigma(a) = \sigma(\tau(\pi_1))/\sigma(\pi_1) - 1 = \sigma\tau\sigma^{-1}(\pi)/\pi - 1 \ ,$$

so $\theta_k(\sigma\tau\sigma^{-1})$ is the class of $\sigma(a)$ modulo $\mathfrak{P}^{k+1}$.

Finally, if we set $b = a/\pi^k$, we have $v_{\mathfrak{P}}(b) \geq 0$, and since $\sigma \in G_0$ we have $\sigma(b) \equiv b \pmod{{}^*\mathfrak{P}}$, so

$$\begin{aligned}\sigma(a) = (\sigma(\pi)/\pi)^k\sigma(b)\pi^k &\equiv \theta_0(\sigma)^k b\pi^k \\ &\equiv \theta_0(\sigma)^k a \equiv \theta_0(\sigma)^k\theta_k(\tau) \pmod{{}^*\mathfrak{P}^{k+1}} \ ,\end{aligned}$$

proving the proposition. □

Corollary 10.1.17. *Let $\sigma \in G_0$ and let $\tau \in G_k$ for some $k \geq 1$. Then $\sigma\tau\sigma^{-1}\tau^{-1} \in G_{k+1}$ if and only if either $\tau \in G_{k+1}$ or $\sigma^k \in G_1$.*

Proof. Indeed, by the above proposition

$$\begin{aligned}\sigma\tau\sigma^{-1}\tau^{-1} \in G_{k+1} &\iff \sigma\tau\sigma^{-1} \in \tau G_{k+1} \\ &\iff \theta_k(\sigma\tau\sigma^{-1}) = \theta_k(\tau) \\ &\iff \theta_0(\sigma)^k\theta_k(\tau) = \theta_k(\tau) \\ &\iff \theta_k(\tau)(\theta_0(\sigma)^k - 1) = 0 \\ &\iff \theta_k(\tau) = 0 \text{ or } \theta_0(\sigma^k) = 1 \\ &\iff \tau \in G_{k+1} \text{ or } \sigma^k \in G_1 \ ,\end{aligned}$$

proving the corollary. □

Corollary 10.1.18. *Let $e_0 = e/(p^\infty, e)$ be the order of G_0/G_1.*

(1) *If G_0 is an Abelian group, then for $k \geq 0$, $G_k \neq G_{k+1}$ implies $e_0 \mid k$.*

(2) *If G_0 is a dihedral group D_n with n odd, then $G_1 = C_n$ and for $k \geq 1$, $G_k \neq G_{k+1}$ implies k odd (note that this is the opposite of the condition in* (1) *if $e_0 = 2$).*

Proof. Let $\sigma \in G_0$ be such that the image of σ in G_0/G_1 generates the cyclic group G_0/G_1.

(1). Assume that $e_0 \nmid k$, and in particular $k \geq 1$. This means that $\sigma^k \notin G_1$. It follows that for any $\tau \in G_k$, $\sigma\tau\sigma^{-1}\tau^{-1} \in G_{k+1}$ if and only if $\tau \in G_{k+1}$. Since G_0 is Abelian, $\sigma\tau\sigma^{-1}\tau^{-1} = 1_G$, so $\tau \in G_{k+1}$ for all $\tau \in G_k$, and hence $G_k = G_{k+1}$.

(2). If $G_0 = D_n$ with n odd, G_1 must be a normal subgroup of G_0 of prime power order such that G_0/G_1 is cyclic, and it is easily seen that the only such subgroup is C_n (see Exercise 2).

Assume that k is even. It follows that $\sigma^k = (\sigma^2)^{k/2} \in G_1$, hence by Corollary 10.1.17, $\tau \in G_k \implies \sigma\tau\sigma^{-1}\tau^{-1} \in G_{k+1}$. Since $\tau \in G_k \subset G_1 = C_n$, $\sigma\tau\sigma^{-1}\tau^{-1} = \tau^{-2}$, hence $\tau \in G_k \implies \tau^2 \in G_{k+1}$, and hence $\tau \in G_{k+1}$ since G_k is of odd order, so $G_k = G_{k+1}$, as claimed. □

Remark. Since G_1 must be a p-group, when $G_0 = D_n$ with n odd, n is necessarily a prime power.

We refer to [Ser] for a more detailed study of the ramification groups (see also Exercises 3 and 4). We now use them for our specific needs.

10.1.4 Application to Different and Conductor Computations

Definition 10.1.19. *For any sub-$\mathbb{Z}_K$-module M of L, we set (when there is no risk of confusion with other notation)*

$$M^* = \{x \in L / \ \mathrm{Tr}_{L/K}(xM) \subset \mathbb{Z}_K\} \ .$$

For example, $\mathfrak{D}(L/K) = \mathbb{Z}_L^*$.

Lemma 10.1.20. *Let π be as above, and let f be the minimal polynomial of π in $\mathbb{Z}_K[X]$. Then*

$$\mathbb{Z}_K[\pi]^* = \frac{1}{f'(\pi)}\mathbb{Z}_K[\pi] \ .$$

Proof. Set $a_{i,j} = \mathrm{Tr}_{L/K}(\pi^i\pi^j/f'(\pi))$, and let A be the $n \times n$ matrix whose entries are the $a_{i,j}$. For any $x \in L$ we can write $x = \sum_i x_i\pi^i/f'(\pi)$ with x_i in K but not necessarily in $\mathbb{Z}_K$. If X is the column vector of the x_i, the condition $x \in \mathbb{Z}_K[\pi]^*$ is equivalent to $AX \in \mathbb{Z}_K^n$, hence to prove the lemma we must prove that $A \in \mathrm{GL}_n(\mathbb{Z}_K)$ — in other words, that A has entries in $\mathbb{Z}_K$ and determinant equal to a unit.

To do this, we first note the rational function identity

$$\frac{1}{f(X)} = \sum_{k=1}^{n} \frac{1}{f'(\alpha_k)(X - \alpha_k)} \ ,$$

where the α_k are the roots of $f(X)$ in $\overline{\mathbb{Q}}$ (see Exercise 5).

Since $f(X)$ is a monic polynomial of degree n, if we expand this identity in powers of $1/X$, all the coefficients of $(1/X)^i$ for $i < n$ vanish and the coefficient of $(1/X)^n$ is equal to 1. Looking at the right-hand side, this gives

$$\sum_{1\le k\le n} \frac{\alpha_k^i}{f'(\alpha_k)} = \begin{cases} 0 & \text{if } i \le n-2, \\ 1 & \text{if } i = n-1; \end{cases}$$

in other words, $\mathrm{Tr}_{L/K}(\pi^i/f'(\pi)) = 0$ for $i \le n-2$ and $\mathrm{Tr}_{L/K}(\pi^{n-1}/f'(\pi)) = 1$.

Let us now prove that $A \in \mathrm{GL}_n(\mathbb{Z}_K)$. Thanks to the above formula, we already see that $a_{i,j} = \mathrm{Tr}_{L/K}(\pi^{i+j}/f'(\pi))$ is equal to 0 for $i+j \le n-2$ and is equal to 1 for $i+j = n-1$. For $i+j \ge n$ we have

$$a_{i,j} = \mathrm{Tr}_{L/K}(\pi^n \pi^{i+j-n}/f'(\pi)) \ .$$

Since f is monic, π^n is equal to a $\mathbb{Z}_K$-linear combination of the π^k for $k < n$, and this shows by induction that $a_{i,j} \in \mathbb{Z}_K$ for all i and j. Finally, up to reversal of rows or columns A is a triangular matrix, and clearly $\det(A) = (-1)^{n(n-1)/2}$ is indeed a unit, proving our claim and the lemma. □

Corollary 10.1.21. *With the above notation,*

$$v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(f'(\pi)) \ .$$

Proof. As in Section 10.1.1, set

$$\mathfrak{a} = \mathrm{Ann}(\mathbb{Z}_L, \mathbb{Z}_K[\pi]) = \{x \in \mathbb{Z}_K / \ x\mathbb{Z}_L \subset \mathbb{Z}_K[\pi]\} \ .$$

Since $\mathbb{Z}_K[\pi]$ is $\mathfrak{p}$-maximal, we know that $\mathfrak{p} \nmid \mathfrak{a}$.

On the other hand,

$$\begin{aligned} x \in \mathfrak{a} &\iff x\mathbb{Z}_L \subset \mathbb{Z}_K[\pi] && \iff x\mathbb{Z}_L/f'(\pi) \subset \mathbb{Z}_K[\pi]^* \\ &\iff \mathrm{Tr}_{L/K}(x\mathbb{Z}_L/f'(\pi)) \subset \mathbb{Z}_K && \iff x/f'(\pi) \in \mathfrak{D}(L/K)^{-1} \\ &\iff x \in f'(\pi)\mathfrak{D}(L/K)^{-1} \ . \end{aligned}$$

It follows that $\mathfrak{D}(L/K) = f'(\pi)\mathfrak{a}^{-1}$, and since $\mathfrak{p} \nmid \mathfrak{a}$, the result follows. □

We are now ready to prove the main theorem that we need, which involves the higher ramification groups.

Theorem 10.1.22. *Let L/K be a normal extension of number fields, set $G = \mathrm{Gal}(L/K)$, let $\mathfrak{P}$ be a prime ideal of L, and let $\mathfrak{p}$ be the prime ideal of K below $\mathfrak{P}$. The valuation at $\mathfrak{P}$ of the relative different is given by the formula*

$$v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = \sum_{k\ge 0}(|G_k| - 1) \ .$$

By Lemma 10.1.8, the apparently infinite sum on k has in fact only a finite number of nonzero terms, so the sum makes sense. Note also that since this formula gives the $\mathfrak{P}$-adic valuation of the different for all $\mathfrak{P}$, it gives the different itself.

Proof. By Lemma 10.1.10, if we set $K' = L^I$, the ideal $\mathfrak{P}$ has the same ramification groups (for $k \geq 0$) in the extension L/K' as in the extension L/K, and $v_{\mathfrak{P}}(\mathfrak{D}(L/K')) = v_{\mathfrak{P}}(\mathfrak{D}(L/K))$. We may thus assume that $\mathfrak{p}$ is totally ramified in the extension L/K, so that we can apply the results proved in that case.

Thus, let $\pi \in \mathfrak{P} \setminus \mathfrak{P}^2$ with minimal monic polynomial $f(X) \in \mathbb{Z}_K[X]$, so that we know that $\mathbb{Z}_K[\pi]$ is a $\mathfrak{p}$-maximal order in $\mathbb{Z}_L$ (by Lemma 10.1.2) and that $v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(f'(\pi))$ (by Corollary 10.1.21). Since $f(\pi) = \prod_{\sigma \in G}(X - \sigma(\pi))$, we have

$$f'(\pi) = \prod_{\sigma \in G \setminus \{1_G\}} (\pi - \sigma(\pi)) \ ,$$

so that

$$v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(f'(\pi)) = \sum_{\sigma \in G \setminus \{1_G\}} i_G(\sigma)$$

by definition of the function $i_G(\sigma)$.

To finish the proof, set $g_k = |G_k| - 1$ and let n be such that $G_n = \{1_G\}$. If $\sigma \in G_{k-1} \setminus G_k$, we have by definition $i_G(\sigma) = k$. It follows that

$$\sum_{\sigma \in G \setminus \{1_G\}} i_G(\sigma) = \sum_{k=0}^{n} k(g_{k-1} - g_k) = \sum_{k=0}^{n} g_k(k+1-k) = \sum_{k=0}^{n} g_k \ ,$$

as was to be proved. Note that in general the intermediate result

$$v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = \sum_{\sigma \in G \setminus \{1_G\}} i_G(\sigma)$$

that we have found is only valid when $\mathfrak{p}$ is totally ramified in L/K. □

As an application of the above result, we prove the following.

Proposition 10.1.23. *Let L/K be a cyclic extension of number fields of prime degree ℓ, let $\mathfrak{p}$ be a prime ideal of K above ℓ, and denote as usual by G_k the ramification groups of $\mathfrak{P}/\mathfrak{p}$ for any prime ideal $\mathfrak{P}$ of L above $\mathfrak{p}$. If k_0 is the largest k such that $G_k \neq \{1\}$, then $k_0 \leq \lfloor \ell e(\mathfrak{p}/\ell)/(\ell-1) \rfloor$.*

Proof. If $\mathfrak{p}$ is unramified in L/K, we have $G_0 = \{1\}$, so the inequality is trivial. Thus, since ℓ is prime, we may assume that $\mathfrak{p}$ is totally ramified in L/K, so that $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}^\ell$. For simplicity set $e = e(\mathfrak{p}/\ell) = v_{\mathfrak{p}}(\ell)$. By Theorem 10.1.22, since $|G_k| = \ell$ for $k \leq k_0$ and $|G_k| = 1$ for $k > k_0$, we have $v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = (k_0+1)(\ell-1)$. Hence

$$\begin{aligned} k_0 > \lfloor \ell e/(\ell-1) \rfloor &\Longrightarrow k_0 > \ell e/(\ell-1) \Longrightarrow (k_0+1)(\ell-1) > \ell e + \ell - 1 \\ &\Longrightarrow (k_0+1)(\ell-1) \geq \ell(e+1) \Longrightarrow \mathfrak{P}^{\ell(e+1)} \mid \mathfrak{D}(L/K) \\ &\Longrightarrow \mathfrak{p}^{e+1} \mid \mathfrak{D}(L/K) \Longrightarrow \mathfrak{p}^{-(e+1)}\mathbb{Z}_L \subset \mathfrak{D}^{-1}(L/K) \\ &\Longrightarrow \mathrm{Tr}_{L/K}(\mathfrak{p}^{-(e+1)}\mathbb{Z}_L) \subset \mathbb{Z}_K \\ &\Longrightarrow \mathfrak{p}^{-(e+1)}\,\mathrm{Tr}_{L/K}(\mathbb{Z}_L) \subset \mathbb{Z}_K \Longrightarrow \mathrm{Tr}_{L/K}(\mathbb{Z}_L) \subset \mathfrak{p}^{e+1} \\ &\Longrightarrow \ell = \mathrm{Tr}_{L/K}(1) \in \mathfrak{p}^{e+1} \Longrightarrow e = v_{\mathfrak{p}}(\ell) \geq e+1\ , \end{aligned}$$

and this contradiction proves the proposition. □

Corollary 10.1.24. *Let L/K be a cyclic extension of number fields of prime degree ℓ and let $\mathfrak{p}$ be a prime ideal of K above ℓ. If $\mathfrak{f}$ denotes the conductor of L/K, then $v_{\mathfrak{p}}(\mathfrak{f}) \leq \lfloor \ell e(\mathfrak{p}/\ell)/(\ell-1) \rfloor + 1$, and this upper bound is best possible.*

Proof. Indeed, we may assume that $\mathfrak{p} \mid \mathfrak{f}$, so $\mathfrak{p}$ is ramified, hence totally ramified in L/K. Thus, using the same notation as that of the proposition, we have $v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = (\ell-1)(k_0+1)$. On the other hand, by Corollary 3.5.12 we have $\mathfrak{d}(L/K) = \mathfrak{f}^{\ell-1}$, so $v_{\mathfrak{p}}(\mathfrak{f}) = k_0 + 1$, and the upper bound thus follows from Proposition 10.1.23. The fact that it is the best possible is an immediate consequence of Hecke's Theorem 10.2.9 (1), which we will prove later. □

10.1.5 Application to Dihedral Extensions of Prime Degree

In this section we choose an odd prime number ℓ, and we let L/K be an extension of number fields of degree ℓ. We assume in this section that the Galois closure N of L/K is a *dihedral* extension with Galois group isomorphic to D_ℓ. The goal of this section is to prove Theorem 9.2.6 in this case (the general case can be proved in a similar but more complicated manner). Thus N has a unique quadratic subfield K_2, and the diagram of fields is similar to the one given for Theorem 9.2.6:

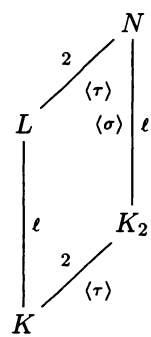

We first prove the following proposition.

Proposition 10.1.25. *Let* $\mathfrak{f}_2 = \mathfrak{f}(N/K_2)$ *be the conductor of* N/K_2. *There exists an ideal* $\mathfrak{f}$ *of* K *such that* $\mathfrak{f}_2 = \mathfrak{f}\mathbb{Z}_{K_2}$.

Proof. Note first that since ℓ is odd, a real place of the normal extension N/K_2 is necessarily unramified, hence the infinite part of the conductor of N/K_2 is empty, so $\mathfrak{f}_2 = \mathfrak{f}(N/K_2)$ is an integral ideal of K_2. Let

$$\mathfrak{f}_2 = \prod_{\mathfrak{P}} \mathfrak{P}^{v_{\mathfrak{P}}(\mathfrak{f}_2)}$$

be the prime ideal decomposition of $\mathfrak{f}_2$ in K_2. Since $\tau(N) = N$, it is clear that $\tau(\mathfrak{f}_2) = \mathfrak{f}_2$, hence if $\mathfrak{P}_1$ and $\mathfrak{P}_2$ are the two prime ideals above a prime ideal of K that splits in K_2/K, we have $v_{\mathfrak{P}_1}(\mathfrak{f}_2) = v_{\mathfrak{P}_2}(\mathfrak{f}_2)$. It follows that

$$\mathfrak{f}_2 = \prod_{\mathfrak{p}} \mathfrak{p}^{x_{\mathfrak{p}}} \prod_{\mathfrak{P}} \mathfrak{P}^{v_{\mathfrak{P}}(\mathfrak{f}_2)} \ ,$$

where the first product is over prime ideals of K which are inert or split in K_2/K and $x_{\mathfrak{p}} = v_{\mathfrak{P}}(\mathfrak{f}_2)$ for any prime ideal $\mathfrak{P}$ of K_2 above $\mathfrak{p}$, and the second product is over prime ideals of K_2 above prime ideals $\mathfrak{p}$ of K which ramify in K_2/K as $\mathfrak{p}\mathbb{Z}_{K_2} = \mathfrak{P}^2$. Thus, to prove the proposition we must show that $v_{\mathfrak{P}}(\mathfrak{f}_2)$ is even, so that

$$\prod_{\mathfrak{P}} \mathfrak{P}^{v_{\mathfrak{P}}(\mathfrak{f}_2)} = \Big(\prod_{\mathfrak{p}} \mathfrak{p}^{x_{\mathfrak{p}}}\Big)\mathbb{Z}_{K_2}$$

with $x_{\mathfrak{p}} = v_{\mathfrak{P}}(\mathfrak{f}_2)/2$.

If $\mathfrak{P} \nmid \mathfrak{f}_2$, this is trivial, so we may assume that $\mathfrak{P} \mid \mathfrak{f}_2$, hence $\mathfrak{P}$ is ramified, and hence totally ramified in N/K_2 since $\ell = [N : K_2]$ is prime. Thus we can write $\mathfrak{P}\mathbb{Z}_N = \mathfrak{P}_N^{\ell}$ for an ideal $\mathfrak{P}_N$ of N, hence $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_N^{2\ell}$, so $\mathfrak{p}$ is totally ramified in N/K. For simplicity set $G_k = G_k(\mathfrak{P}_N/\mathfrak{p})$. Since $\mathfrak{p}$ is totally ramified in N/K, we have $G_0 = \mathrm{Gal}(N/K) \simeq D_\ell$. Thus, as we have already seen, we must have $G_1 \simeq C_\ell$, hence $p = \ell$.

Since the only subgroups of C_ℓ are itself and $\{1\}$, it follows that the sequence of ramification groups is $G_0 \simeq D_\ell$, $G_k \simeq C_\ell$ for $1 \le k \le k_0$, and $G_k = \{1\}$ for $k > k_0$. By Corollary 10.1.18, we must have k_0 odd. This is of course the crucial point in this whole proof. Applying Exercise 3 (which for this situation is trivial), we know that $G_k(\mathfrak{P}_N/\mathfrak{P}) = G_k \cap \mathrm{Gal}(N/K_2) = G_k \cap C_\ell$, hence that $G_k(\mathfrak{P}_N/\mathfrak{P}) \simeq C_\ell$ for $0 \le k \le k_0$ and $G_k(\mathfrak{P}_N/\mathfrak{P}) = \{1\}$ for $k > k_0$ with k_0 odd. By Theorem 10.1.22, we thus have

$$v_{\mathfrak{P}_N}(\mathfrak{D}(N/K_2)) = \sum_{k \ge 0}(|G_k(\mathfrak{P}_N/\mathfrak{P})| - 1) = (k_0 + 1)(\ell - 1) \ .$$

Since the residual degree of $\mathfrak{P}_N$ over $\mathfrak{P}$ is equal to 1, we have

$$v_{\mathfrak{P}}(\mathfrak{d}(N/K_2)) = v_{\mathfrak{P}_N}(\mathfrak{D}(N/K_2)) = (k_0+1)(\ell-1) \ .$$

By Corollary 3.5.12 we have $\mathfrak{d}(N/K_2) = \mathfrak{f}_2^{\ell-1}$, so $v_{\mathfrak{P}}(\mathfrak{f}_2) = k_0+1 \equiv 0 \pmod 2$, finishing the proof of the proposition. □

The following proposition completely describes the splitting and ramification behavior of primes in L/K.

Proposition 10.1.26. *Keep the above hypotheses and notation, and let* $\mathfrak{p}$ *be a prime ideal of* K.

(1) *The prime ideal* $\mathfrak{p}$ *cannot be inert in* N/K.
(2) *If* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_{N,1}\mathfrak{P}_{N,2}$ *with prime ideals* $\mathfrak{P}_{N,i}$ *of* N *of degree* ℓ *over* $\mathfrak{p}$, *then* $\mathfrak{p}$ *is inert in* L/K.
(3) *If* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_{N,1}\mathfrak{P}_{N,2}\dots\mathfrak{P}_{N,\ell}$ *with prime ideals* $\mathfrak{P}_{N,i}$ *of* N *of degree* 2 *over* $\mathfrak{p}$, *then*
$$\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}_{L,1}\mathfrak{P}_{L,2}\dots\mathfrak{P}_{L,(\ell+1)/2} \ ,$$
where $\mathfrak{P}_{L,1}$ *has degree* 1 *over* $\mathfrak{p}$ *and* $\mathfrak{P}_{L,i}$ *has degree* 2 *over* $\mathfrak{p}$ *for* $2 \le i \le (\ell+1)/2$.
(4) *If* $\mathfrak{p}$ *is totally split in* N/K, *it is totally split in* L/K.
(5) *We cannot have* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_N^2$ *with a prime ideal* $\mathfrak{P}_N$ *of* N *of degree* ℓ *over* $\mathfrak{p}$.
(6) *If* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_{N,1}^2\mathfrak{P}_{N,2}^2\dots\mathfrak{P}_{N,\ell}^2$ *with prime ideals* $\mathfrak{P}_{N,i}$ *of* N *of degree* 1 *over* $\mathfrak{p}$, *then*
$$\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}_{L,1}\mathfrak{P}_{L,2}^2\dots\mathfrak{P}_{L,(\ell+1)/2}^2 \ ,$$
where the $\mathfrak{P}_{L,i}$ *have degree* 1 *over* $\mathfrak{p}$.
(7) *If* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_N^\ell$ *with a prime ideal* $\mathfrak{P}_N$ *of* N *of degree* 2 *over* $\mathfrak{p}$, *then* $\mathfrak{p}$ *is totally ramified in* L/K.
(8) *If* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_{N,1}^\ell\mathfrak{P}_{N,2}^\ell$ *with prime ideals* $\mathfrak{P}_{N,i}$ *of* N *of degree* 1 *over* $\mathfrak{p}$, *then* $\mathfrak{p}$ *is totally ramified in* L/K.
(9) *If* $\mathfrak{p}$ *is totally ramified in* N/K, *in other words if* $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_N^{2\ell}$, *then* $\mathfrak{p}$ *is totally ramified in* L/K, *and in addition* $\mathfrak{p} \mid \ell$.

Proof. Note first that if g is the number of prime ideals of N above $\mathfrak{p}$ and if $\mathfrak{P}_N$ is one of them, we have $e(\mathfrak{P}_N/\mathfrak{p})f(\mathfrak{P}_N/\mathfrak{p})g = [N:K] = 2\ell$, hence the possibilities listed in the proposition are exhaustive.

(1). This follows immediately from Corollary 10.1.7 since D_ℓ is not a cyclic group.

(2). By transitivity of residual degrees and since ℓ is an odd prime, if $\ell \mid f(\mathfrak{P}_N/\mathfrak{p})$ for some prime ideal $\mathfrak{P}_N$ of N, we must have $\ell \mid f((\mathfrak{P}_N \cap L)/\mathfrak{p})$; in other words, $\mathfrak{p}$ is inert in L/K, proving (2).

(3). Since the $\mathfrak{P}_{N,i}$ are prime ideals of degree 2 over $\mathfrak{p}$, it follows that $G_{-1}(\mathfrak{P}_{N,i}/\mathfrak{p})$ is a subgroup of order 2 in D_ℓ. Since the Galois group of N/K

permutes transitively the $\mathfrak{P}_{N,i}$, and since the Galois group acts by conjugation on the decomposition groups, it follows that when $1 \leq i \leq \ell$, the decomposition groups $G_{-1}(\mathfrak{P}_{N,i}/\mathfrak{p})$ span the ℓ subgroups of order 2 of D_ℓ. Thus, exactly one of these groups, say $G_{-1}(\mathfrak{P}_{N,1}/\mathfrak{p})$, will be equal to $\mathrm{Gal}(N/L)$, and the others will have a trivial intersection. Since the residual degrees are transitive, it follows that the prime ideal of L below $\mathfrak{P}_{N,1}$ will be of degree 1 over $\mathfrak{p}$, and the prime ideals of L below the $\mathfrak{P}_{N,i}$ for $2 \leq i \leq \ell$ will be of degree 2, proving (3).

(4). Trivial.

(5). If $\mathfrak{p}\mathbb{Z}_N = \mathfrak{P}_N^2$, then $G_{-1}(\mathfrak{P}_N/\mathfrak{p}) \simeq D_\ell$ and $G_0(\mathfrak{P}_N/\mathfrak{p}) \simeq C_2$, which is impossible since no subgroup of D_ℓ isomorphic to C_2 is normal in D_ℓ.

(6). The proof of (6) is identical to that of (3), replacing the decomposition groups G_{-1} by the inertia groups G_0, and the residual degrees by the ramification indices.

(7) and (8). Same proof as for (2), replacing residual degrees by ramification indices.

(9). The first statement of (9) is proved as (7) and (8). The second has been proved during the proof of Proposition 10.1.25. □

Remark. By giving explicit examples for $\ell = 3$, it is easy to show that all possibilities not excluded by this proposition can occur (Exercise 7).

Lemma 10.1.27. *Keep the above hypotheses and notation. We have*

$$\mathcal{N}_{L/K}(\mathfrak{d}(N/L)) = \mathfrak{d}(K_2/K) \ .$$

Proof. We are going to show that for every prime ideal $\mathfrak{p}$ of K, the $\mathfrak{p}$-adic valuations of both sides are equal. Thus, let $\mathfrak{p}$ be a prime ideal of K, and assume first that $\mathfrak{p} \nmid 2$. If $\mathfrak{p} \mid \mathfrak{d}(K_2/K)$, then $\mathfrak{p}\mathbb{Z}_{K_2} = \mathfrak{P}_{K_2}^2$, hence we are necessarily in cases (6) or (9) of the above proposition, since case (5) cannot occur. In case (9), $\mathfrak{p}$ is totally ramified everywhere, and since $\mathfrak{p} \mid \ell$ and $\ell \neq 2$, it is easily seen by applying Proposition 3.3.21 that $v_\mathfrak{p}(\mathfrak{d}(N/L)) = v_\mathfrak{p}(\mathfrak{d}(K_2/K)) = 1$.

In case (6), we have $\mathfrak{p}\mathbb{Z}_L = \mathfrak{P}_{L,1}\mathfrak{P}_{L,2}^2 \cdots \mathfrak{P}_{L,(\ell+1)/2}^2$, $\mathfrak{P}_{L,1}$ is ramified in N/L, and the $\mathfrak{P}_{L,i}$ are split in N/L for $i > 1$. It follows that $\mathfrak{P}_{L,i} \nmid \mathfrak{d}(N/L)$ for $i > 1$. By Proposition 3.3.21 once again, we know that $v_\mathfrak{p}(\mathfrak{d}(K_2/K)) = v_{\mathfrak{P}_{L,1}}(\mathfrak{d}(N/L)) = 1$, and since $\mathfrak{P}_{L,1}$ is of degree 1 above $\mathfrak{p}$ and is the only prime ideal of L above $\mathfrak{p}$ ramified in N/L, we have

$$v_\mathfrak{p}(\mathfrak{d}(K_2/K)) = v_\mathfrak{p}(\mathcal{N}_{L/K}(\mathfrak{d}(N/L))) = 1 \ .$$

Conversely, assume that $\mathfrak{p} \nmid \mathfrak{d}(K_2/K)$. Then for any prime ideal $\mathfrak{P}_N$ of N above $\mathfrak{p}$, $e(\mathfrak{P}_N/\mathfrak{p})$ is odd (it must be equal to 1 or ℓ), hence no prime ideal of L above $\mathfrak{p}$ can ramify in N/L by transitivity of ramification indices. Thus $v_\mathfrak{p}(\mathcal{N}_{L/K}(\mathfrak{d}(N/L))) = v_\mathfrak{p}(\mathfrak{d}(K_2/K)) = 0$. We have thus proved that if $\mathfrak{p} \nmid 2$ (the "tame" case), then $v_\mathfrak{p}(\mathcal{N}_{L/K}(\mathfrak{d}(N/L))) = v_\mathfrak{p}(\mathfrak{d}(K_2/K))$.

Assume now that $\mathfrak{p} \mid 2$ (the "wild" case). The reasoning we have just made is still valid in that case, hence we may assume that $\mathfrak{p} \mid \mathfrak{d}(K_2/K)$. Thus $\mathfrak{p}\mathbb{Z}_{K_2} = \mathfrak{P}_{K_2}^2$, so we are in case (6) (case (9) is impossible here since it would imply $\mathfrak{p} \mid \ell$). We have $G_k(\mathfrak{P}_{N,i}/\mathfrak{p}) \simeq C_2$ for $0 \leq k \leq k_0$ and $G_k(\mathfrak{P}_{N,i}/\mathfrak{p}) = \{1\}$ for $k > k_0$ for some $k_0 \geq 0$. It follows from this and Theorem 10.1.22 that $v_{\mathfrak{P}_{N,i}}(\mathfrak{D}(N/K)) = k_0 + 1$. Intersecting the ramification groups with $\mathrm{Gal}(N/L)$ and $\mathrm{Gal}(N/K_2)$, respectively, and applying Theorem 10.1.22 once again, we find that $v_{\mathfrak{P}_{N,1}}(\mathfrak{D}(N/L)) = k_0+1$, $v_{\mathfrak{P}_{N,i}}(\mathfrak{D}(N/L)) = 0$ for $i > 1$, and $v_{\mathfrak{P}_{N,i}}(\mathfrak{D}(N/K_2)) = 0$ for all i. Since the different is transitive, we have $\mathfrak{D}(N/K) = \mathfrak{D}(N/K_2)\mathfrak{D}(K_2/K)$. Hence for all i,

$$k_0 + 1 = v_{\mathfrak{P}_{N,i}}(\mathfrak{D}(N/K)) = v_{\mathfrak{P}_{N,i}}(\mathfrak{D}(K_2/K)\mathbb{Z}_N) \ .$$

This implies that $v_{\mathfrak{p}}(\mathfrak{d}(K_2/K)) = k_0+1$. On the other hand, $v_{\mathfrak{P}_{L,1}}(\mathfrak{d}(N/L)) = k_0+1$ and $v_{\mathfrak{P}_{L,i}}(\mathfrak{d}(N/L)) = 0$ for $i > 1$, hence $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(\mathfrak{d}(N/L))) = k_0+1 = v_{\mathfrak{p}}(\mathfrak{d}(K_2/K))$, finishing the proof of the lemma. □

We can now easily prove the case $n = \ell$ prime of Martinet's Theorem 9.2.6.

Proposition 10.1.28. *Keep the above hypotheses and notation. In particular, recall that* $\mathfrak{f}$ *is an ideal of* K *such that* $\mathfrak{f}\mathbb{Z}_{K_2}$ *is the conductor of* N/K_2. *Let* $\mathfrak{p}$ *be a prime ideal of* K.

(1) *We have*

$$\mathfrak{d}(L/K) = \mathfrak{d}(K_2/K)^{(\ell-1)/2}\mathfrak{f}^{\ell-1} = \left(\mathfrak{d}(K_2/K)\mathfrak{f}^2\right)^{(\ell-1)/2} \ .$$

(2) *The ideal* $\mathfrak{p}$ *is totally ramified in* L/K *if and only if* $\mathfrak{p} \mid \mathfrak{f}$.
(3) *If* $\mathfrak{p}^2 \mid \mathfrak{f}$, *then* $\mathfrak{p} \mid \ell$.
(4) *If* $\mathfrak{p} \mid (\mathfrak{d}(K_2/K), \mathfrak{f})$, *then* $\mathfrak{p} \mid \ell$.

Proof. (1). By Theorem 2.5.1 we know that

$$\begin{aligned}\mathfrak{d}(N/K) &= \mathfrak{d}(L/K)^2\,\mathcal{N}_{L/K}(\mathfrak{d}(N/L)) = \mathfrak{d}(K_2/K)^{\ell}\,\mathcal{N}_{K_2/K}(\mathfrak{d}(N/K_2)) \\ &= \mathfrak{d}(K_2/K)^{\ell}\,\mathcal{N}_{K_2/K}\left(\mathfrak{f}_2^{\ell-1}\right) = \mathfrak{d}(K_2/K)^{\ell}\mathfrak{f}^{2(\ell-1)} \ ,\end{aligned}$$

where we have used Corollary 3.5.12 and Proposition 10.1.25. By Lemma 10.1.27, we deduce that

$$\mathfrak{d}(L/K)^2 = \mathfrak{d}(K_2/K)^{\ell-1}\mathfrak{f}^{2(\ell-1)} \ ,$$

proving (1).

(2). The prime ideal $\mathfrak{p}$ is totally ramified in L/K if and only if we are in cases (7), (8), or (9) of Proposition 10.1.26, and these are the cases for which some prime ideal of K_2 above $\mathfrak{p}$ is totally ramified in N/K_2, hence for which $\mathfrak{p} \mid \mathfrak{f}$.

(3). This is a restatement of Proposition 3.3.21.

(4). If $\mathfrak{p} \mid (\mathfrak{d}(K_2/K, \mathfrak{f}))$, then $\mathfrak{p}$ is totally ramified in K_2/K and the prime ideal of K_2 above $\mathfrak{p}$ is totally ramified in N/K_2, hence $\mathfrak{p}$ is totally ramified in N/K, so we are in case (9) of Proposition 10.1.26, and hence $\mathfrak{p} \mid \ell$. □

Remarks

(1) Statements (2), (3), and (4) are easy to prove. On the other hand, we have had some trouble proving (1). It is possible to prove (1) in a more natural and shorter way by using the invariance of Artin L-functions under induction. This would have carried us too far afield, however, hence I have preferred to give the above admittedly heavier proof, which has the added advantage of giving complete information on the ramification and splitting behavior of prime ideals in a dihedral extension L/K.

(2) We leave to the reader the proofs of statements (3) and (5) of Theorem 9.2.6 in the case $n = \ell$ prime (Exercise 8).

10.2 Kummer Theory

In this section, we state and prove in detail a number of results we need in an essential way in Chapter 5 in order to compute defining polynomials of number field extensions using Kummer theory.

10.2.1 Basic Lemmas

Lemma 10.2.1 (Dirichlet's Character Independence Theorem). *Let G be a group, let L be a field, and let $\chi_1, \ldots, \chi_m$ be* distinct *characters of G with values in L^*. The characters χ_i are L-linearly independent, in other words a relation $\sum_{1 \le i \le m} a_i \chi_i = 0$ for $a_i \in L$ implies that $a_i = 0$ for all i.*

Proof. Assume that the characters are L-linearly dependent. Choose a dependence relation of *minimal length*, so that, up to reordering of the χ_i,

$$\forall h \in G \quad \sum_{1 \le i \le n} a_i \chi_i(h) = 0 \tag{1}$$

with n minimal. For any $g \in G$, we have for all h, $\sum_{1 \le i \le n} a_i \chi_i(gh) = 0$. Multiplying relation (1) by $\chi_1(g)$ and subtracting, we obtain that for all g and h in G we have

$$\sum_{1 \le i \le n} a_i (\chi_i(g) - \chi_1(g)) \chi_i(h) \ ,$$

and since the first coefficient vanishes, this is a relation of length $n-1$ between the characters. By the minimality of n, this must be the trivial relation, and

again by minimality the a_i are nonzero, hence $\chi_i(g) = \chi_1(g)$ for all $i \leq n$ and all $g \in G$. Since the characters are distinct, this implies $n = 1$, hence $\chi_1 = 0$, which is absurd. □

Corollary 10.2.2. *Let K and L be fields, and let $\sigma_1, \ldots, \sigma_m$ be* distinct *homomorphisms from K to L. Then the σ_i are L-linearly independent.*

Proof. Simply apply the preceding lemma to $G = K^*$ and to χ_i equal to the restriction of σ_i to K^*. □

Lemma 10.2.3 (Noether's Theorem). *Let L/K be a normal extension with Galois group G, and let ϕ be a map from G to L^*. We will say that ϕ satisfies the* cocycle condition *if for all g, h in G we have*

$$\phi(gh) = \phi(g) \cdot g(\phi(h)) \ .$$

Then ϕ satisfies the cocycle condition if and only if there exists $\alpha \in L^*$ such that

$$\forall g \in G, \ \phi(g) = \frac{\alpha}{g(\alpha)} \ .$$

Proof. If $\phi(g) = \alpha/g(\alpha)$, we have

$$\phi(g) \cdot g(\phi(h)) = \frac{\alpha}{g(\alpha)} g\left(\frac{\alpha}{h(\alpha)}\right) = \frac{\alpha}{g(h(\alpha))} = \phi(gh) \ ,$$

so ϕ satisfies the cocycle condition. Conversely, assume that ϕ satisfies the cocycle condition. For $x \in L$, set

$$\sigma(x) = \sum_{h \in G} \phi(h) h(x) \ .$$

Then σ is an additive map from L to L. Applying Lemma 10.2.2 to the distinct homomorphisms $h \in G$, we deduce that σ is not identically zero (recall that $\phi(h) \neq 0$ for all h by assumption). Hence, let $x \in L$ such that $\alpha = \sigma(x) \neq 0$. We have

$$g(\alpha) = g\left(\sum_{h \in G} \phi(h) h(x)\right) = \sum_{h \in G} g(\phi(h)) gh(x) \ ;$$

hence by the cocycle condition

$$g(\alpha) = \phi(g)^{-1} \sum_{h \in G} \phi(gh) gh(x) = \phi(g)^{-1} \sum_{h \in G} \phi(h) h(x) = \phi(g)^{-1} \alpha \ ,$$

proving the lemma. □

Lemma 10.2.4 (Hilbert's Theorem 90). *Let L/K be a* cyclic *extension with Galois group G generated by an element σ. Then $\alpha \in L$ is an element of relative norm equal to* 1 *if and only if there exists $\beta \in L$ such that $\alpha = \beta/\sigma(\beta)$.*

Proof. Clearly $\mathcal{N}_{L/K}(\sigma(\beta)) = \mathcal{N}_{L/K}(\beta)$, hence the relative norm of $\beta/\sigma(\beta)$ is equal to 1. Conversely, assume that $\mathcal{N}_{L/K}(\alpha) = 1$. Let $n = |G|$, and for $0 \leq i < n$, set

$$\phi(\sigma^i) = \prod_{0 \leq k < i} \sigma^k(\alpha) \ .$$

I claim that ϕ satisfies the cocycle condition. Indeed,

$$\phi(\sigma^i)\sigma^i(\phi(\sigma^j)) = \prod_{0 \leq k < i} \sigma^k(\alpha) \prod_{i \leq k < i+j} \sigma^k(\alpha) = \prod_{0 \leq k < i+j} \sigma^k(\alpha) \ .$$

Hence if $i + j < n$, this is equal to $\phi(\sigma^{i+j})$, while if $i + j \geq n$, this is equal to

$$\prod_{0 \leq k < n} \sigma^k(\alpha) \prod_{0 \leq k < i+j-n} \sigma^k(\alpha) = \mathcal{N}_{L/K}(\alpha)\phi(\sigma^{i+j}) = \phi(\sigma^{i+j})$$

once again, since $\mathcal{N}_{L/K}(\alpha) = 1$. Thus ϕ satisfies the cocycle condition. Hence by Lemma 10.2.3, there exists $\beta \in L^*$ such that $\phi(\sigma^i) = \beta/\sigma^i(\beta)$ for all i, and in particular $\alpha = \phi(\sigma) = \beta/\sigma(\beta)$, as desired. Note that by choosing $\gamma = \sigma^{n-1}(\beta)$, we would get $\alpha = \sigma(\gamma)/\gamma$. □

Remarks

(1) The above construction is algorithmic. Indeed, if we retrace our steps, for all $x \in L$, we have $\beta = \sigma(\beta)\alpha$ with

$$\beta = \sum_{0 \leq i < n} \sigma^i(x) \prod_{0 \leq k < i} \sigma^k(\alpha) \ ,$$

which can, of course, be checked directly. The point of the proof is to show that x can be chosen so that $\beta \neq 0$.

(2) There is a simpler *additive* version of Hilbert's Theorem 90, as well as a version for ideals (see Exercise 9).

(3) If $\beta \in L^*$ is such that $\alpha = \beta/\sigma(\beta)$, then by Galois theory all other possible β are of the form $\gamma\beta$ for $\gamma \in K^*$.

(4) Even though Hilbert's Theorem 90 is not true as written for an arbitrary Abelian extension L/K, there exist suitable generalizations to this case (see Exercise 10).

10.2.2 The Basic Theorem of Kummer Theory

Let K be a number field and $\overline{K}$ a fixed algebraic closure of K. We will assume that all algebraic extensions of K are in $\overline{K}$. Let $n \geq 1$ be an integer, and denote by ζ_n a primitive nth root of unity. In this section, we make the fundamental assumption that $\zeta_n \in K$.

Theorem 10.2.5. *Let $n \geq 1$ be an integer, and let K be a number field such that $\zeta_n \in K$. There is a natural bijection between finite subgroups of K^*/K^{*n} and finite Abelian extensions of K whose Galois group is of exponent dividing n. This bijection is obtained as follows. If B is a finite subgroup of K^*/K^{*n}, the corresponding Abelian extension is obtained by adjoining to K all nth roots of lifts of elements of B. If L is a finite Abelian extension of K whose Galois group is of exponent dividing n, then $B = (L^{*n} \cap K^*)/K^{*n}$. In addition, under this correspondence the Galois group $\mathrm{Gal}(L/K)$ is isomorphic to B.*

Proof. Let B be a finite subgroup of K^*/K^{*n}, and let $S = \{s_1, \ldots, s_k\}$ be a set such that the classes of elements of S in K^* generate the group B (for example, the representatives of all the elements of B). Note that conversely, if S is a finite set, the subgroup of K^*/K^{*n} generated by the classes of the elements of S is also finite, since it has at most $n^{|S|}$ elements. We let

$$K_B = K\left(\sqrt[n]{s_1}, \ldots, \sqrt[n]{s_k}\right) .$$

This makes sense since $\zeta_n \in K$ (it could also be made to make sense otherwise). Note for future reference that for all $b \in K^*$ such that $\overline{b} \in B$, the equation $x^n - b = 0$ has a solution in K_B^* (and, in fact, n solutions since $\zeta_n \in K$). Indeed, if $\overline{b} = \prod_i \overline{s_i}^{a_i}$ for some integers a_i, then $x = \prod_i (\sqrt[n]{s_i})^{a_i}$ is a solution.

We are going to prove that the map $B \mapsto K_B$ is the desired bijection. Note first that K_B/K is a finite Abelian extension. Indeed, it is the compositum of the extensions $K(\sqrt[n]{s_i})/K$, and these are Abelian extensions since $\zeta_n \in K$, and so all the roots of the polynomial $X^n - s_i = 0$ belong to $K(\sqrt[n]{s_i})$ for any choice of the root. In fact, all these extensions are cyclic extensions of degree dividing n; hence the Galois group of their compositum is isomorphic to a subgroup of $(\mathbb{Z}/n\mathbb{Z})^k$, hence in particular has an exponent dividing n.

Let G be the Galois group of K_B/K, and denote by $\mu_n = \mu_n(K)$ the subgroup of K^* of nth roots of unity. We define the following pairing $< , >$ from $G \times B$ to μ_n as follows. Let $\sigma \in G$ and $\overline{b} \in B$. As we have seen, there exists $\beta \in K_B$ such that $\beta^n = b$. We will set

$$< \sigma, \overline{b} > = \frac{\sigma(\beta)}{\beta} .$$

First, note that this is indeed an nth root of unity. In fact, $(\sigma(\beta)/\beta)^n = \sigma(b)/b = 1$ since $b \in K^*$. Second, the definition does not depend on the choice of β. Indeed, if β' is such that $\beta'^n = b\gamma^n$ for some $\gamma \in K^*$, then for some j we have $\beta'/\beta = \zeta_n^j \gamma \in K^*$, and so $\sigma(\beta')/\beta' = \sigma(\beta)/\beta$.

Furthermore, we evidently have

$$< \sigma, \overline{bb'} > = < \sigma, \overline{b} >< \sigma, \overline{b'} > \quad \text{and}$$

$$< \sigma\tau, \overline{b} > = \frac{\sigma\tau(\beta)}{\beta} = \frac{\sigma(\tau(\beta))}{\tau(\beta)} \frac{\tau(\beta)}{\beta} = \tau(< \sigma, \overline{b} >) < \tau, \overline{b} > ,$$

and since τ acts trivially on K, hence on μ_n, we have

$$< \sigma\tau, \overline{b} >=< \sigma, \overline{b} >< \tau, \overline{b} > \ .$$

This means that $< \ , \ >$ is a $\mathbb{Z}$-bilinear pairing. In other words, the map $\sigma \mapsto < \sigma, \cdot >$ is a group homomorphism from G to $\mathrm{Hom}(B, \mu_n)$, and the map $\overline{b} \mapsto < \cdot, \overline{b} >$ is a group homomorphism from B to $\mathrm{Hom}(G, \mu_n)$. We are going to compute the kernels of these two homomorphisms.

First, fix $\sigma \in G$, and assume that $< \sigma, \overline{b} >= 1$ for all $b \in B$. Thus, if $\beta^n = b$, then $\sigma(\beta) = \beta$. This implies that for all our generators s_i we have $\sigma(\sqrt[n]{s_i}) = \sqrt[n]{s_i}$, and so $\sigma(x) = x$ for all $x \in K_B$; hence $\sigma = 1$, so the left kernel is trivial.

Second, fix $\overline{b} \in B$, and assume that $< \sigma, \overline{b} >= 1$ for all $\sigma \in G$. If $\beta^n = b$, we thus have $\sigma(\beta) = \beta$ for all $\sigma \in G$, and hence by Galois theory, $\beta \in K^*$. Thus, $b \in K^{*n}$, so $\overline{b} = \overline{1}$ in B, and the kernel is again trivial. Therefore, we obtain what is called a *perfect pairing* between G and B.

Thus, the two maps we deduce from the pairing are injective, and in particular we obtain

$$|G| \leq |\mathrm{Hom}(B, \mu_n)| \quad \text{and} \quad |B| \leq |\mathrm{Hom}(G, \mu_n)| \ .$$

On the other hand, if A is a finite Abelian group of exponent dividing n then $\mathrm{Hom}(A, \mu_n) \simeq A$ noncanonically (see Exercise 11). Hence $|\mathrm{Hom}(G, \mu_n)| = |G|$ and $|\mathrm{Hom}(B, \mu_n)| = |B|$, so both our injective homomorphisms are also surjective, from which we deduce that

$$B \simeq \mathrm{Hom}(G, \mu_n) \simeq G \ .$$

Thus, to each finite subgroup B of K^*/K^{*n} we have associated a finite Abelian extension K_B of K whose Galois group G is of exponent n and isomorphic to B.

Conversely, let L be such an Abelian extension. We must show that $L = K_B$ for a suitable B. Let G be the Galois group of L/K. We are going to show that $B = (L^{*n} \cap K^*)/K^{*n}$ is such that $L = K_B$. Clearly, B is a subgroup of K^* of exponent dividing n. Let us show that B is finite. Using the same pairing $< \ , \ >$ as before, we see that the proof of the injectivity of the map $B \to \mathrm{Hom}(G, \mu_n)$ did not use the finiteness of B. Thus this map is still injective, and since G is a finite group, we deduce that B is finite.

Lemma 10.2.6. *Any homomorphism from G to μ_n is of the form*

$$\sigma \longmapsto < \sigma, \overline{b} >$$

for some $\overline{b} \in B$.

Assuming this lemma, it follows that the map $B \to \mathrm{Hom}(G, \mu_n)$ is a bijection and hence that $|B| = |G|$. By definition of B we have $K \subset K_B \subset L$.

Since $\mathrm{Gal}(K_B/K) \simeq B$, $\mathrm{Gal}(L/K) = G$ and $|B| = |G|$, it follows that $[K_B : K] = [L : K]$ and so $L = K_B$, as claimed.

To prove the lemma, let ϕ be a homomorphism from G to $\boldsymbol{\mu}_n$. Recall that $\boldsymbol{\mu}_n \subset K$, hence that any element of $G = \mathrm{Gal}(L/K)$ fixes $\boldsymbol{\mu}_n$ pointwise. Thus, for all σ and τ in G we have

$$\phi(\sigma\tau) = \phi(\sigma)\phi(\tau) = \phi(\sigma)\sigma(\phi(\tau)) \ .$$

Thus the map ϕ considered as a map from G to L^* satisfies the conditions of Noether's theorem (Lemma 10.2.3); therefore there exists $\alpha \in L^*$ such that $\phi(\sigma) = \sigma(\alpha)/\alpha$ for all $\sigma \in G$. Since we also have $\phi(\sigma)^n = 1$, we obtain $\sigma(\alpha)^n = \alpha^n$ for all $\sigma \in G$. Hence by Galois theory $\alpha^n \in K^*$, and so $\alpha^n \in L^{*n} \cap K^*$. It is clear that $\overline{b} = \overline{\alpha^n}$ is such that $\phi(\sigma) =< \sigma, \overline{b} >$, proving the lemma and hence the theorem. □

Corollary 10.2.7. *Let K be a number field and $n \geq 1$ be an integer such that $\zeta_n \in K$.*

(1) *An extension L/K is a cyclic extension of degree n if and only if there exists $\alpha \in K^*$ such that $\overline{\alpha}$ is exactly of order n in K^*/K^{*n} and such that $L = K(\sqrt[n]{\alpha})$.*
(2) *The cyclic extensions $L_1 = K(\sqrt[n]{\alpha_1})$ and $L_2 = K(\sqrt[n]{\alpha_2})$ are K-isomorphic if and only if there exists an integer j coprime to n and $\gamma \in K^*$, such that $\alpha_2 = \alpha_1^j \gamma^n$.*

Proof. (1) Let L/K be a cyclic extension of degree n. By Theorem 10.2.5, there exists a subgroup B of K^*/K^{*n} such that $L = K_B$ and $B \simeq \mathrm{Gal}(L/K) \simeq \mathbb{Z}/n\mathbb{Z}$. If $\overline{\alpha}$ is a generator of B, it is clear that $K_B = K(\sqrt[n]{\alpha})$. Conversely, if $L = K(\sqrt[n]{\alpha})$ with $\alpha \in K^*$, then L/K is a cyclic extension of degree n if and only if $\overline{\alpha}$ generates a subgroup of order n of K^*/K^{*n}.

(2) Let ϕ be a K-isomorphism from L_1 to L_2, and let B_1 and B_2 be the subgroups of K^*/K^{*n} corresponding to L_1 and L_2, respectively. If $z = \phi(\sqrt[n]{\alpha_1})$, we thus have $z^n \in L_2^{*n} \cap K^*$, hence $\overline{z^n} \in B_2$, and since ϕ is a K-isomorphism, we have $\overline{z^n} = \overline{\phi(\alpha_1)} = \overline{\alpha_1} \in B_2$, and similarly $\overline{\alpha_2} \in B_1$. Since $\overline{\alpha_i}$ is a generator of B_i, it follows that $\alpha_2 = \alpha_1^j \gamma^n$ and $\alpha_1 = \alpha_2^k \delta^n$. Hence $\alpha_1^{kj-1} \in K^{*n}$, and since α_1 is exactly of order n in K^*/K^{*n}, this implies that $kj \equiv 1 \pmod{n}$, hence j is coprime to n, as claimed. □

Definition 10.2.8. *Let K be a number field and $n \geq 1$ be an integer such that $\zeta_n \in K$. Let α_1 and α_2 be elements of K^* of order exactly equal to n in K^*/K^{*n}. We will say that α_1 and α_2 are* n-Kummer-equivalent *(or simply* Kummer-equivalent *if n is understood) if $K(\sqrt[n]{\alpha_1})$ is K-isomorphic to $K(\sqrt[n]{\alpha_2})$, hence by the above corollary, if there exists an integer j coprime to n and $\gamma \in K^*$ such that $\alpha_2 = \alpha_1^j \gamma^n$.*

Since any finite Abelian extension of K can be obtained as a compositum of cyclic extensions of prime power degree, to build finite Abelian extensions it suffices to build cyclic extensions of prime power degree. In turn, these extensions can be built as towers of extensions of prime degree (although this is not a nice way to look at such extensions). Thus we now consider in detail such extensions.

10.2.3 Hecke's Theorem

In view of our applications to the explicit construction of ray class fields, we now want to study in detail ramification and discriminants of Kummer extensions. In this section, we consider the case of cyclic Kummer extensions of prime degree ℓ, which is the only case that can be treated relatively easily.

We use the following notation. Let K be a number field, let ℓ be a prime, and assume that $\zeta_\ell \in K$. Let L be a cyclic extension of K of degree ℓ. By Corollary 10.2.7, there exists $\alpha \in K^*$, $\alpha \notin K^{*\ell}$ such that $L = K(\sqrt[\ell]{\alpha})$. We let $\mathfrak{d}(L/K)$ be the relative discriminant ideal of L/K and $\mathfrak{D}(L/K)$ be the relative different of L/K (see Definition 2.3.16). Finally, $\mathfrak{p}$ denotes a prime ideal of $\mathbb{Z}_K$.

In this section, our goals are to find the decomposition of $\mathfrak{p}$ in the extension L/K and to compute the valuation $v_\mathfrak{p}(\mathfrak{d}(L/K))$, so as to obtain $\mathfrak{d}(L/K)$ (recall that by Corollary 3.5.12 we have $\mathfrak{d}(L/K) = \mathfrak{f}_0(L/K)^{\ell-1}$, where $\mathfrak{f}(L/K)$ is the conductor of L/K). The final result, due to Hecke, is as follows.

Theorem 10.2.9. *Let K be a number field, ℓ a prime number such that $\zeta_\ell \in K$, and $L = K(\sqrt[\ell]{\alpha})$, where $\alpha \in K^* \setminus K^{*\ell}$. If $\mathfrak{p}$ is a prime ideal of $\mathbb{Z}_K$, we set*

$$e(\mathfrak{p}/\ell) = v_\mathfrak{p}(\ell) = (\ell-1)v_\mathfrak{p}(1-\zeta_\ell) \ ,$$

so that $e(\mathfrak{p}/\ell)$ is the absolute ramification index of $\mathfrak{p}$ if $\mathfrak{p}$ is above ℓ and 0 otherwise, and we also set

$$z(\mathfrak{p},\ell) = \ell\frac{e(\mathfrak{p}/\ell)}{\ell-1} + 1 = \ell v_\mathfrak{p}(1-\zeta_\ell) + 1 \ .$$

(1) *Assume that $\ell \nmid v_\mathfrak{p}(\alpha)$. Then $\mathfrak{p}$ is totally ramified in L/K and*

$$v_\mathfrak{p}(\mathfrak{d}(L/K)) = \ell - 1 + \ell e(\mathfrak{p}/\ell) = (\ell-1)z(\mathfrak{p},\ell) \ .$$

In particular, $v_\mathfrak{p}(\mathfrak{d}(L/K)) = \ell - 1$ if $\ell \nmid v_\mathfrak{p}(\alpha)$ and $\mathfrak{p} \nmid \ell$.

(2) *Assume that $\ell \mid v_\mathfrak{p}(\alpha)$ and that $\mathfrak{p} \nmid \ell$. Then $\mathfrak{p}$ is unramified in L/K, and hence $v_\mathfrak{p}(\mathfrak{d}(L/K)) = 0$. In addition, $\mathfrak{p}$ is totally split in L/K if and only if the congruence*

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{1+v_\mathfrak{p}(\alpha)}}$$

has a solution in K; otherwise, $\mathfrak{p}$ is inert in L/K.

(3) *Finally, assume that* $\ell \mid v_{\mathfrak{p}}(\alpha)$ *and that* $\mathfrak{p} \mid \ell$. *Let* a *be the largest (possibly infinite) exponent* k *such that the congruence*

$$x^{\ell} \equiv \alpha \pmod{\mathfrak{p}^{k+v_{\mathfrak{p}}(\alpha)}}$$

has a solution. Then:

a) $\mathfrak{p}$ *is totally split in* L/K *if and only if* $a \geq z(\mathfrak{p},\ell)$, *in which case we have in fact* $a = \infty$ — *in other words, the congruence is soluble for all* k;

b) $\mathfrak{p}$ *is inert in* L/K *if and only if* $a = z(\mathfrak{p},\ell) - 1$;

c) $\mathfrak{p}$ *is totally ramified in* L/K *if and only if* $a \leq z(\mathfrak{p},\ell) - 2$; *in that case, we have* $a \geq 1$, $\ell \nmid a$, *and*

$$v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = (\ell - 1)(z(\mathfrak{p},\ell) - a) \ .$$

Remark. Here and in the sequel, we make an abuse of notation. When we write $x^{\ell} \equiv \alpha \pmod{\mathfrak{p}^k}$, we mean in fact $v_{\mathfrak{p}}(x^{\ell} - \alpha) \geq k$, in other words, $x^{\ell} \equiv \alpha \pmod{{}^*\mathfrak{p}^k}$.

Proof. To simplify notation, write ζ for ζ_{ℓ}. For (1), assume that $\ell \nmid v_{\mathfrak{p}}(\alpha)$. Let π be a uniformizer of $\mathfrak{p}$ in K. There exist integers x and y such that $x\ell + yv_{\mathfrak{p}}(\alpha) = 1$. Thus, if $\beta = \alpha^y \pi^{x\ell}$ we have $v_{\mathfrak{p}}(\beta) = 1$. Furthermore, since $(y,\ell) = 1$, by Corollary 10.2.7 we have $K(\sqrt[\ell]{\beta}) = K(\sqrt[\ell]{\alpha})$. Thus, replacing α by β, we may assume that $v_{\mathfrak{p}}(\alpha) = 1$.

Set $\theta = \sqrt[\ell]{\alpha}$. Thus, θ is a root of the polynomial $X^{\ell} - \alpha = 0$, which is an Eisenstein polynomial since $v_{\mathfrak{p}}(\alpha) = 1$. By Eisenstein's criterion (see [Coh0, Corollary 6.2.4]), or more precisely by its extension to the relative case, it follows that $\mathfrak{p}$ is totally ramified in L/K. This is, however, easily checked directly. Set $\mathfrak{P} = \mathfrak{p}\mathbb{Z}_L + \theta\mathbb{Z}_L$. Then, by looking at valuations, we see that

$$\mathfrak{P}^{\ell} = \mathfrak{p}^{\ell}\mathbb{Z}_L + \theta^{\ell}\mathbb{Z}_L = (\mathfrak{p}^{\ell} + \alpha\mathbb{Z}_K)\mathbb{Z}_L \ .$$

Since $v_{\mathfrak{p}}(\alpha) = 1$, we have $\mathfrak{p}^{\ell} + \alpha\mathbb{Z}_K = \mathfrak{p}$; hence $\mathfrak{P}^{\ell} = \mathfrak{p}\mathbb{Z}_L$, as claimed.

Thus

$$\ell v_{\mathfrak{P}}(\theta) = v_{\mathfrak{P}}(\alpha) = \ell v_{\mathfrak{p}}(\alpha) = \ell \ ,$$

so $\theta \in \mathfrak{P} \setminus \mathfrak{P}^2$ is a uniformizer for $\mathfrak{P}$. Since $\mathfrak{p}$ is totally ramified, Corollary 10.1.21 tells us that $v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(f'(\theta))$, where $f(X)$ is the minimal polynomial of θ in $\mathbb{Z}_K[X]$, which here is simply $f(X) = X^{\ell} - \alpha$. Hence

$$v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = v_{\mathfrak{P}}(\ell\theta^{\ell-1}) = \ell - 1 + \ell v_{\mathfrak{p}}(\ell) = \ell - 1 + \ell e(\mathfrak{p}/\ell) \ ,$$

proving (1).

For (2), we assume that $\ell \mid v_{\mathfrak{p}}(\alpha)$ and $\mathfrak{p} \nmid \ell$. Replacing if necessary α by $(\pi^{-v_{\mathfrak{p}}(\alpha)/\ell})^{\ell}\alpha$, we may assume that $v_{\mathfrak{p}}(\alpha) = 0$. Since a trivial computation shows that the discriminant of the *polynomial* $X^{\ell} - \alpha$ is equal to $(-1)^{(\ell-1)(\ell-2)/2}\ell^{\ell}\alpha^{\ell-1}$, the valuation at $\mathfrak{p}$ of this discriminant is equal to

zero, hence $\mathfrak{p}$ is unramified in L/K. Furthermore, by Proposition 2.3.9, the decomposition type of $\mathfrak{p}$ in $\mathbb{Z}_L$ is the same as that of the polynomial $X^\ell - \alpha$ in $(\mathbb{Z}_K/\mathfrak{p})[X]$, and since $\zeta \in \mathbb{Z}_K$, (2) is clear.

The proof of (3) is the longest, although it is not more difficult than the other proofs. We assume here that $\ell \mid v_\mathfrak{p}(\alpha)$ and $\mathfrak{p} \mid \ell$. As in (2), we may assume that $v_\mathfrak{p}(\alpha) = 0$. To simplify notation, we write e instead of $e(\mathfrak{p}/\ell)$ (any other ramification index will be written explicitly). We first note the following lemma.

Lemma 10.2.10. *Assume that $v_\mathfrak{p}(\alpha) = 0$ and $\mathfrak{p} \mid \ell$. The congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ has a solution in K for $k = z(\mathfrak{p},\ell)$ if and only if it has a solution for* all *$k \geq z(\mathfrak{p},\ell)$.*

Proof. Assume that the congruence has a solution x for some $k \geq z(\mathfrak{p},\ell)$. We apply a Newton–Hensel iteration: in other words, we set

$$x_1 = x + y \quad \text{with} \quad y = -\frac{x^\ell - \alpha}{\ell x^{\ell-1}} \ .$$

Then $v_\mathfrak{p}(y) \geq k - e > 0$. On the other hand,

$$x_1^\ell - \alpha = x^\ell - \alpha + \ell x^{\ell-1} y + \sum_{j \geq 2} \binom{\ell}{j} x^{\ell-j} y^j = \sum_{2 \leq j \leq \ell-1} \binom{\ell}{j} x^{\ell-j} y^j + y^\ell \ .$$

By our assumption on k, for $2 \leq j \leq \ell - 1$ we have

$$v_\mathfrak{p}\left(\binom{\ell}{j} x^{\ell-j} y^j\right) = e + j v_\mathfrak{p}(y) \geq e + 2(k-e) = 2k - e \geq k + 1 \ .$$

Also, $v_\mathfrak{p}(y^\ell) = \ell v_\mathfrak{p}(y) \geq \ell(k-e)$, and the inequality $\ell(k-e) \geq k+1$ is equivalent to $k \geq 1/(\ell-1) + e\ell/(\ell-1)$, hence to $k \geq z(\mathfrak{p},\ell)$, which is true by assumption. Thus all the terms have a $\mathfrak{p}$-adic valuation greater than or equal to $k+1$; hence $x_1^\ell \equiv \alpha \pmod{\mathfrak{p}^{k+1}}$, proving the lemma. □

Resuming the proof of (3), let k be an exponent such that $k \leq z(\mathfrak{p},\ell)$ and such that the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ has a solution. We will use the following construction. Let $k = q\ell + r$ with $0 \leq r < \ell$ be the Euclidean division of k by ℓ, so that $q = \lfloor k/\ell \rfloor$.

If x is a solution of the above congruence, we set $z = \rho(\theta - x)$, where $\theta = \sqrt[\ell]{\alpha}$, $\rho = \pi^{-q}$, and $\pi \in \mathfrak{p} \setminus \mathfrak{p}^2$. Then $(z + \rho x)^\ell - \rho^\ell \alpha = 0$, hence z is a root of the monic polynomial equation $P(Z) = 0$, where

$$P(Z) = Z^\ell + \sum_{1 \leq j \leq \ell-1} \binom{\ell}{j} \rho^j x^j Z^{\ell-j} + \rho^\ell (x^\ell - \alpha) \ .$$

Since $v_{\mathfrak{p}}(\alpha) = 0$, for $1 \leq j \leq \ell - 1$ we have

$$v_{\mathfrak{p}}\left(\binom{\ell}{j}\rho^j x^j\right) = e - jq \geq e - (\ell - 1)q$$

and

$$v_{\mathfrak{p}}(\rho^\ell(x^\ell - \alpha)) \geq -q\ell + k = r \ .$$

Since $k \leq z(\mathfrak{p}, \ell) = 1 + e\ell/(\ell - 1)$, it follows that $q \leq e/(\ell - 1)$; hence the $\mathfrak{p}$-adic valuations of all the coefficients of $P(Z)$ are nonnegative, so $v_{\mathfrak{P}}(z) \geq 0$ for all prime ideals $\mathfrak{P}$ above $\mathfrak{p}$ (see Exercise 12). On the other hand, $\sigma(z) - z = -\rho\theta(1 - \zeta)$, and $v_{\mathfrak{P}}(\theta) = 0$ since $v_{\mathfrak{p}}(\alpha) = 0$, hence

$$v_{\mathfrak{P}}(\sigma(z) - z) = e(\mathfrak{P}/\mathfrak{p})v_{\mathfrak{p}}(\rho(1 - \zeta)) = e(\mathfrak{P}/\mathfrak{p})\left(-q + \frac{e}{\ell - 1}\right) \ .$$

Assume first that our congruence is soluble for some $k \geq z(\mathfrak{p}, \ell) - 1$. Using $k = z(\mathfrak{p}, \ell) - 1$ in the construction above, we obtain $q = e/(\ell - 1)$, and so $v_{\mathfrak{P}}(\sigma(z) - z) = 0$ for all prime ideals $\mathfrak{P}$ above $\mathfrak{p}$. Since the only subgroups of G are G and $\{1\}$, it follows that the inertia group $I = G_0$ is trivial, hence that $\mathfrak{p}$ is unramified in L/K.

Assume in addition that the congruence is soluble for some $k \geq z(\mathfrak{p}, \ell)$. By Lemma 10.2.10, this is equivalent to the solubility for $k = z(\mathfrak{p}, \ell)$. Choosing this value of k in the preceding construction gives an element $z \in L$ such that, for all prime ideals $\mathfrak{P}$ above $\mathfrak{p}$, $v_{\mathfrak{P}}(\sigma(z) - z) = 0$ and $v_{\mathfrak{P}}(z) \geq 0$. Furthermore, up to sign the relative norm of z is equal to the constant term of the polynomial $P(Z)$; in other words,

$$\mathcal{N}_{L/K}(z) = \pm\rho^\ell(x^\ell - \alpha) \ .$$

In particular, $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(z)) \geq 1$.

Fix a prime ideal $\mathfrak{P}$ above $\mathfrak{p}$. Using the same proof as before, we see that for all i and j with $i \not\equiv j \pmod{\ell}$ we have $v_{\mathfrak{P}}(\sigma^i(z) - \sigma^j(z)) = 0$. Therefore, we cannot have $v_{\mathfrak{P}}(\sigma^i(z)) = v_{\mathfrak{P}}(\sigma^j(z)) \geq 1$. Since all the valuations are nonnegative, it follows that there exists at most one i such that $0 \leq i \leq \ell - 1$ with $v_{\mathfrak{P}}(\sigma^i(z)) > 0$. Since $v_{\mathfrak{P}}(\mathcal{N}_{L/K}(z)) = v_{\mathfrak{p}}(\mathcal{N}_{L/K}(z)) \geq 1$, there must indeed be such an i. We have thus defined a map from the set S of prime ideals above $\mathfrak{p}$ to the interval $[0, \ell - 1]$. But conversely, if i is in this interval, we cannot have $v_{\mathfrak{P}}(\sigma^i(z)) = 0$ for all prime ideals $\mathfrak{P}$ above $\mathfrak{p}$, since this would imply that the norm of $\sigma^i(z)$, which is equal to the norm of z, is coprime to $\mathfrak{p}$. Thus our map is surjective, and hence S has at least ℓ elements, and hence exactly ℓ since this is the degree of the extension L/K, so $\mathfrak{p}$ is totally split. The map defined above is thus a bijection, so it follows that the prime ideals above $\mathfrak{p}$ are the ideals

$$\mathfrak{P}_i = \mathfrak{p}\mathbb{Z}_L + \sigma^i(z)\mathbb{Z}_L \ .$$

Conversely, assume that $\mathfrak{p}$ is totally split in L/K. Let $\mathfrak{P}$ be a prime ideal above $\mathfrak{p}$. Then $\mathfrak{P}$ is of relative degree 1. Let $k \geq 1$ be an integer, and consider

the natural homomorphism ϕ from $\mathbb{Z}_K/\mathfrak{p}^k$ to $\mathbb{Z}_L/\mathfrak{P}^k$ obtained by sending the class of x modulo $\mathfrak{p}^k$ to the class of x modulo $\mathfrak{P}^k$. Since $\mathfrak{p}$ is unramified, we have $\mathfrak{P}^k \cap \mathbb{Z}_K = \mathfrak{p}^k$ (see Exercise 14). Therefore, the map ϕ is injective. By multiplicativity of the norm, we have $\left|\mathbb{Z}_L/\mathfrak{P}^k\right| = \mathcal{N}_{L/\mathbb{Q}}(\mathfrak{P})^k$ and $\left|\mathbb{Z}_K/\mathfrak{p}^k\right| = \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{p})^k$. Since $\mathfrak{P}$ is of relative degree 1, we have $\mathcal{N}_{L/\mathbb{Q}}(\mathfrak{P}) = \mathcal{N}_{K/\mathbb{Q}}(\mathfrak{p})$; therefore, $\left|\mathbb{Z}_L/\mathfrak{P}^k\right| = \left|\mathbb{Z}_K/\mathfrak{p}^k\right|$. It follows that our map ϕ is also surjective or, in other words, that any element of $\mathbb{Z}_L$ is congruent to an element of $\mathbb{Z}_K$ modulo $\mathfrak{P}^k$.

In particular, there exists $x \in \mathbb{Z}_K$ such that $x - \theta \in \mathfrak{P}^k$. Taking the relative norm of this, we obtain $x^\ell - \alpha \in \mathfrak{p}^k$, and so the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ is soluble for all k. It follows from above that if the congruence is soluble for $a = z(\mathfrak{p},\ell) - 1 = e(\mathfrak{p}/\ell)\ell/(\ell-1)$ but not for $a+1 = z(\mathfrak{p},\ell)$, then $\mathfrak{p}$ is unramified and not split, so $\mathfrak{p}$ is inert.

We now assume that the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ is not soluble in k for $k = z(\mathfrak{p},\ell) - 1$, and we let a be the largest exponent for which it does have a solution, so that $a \leq z(\mathfrak{p},\ell) - 2$.

Lemma 10.2.11. *With this notation, we have $a \geq 1$ and $\ell \nmid a$.*

Proof. This follows immediately from Proposition 10.2.13, which we will prove in the next section. □

Thus, if we write $a = \ell q + r$ with $0 \leq r < \ell$ as above, we know that $r \geq 1$. If we use our construction, we obtain an element z such that $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(z)) = r \geq 1$. Multiplying by an element prime to $\mathfrak{p}$, we may assume that $z \in \mathbb{Z}_L$. If we set $I = \mathfrak{p}\mathbb{Z}_L + z\mathbb{Z}_L$, so that $\mathfrak{p}\mathbb{Z}_L \subset I \subset \mathbb{Z}_L$, then $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(I)) \geq 1$. Hence $I \neq \mathbb{Z}_L$, and $I \neq \mathfrak{p}\mathbb{Z}_L$ since otherwise $z \in \mathfrak{p}$, and so $v_{\mathfrak{p}}(\mathcal{N}_{L/K}(z)) = v_{\mathfrak{p}}(z^\ell) \geq \ell$, in contradiction with $r < \ell$. So $\mathfrak{p}\mathbb{Z}_L$ is not a maximal ideal, hence $\mathfrak{p}$ is not inert, and since we have seen that $\mathfrak{p}$ cannot be split, it follows that $\mathfrak{p}$ is totally ramified in L/K, say $\mathfrak{p} = \mathfrak{P}^\ell$.

To compute the valuation of the relative discriminant, we use Theorem 10.1.22. Since $\mathfrak{d}(L/K) = \mathcal{N}_{L/K}(\mathfrak{D}(L/K))$ and $\mathcal{N}_{L/K}(\mathfrak{P}) = \mathfrak{p}$ (the ideal $\mathfrak{p}$ being totally ramified in L/K), we obtain the formula

$$v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = \sum_{i \geq 0} (|G_k| - 1) \ .$$

Let us compute the cardinalities of the ramification groups G_k. Recall that if $\theta^\ell = \alpha$ and $\rho = \pi^{-q}$, then $z = \rho(\theta - x)$. We know that $v_{\mathfrak{P}}(\mathcal{N}_{L/K}(z)) = \ell v_{\mathfrak{p}}(\mathcal{N}_{L/K}(z)) = \ell r$. I claim that $v_{\mathfrak{P}}(z) = r$. Indeed, for $1 \leq j \leq \ell - 1$ we have

$$v_{\mathfrak{P}}(\sigma^j(z) - z) = v_{\mathfrak{P}}(\rho\theta(1-\zeta^j)) = \ell\left(-q + \frac{e}{\ell-1}\right) = z(\mathfrak{p},\ell) - 1 - (a-r) > r \ .$$

Thus, if $v_{\mathfrak{P}}(z) > r$, we would have $v_{\mathfrak{P}}(\sigma^j(z)) > r$ for all j, so $v_{\mathfrak{P}}(\mathcal{N}_{L/K}(z)) > \ell r$, which would be a contradiction. Thus, $v_{\mathfrak{P}}(z) \leq r$; hence for all j, $v_{\mathfrak{P}}(\sigma^j(z)) = v_{\mathfrak{P}}(z)$, from which it follows that $v_{\mathfrak{P}}(z) = r$, as claimed.

Since r is coprime to ℓ, let u and v be such that

$$ur + v\ell = 1 \ , \quad 1 \leq u < \ell \ ,$$
$$\psi = z^u \pi^v = \pi^{(1-ua)/\ell}(\theta - x)^u \ .$$

Then $v_{\mathfrak{P}}(\psi) = uv_{\mathfrak{P}}(z) + v\ell = 1$, so ψ is a uniformizer for the ideal $\mathfrak{P}$.

Since ℓ is prime, $G_k = G$ or $G_k = \{1\}$ and Lemma 10.1.11 implies that $G_k = G$ if and only if $v_{\mathfrak{P}}(\sigma(\psi) - \psi) \geq k + 1$. We now compute this quantity. We have

$$v_{\mathfrak{P}}(\sigma(\psi) - \psi) = 1 - ua + v_{\mathfrak{P}}((\theta\zeta - x)^u - (\theta - x)^u) \ .$$

Set $y = \theta - x = \pi^q z$. We then have

$$(\theta\zeta - x)^u - (\theta - x)^u = (y + \theta(\zeta - 1))^u - y^u = \sum_{j=1}^{u} \binom{u}{j} y^{u-j}\theta^j(\zeta - 1)^j \ .$$

Since $v_{\mathfrak{P}}(y) = q\ell + r = a$, $v_{\mathfrak{P}}(\zeta - 1) = e\ell/(\ell - 1) = z(\mathfrak{p}, \ell) > a$, and $v_{\mathfrak{P}}\left(\binom{u}{j}\right) = 0$ since $1 \leq u \leq \ell - 1$, the valuation at $\mathfrak{P}$ of the term of degree j is equal to

$$(u - j)a + j(z(\mathfrak{p}, \ell) - 1) = ua + j(z(\mathfrak{p}, \ell) - 1 - a) \ .$$

Since $a < z(\mathfrak{p}, \ell) - 1$, these valuations are strictly increasing with j, and so the valuation of the sum is equal to the lowest valuation, obtained for $j = 1$. Hence,

$$v_{\mathfrak{P}}(\sigma(\psi) - \psi) = 1 - ua + (ua + z(\mathfrak{p}, \ell) - 1 - a) = z(\mathfrak{p}, \ell) - a \ .$$

If follows that $G_k = G$ if and only if $0 \leq k \leq z(\mathfrak{p}, \ell) - 1 - a$, hence $v_{\mathfrak{P}}(\mathfrak{D}(L/K)) = (\ell - 1)(z(\mathfrak{p}, \ell) - a)$, and since $\mathfrak{p}$ is totally ramified, we have $v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = v_{\mathfrak{P}}(\mathfrak{D}(L/K))$. Thus,

$$v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = (\ell - 1)(z(\mathfrak{p}, \ell) - a) = \ell - 1 + \ell e - (\ell - 1)a \ ,$$

as claimed, thus finishing the proof of Theorem 10.2.9. □

Corollary 10.2.12. *Let K be a number field, ℓ a prime number such that $\zeta_\ell \in K$, and $L = K(\sqrt[\ell]{\alpha})$, where $\alpha \in K^* \setminus K^{*\ell}$. Let $\mathfrak{p}$ be a prime ideal of $\mathbb{Z}_K$. Then we have the following results.*

(1) *The ideal $\mathfrak{p}$ is unramified in L/K if and only if $\ell \mid v_{\mathfrak{p}}(\alpha)$, and in addition, either $\mathfrak{p} \nmid \ell$ or $\mathfrak{p} \mid \ell$ and the congruence*

$$x^\ell \equiv \alpha \pmod{\mathfrak{p}^{z(\mathfrak{p},\ell)-1+v_{\mathfrak{p}}(\alpha)}} \tag{1}$$

has a solution in K.

(2) *The ideal* $\mathfrak{p}$ *is ramified in* L/K *if and only if either* $\ell \nmid v_{\mathfrak{p}}(\alpha)$ *or if* $\ell \mid v_{\mathfrak{p}}(\alpha)$ *and the congruence* (1) *has no solution in* K.
(3) *We always have* $(\ell - 1) \mid v_{\mathfrak{p}}(\mathfrak{d}(L/K))$, *and* $v_{\mathfrak{p}}(\mathfrak{d}(L/K)) = \ell - 1$ *if and only if* $\ell \nmid v_{\mathfrak{p}}(\alpha)$ *and* $\mathfrak{p} \nmid \ell$.

Proof. This is an immediate consequence of Theorem 10.2.9. □

10.2.4 Algorithms for ℓth Powers

In view of Hecke's Theorem 10.2.9, we must be able to check the solubility of congruences of the type

$$x^{\ell} \equiv \alpha \pmod{\mathfrak{p}^{k+v_{\mathfrak{p}}(\alpha)}}$$

when $\mathfrak{p} \mid \ell$ and $\ell \mid v_{\mathfrak{p}}(\alpha)$, for $1 \le k \le z(\mathfrak{p},\ell) - 1 = \ell e(\mathfrak{p}/\ell)/(\ell - 1)$. If π is a uniformizer of $\mathfrak{p}$, we can replace α by $\alpha/\pi^{v_{\mathfrak{p}}(\alpha)}$ and x by $x/\pi^{v_{\mathfrak{p}}(\alpha)/\ell}$, so we may assume without loss of generality that $v_{\mathfrak{p}}(\alpha) = 0$, an assumption we make in the rest of this section.

There are several methods to check the solubility of the congruence. The most straightforward is to use Algorithm 4.2.17 directly, which gives the structure of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$, together with Algorithm 4.2.18, which gives the discrete logarithm of α. The congruence is soluble if and only if the components of the discrete logarithm of α corresponding to the cyclic components of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ that are divisible by ℓ are themselves divisible by ℓ. This method can take some time since the computation of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$ and/or of discrete logarithms in this group may be an expensive operation.

In our specific situation, we can considerably improve on this general method by using the following proposition.

Proposition 10.2.13. *Keep the above notation and hypotheses, and let* π *be a uniformizer of* $\mathfrak{p}$. *Let* k *be an integer such that* $1 \le k \le z(\mathfrak{p},\ell) - 1 = \ell e(\mathfrak{p}/\ell)/(\ell - 1)$, *and let* $x_{k-1} \in \mathbb{Z}_K$ *be such that* $x_{k-1}^{\ell} \equiv \alpha \pmod{\mathfrak{p}^{k-1}}$. *The congruence* $x_k^{\ell} \equiv \alpha \pmod{\mathfrak{p}^k}$ *is soluble if and only if one of the following two conditions holds:*

(1) $v_{\mathfrak{p}}(x_{k-1}^{\ell} - \alpha) \ge k$, *in which case we can take* $x_k = x_{k-1}$;
(2) $v_{\mathfrak{p}}(x_{k-1}^{\ell} - \alpha) = k - 1$ *and* $k \equiv 1 \pmod{\ell}$, *in which case we can take* $x_k = x_{k-1} + \pi^{(k-1)/\ell} y$, *where* y *is a solution of the congruence* $y^{\ell} \equiv (\alpha - x_{k-1}^{\ell})/\pi^{k-1} \pmod{{}^*\mathfrak{p}}$.

Proof. Assume first that the congruence modulo $\mathfrak{p}^k$ is satisfied. Since $\mathbb{Z}_K/\mathfrak{p}$ is a perfect field of characteristic ℓ, the map $x \mapsto x^{\ell}$ is a bijection (hence an injection) from $\mathbb{Z}_K/\mathfrak{p}$ to itself, so if we write $x_k = x_{k-1} + u$, we know that $u \in \mathfrak{p}$. If $u = 0$ — in other words, if $v_{\mathfrak{p}}(x_{k-1}^{\ell} - \alpha) \ge k$ — we are in case (1). Thus, assume that $v_{\mathfrak{p}}(x_{k-1}^{\ell} - \alpha) = k - 1$.

By the binomial theorem we can write

$$x_k^\ell - \alpha = x_{k-1}^\ell - \alpha + s + u^\ell$$

with

$$s = \sum_{1 \le j \le \ell-1} \binom{\ell}{j} x_{k-1}^{\ell-j} u^j \ ,$$

and in particular $v_\mathfrak{p}(s) = e(\mathfrak{p}/\ell) + v_\mathfrak{p}(u)$.

If $v_\mathfrak{p}(u) \ge e(\mathfrak{p}/\ell)/(\ell-1)$, we have $v_\mathfrak{p}(u^\ell) \ge z(\mathfrak{p},\ell) - 1$ and $v_\mathfrak{p}(s) \ge z(\mathfrak{p},\ell) - 1$, hence since $k \le z(\mathfrak{p},\ell) - 1$ we have $v_\mathfrak{p}(x_k^\ell - \alpha) = v_\mathfrak{p}(x_{k-1}^\ell - \alpha) = k-1$, which is a contradiction. Thus if we are not in case (1), we have $1 \le v_\mathfrak{p}(u) < e(\mathfrak{p}/\ell)/(\ell-1)$. In this case we have $v_\mathfrak{p}(s + u^\ell) = \ell v_\mathfrak{p}(u)$; hence a necessary condition for the solubility of the congruence modulo $\mathfrak{p}^k$ is that $\ell v_\mathfrak{p}(u) = v_\mathfrak{p}(x_{k-1}^\ell - \alpha) = k-1$, that is, $k \equiv 1 \pmod{\ell}$, in which case we must choose u such that $v_\mathfrak{p}(u) = (k-1)/\ell$. Since $k \equiv 1 \pmod{\ell}$ and $k \le z(\mathfrak{p},\ell) - 1 = \ell e(\mathfrak{p}/\ell)/(\ell-1)$, we have $k \le \ell e(\mathfrak{p}/\ell)/(\ell-1) - \ell + 1$; hence $e(\mathfrak{p}/\ell) \ge (\ell-1)((k-1)/\ell) + \ell - 1$. Thus

$$v_\mathfrak{p}(s) = e(\mathfrak{p}/\ell) + v_\mathfrak{p}(u) = e(\mathfrak{p}/\ell) + \frac{k-1}{\ell} \ge k + \ell - 2 \ge k \ ,$$

and hence

$$(x_{k-1} + u)^\ell - \alpha \equiv x_{k-1}^\ell - \alpha + u^\ell \pmod{\mathfrak{p}^k} \ .$$

Thus, if we let $y \in \mathbb{Z}_K$ be such that $u/\pi^{(k-1)/\ell} \equiv y \pmod{{}^*\mathfrak{p}}$, we will have $u^\ell \equiv y^\ell \pi^{k-1} \pmod{\mathfrak{p}^k}$, so $y^\ell \equiv (\alpha - x_{k-1}^\ell)/\pi^{k-1} \pmod{{}^*\mathfrak{p}}$, finishing the proof of the necessity of the conditions.

Conversely, if condition (1) is satisfied, we take $x_k = x_{k-1}$. If it is not satisfied we are in case (2), and our construction shows how to construct x_k if we can find a $y \in \mathbb{Z}_K$ such that $y^\ell \equiv (\alpha - x_{k-1}^\ell)/\pi^{k-1} \pmod{{}^*\mathfrak{p}}$. Such a y exists (and is unique modulo $\mathfrak{p}$) since $\mathbb{Z}_K/\mathfrak{p}$ is a perfect field, finishing the proof of the proposition. □

From this proposition, we immediately deduce the following algorithm for solving ℓth power congruences. The proof of its validity, which is immediate from the above proposition, is left to the reader (Exercise 17).

Algorithm 10.2.14 (Solving ℓth Power Congruences). Let K be a number field, let ℓ be a prime number such that $\zeta_\ell \in K$, let $\mathfrak{p}$ be a prime ideal of K above ℓ, and let $\alpha \in \mathbb{Z}_K \setminus \mathbb{Z}_K^\ell$ be an integral element of K such that $v_\mathfrak{p}(\alpha) = 0$. (As explained earlier, it is easy to reduce to this case if we have more the more general condition $\ell \mid v_\mathfrak{p}(\alpha)$.) If k is an integer such that $1 \le k \le z(\mathfrak{p},\ell) - 1 = \ell e(\mathfrak{p}/\ell)/(\ell-1)$, this algorithm either determines an $x \in \mathbb{Z}_K$ such that $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ or says that there is no such x. We assume as usual that the prime ideal $\mathfrak{p}$ is given by a two-element representation $\mathfrak{p} = \ell\mathbb{Z}_K + \pi\mathbb{Z}_K$ with $v_\mathfrak{p}(\pi) = 1$ (otherwise, change π into $\pi + \ell$).

1. [Initialize] Set $v \leftarrow 0$ and $x \leftarrow 0$.
2. [Compute α_2] (Here $v = v_{\mathfrak{p}}(x^\ell - \alpha)$, $v < k$ and $\ell \mid v$.) Set $\alpha_1 \leftarrow (\alpha - x^\ell)/\pi^v$. This will be of the form $\alpha_1 = \beta/d$ with $\beta \in \mathbb{Z}_K$ and $d \in \mathbb{Z}$ coprime to ℓ. Let d_1 be an inverse of d modulo ℓ, and set $\alpha_2 \leftarrow d_1\beta \bmod \ell$.
3. [Compute y] Compute $y \in \mathbb{Z}_K$ such that $y^\ell \equiv \alpha_2 \pmod{\mathfrak{p}}$. Such a y will be unique modulo $\mathfrak{p}$ and can be computed by simple enumeration if $\mathcal{N}(\mathfrak{p})$ is small or by the formula $y \equiv \alpha_2^{\mathcal{N}(\mathfrak{p})/\ell} \pmod{\mathfrak{p}}$ otherwise.
4. [Loop on x] Set $x \leftarrow x + \pi^{v/\ell} y \bmod \ell$ and $v \leftarrow v_{\mathfrak{p}}(x^\ell - \alpha)$. If $v < k$ and $\ell \mid v$, go to step 2.
5. [Terminate] If $v \geq k$, output x; otherwise output a message saying that there is no solution. Terminate the algorithm.

Remarks

(1) By our choice of π, we have $v_{\mathfrak{p}}(\pi) = 1$ and $v_{\mathfrak{q}}(\pi) = 0$ for every prime ideal $\mathfrak{q}$ above ℓ and different from $\mathfrak{p}$. It follows that if $\gamma \in \mathbb{Z}_K$, we have $v_{\mathfrak{p}}(\gamma/\pi^{v_{\mathfrak{p}}(\gamma)}) = 0$ and $v_{\mathfrak{q}}(\gamma/\pi^{v_{\mathfrak{p}}(\gamma)}) \geq 0$ for every prime ideal $\mathfrak{q}$ above ℓ and different from $\mathfrak{p}$. Proposition 4.2.23 shows that $\gamma/\pi^{v_{\mathfrak{p}}(\gamma)}$ is of the form β/d with $\beta \in \mathbb{Z}_K$ and d coprime to ℓ, as claimed in step 2 of the algorithm.

(2) The usual way to compute the $\mathfrak{p}$-adic valuation of an element $\gamma \in \mathbb{Z}_K$ is to use the five-element representation of the prime ideal $\mathfrak{p}$ as explained in [Coh0, Algorithm 4.8.17] and the remark following it (see Algorithms 2.3.13 and 2.3.14 of the present book in the relative case). In the present case, however, we can also compute γ/π^v on some integral basis for $v = 0$, 1, ... as long as the denominator is not divisible by ℓ. For the same reason as in (1), the largest possible value of v will be equal to $v_{\mathfrak{p}}(\gamma)$. The advantage of this method is that the value of α_1 used in step 2 will have already been computed. The disadvantage is that we do successive divisions (by π) instead of successive multiplications, which is usually more expensive.

(3) If we want to compute the integer a that occurs in Theorem 10.2.9 (in other words, the largest $k \leq z(\mathfrak{p}, \ell) - 1$ such that the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ is soluble), we use the above algorithm applied to $k = z(\mathfrak{p}, \ell) - 1$, and we replace step 5 by the following:

5'. [Terminate] Output $a \leftarrow \min(z(\mathfrak{p}, \ell) - 1, v)$ and terminate the algorithm.

Let us consider two special cases of Algorithm 10.2.14.

• $k \leq \ell$. Note that this *always* happens when $e(\mathfrak{p}/\ell) = \ell - 1$, which is a very common occurrence. In this case, there is a single loop: we compute y such that $y^\ell \equiv \alpha \pmod{\mathfrak{p}}$ by one of the two methods indicated in step 3. The congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ is soluble if and only if $v_{\mathfrak{p}}(y^\ell - \alpha) \geq k$, in which case we can, of course, take $x = y$.

• $e(\mathfrak{p}/\ell) < \ell(\ell - 1)$ and $k > e(\mathfrak{p}/\ell)$. In this case, we first check using Algorithm 10.2.14 or another (see, for example, Algorithm 10.2.15) whether

or not the congruence $x_e^\ell \equiv \alpha \pmod{\mathfrak{p}^{e(\mathfrak{p}/\ell)}}$ has a solution. If a solution does exist, then the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ is soluble if and only if $v_\mathfrak{p}(x_e^\ell - \alpha) \geq k$, in which case we can of course take $x = x_e$. Indeed, in that case the largest multiple of ℓ that is strictly less than $z(\mathfrak{p},\ell) - 1$ is equal to $z(\mathfrak{p},\ell) - \ell - 1$, and we have

$$z(\mathfrak{p},\ell) - \ell - 1 = \frac{\ell e(\mathfrak{p}/\ell)}{\ell - 1} - \ell = e(\mathfrak{p}/\ell) + \frac{e(\mathfrak{p}/\ell)}{\ell - 1} - \ell < e(\mathfrak{p}/\ell)$$

since $e(\mathfrak{p}/\ell) < \ell(\ell - 1)$. Thus, according to Proposition 10.2.13, if the congruence has a solution x_e modulo $\mathfrak{p}^{e(\mathfrak{p}/\ell)}$, it will have a solution modulo $\mathfrak{p}^k$ if and only if $v_\mathfrak{p}(x_e^\ell - \alpha) \geq k$, as claimed.

If, in addition, $k \leq e(\mathfrak{p}/\ell)$, we may also use the following algorithm (see [Dab2]).

Algorithm 10.2.15 (Solving ℓth Power Congruences When $k \leq e(\mathfrak{p}/\ell)$). Keep the notation and hypotheses of Algorithm 10.2.14, and assume in addition that $k \leq e(\mathfrak{p}/\ell)$. This algorithm determines an $x \in \mathbb{Z}_K$ such that $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ or says that there is no such x. We assume given an integral basis $(\omega_i)_{1\leq i\leq n}$ of K as well as the prime ideal $\mathfrak{p}$ by a two-element representation $\mathfrak{p} = \ell\mathbb{Z}_K + \pi\mathbb{Z}_K$.

1. [Compute the matrix of the ω_j^ℓ] Compute the $n \times n$ matrix M whose jth column expresses ω_j^ℓ on the ω_i for $1 \leq j \leq n$. The coefficients of M can be reduced modulo ℓ.
2. [Compute the matrix of $\mathfrak{p}^k$] Set $\beta \leftarrow \pi^k$. Compute the $n \times n$ matrix P whose jth column expresses $\beta\omega_j$ on the ω_i for $1 \leq j \leq n$. The coefficients of P can be reduced modulo ℓ. Finally, replace P by the HNF of $(P|\ell I_n)$ (P will be the HNF matrix of $\mathfrak{p}^k$ on the ω_i).
3. [Compute HNF] Apply an HNF algorithm to the matrix $(M|P)$, and let $U = \left(\begin{smallmatrix} U_1 & U_2 \\ U_3 & U_4 \end{smallmatrix}\right)$ be a unimodular matrix and H an HNF matrix such that $(M|P)U = (0|H)$. We can discard the matrices U_2, U_3, and U_4.
4. [Terminate] Let A be the column vector expressing α on the ω_i. If $H^{-1}A \notin \mathbb{Z}^n$, output a message saying that the congruence is not soluble. Otherwise, output the element x whose coefficients on the ω_i are the entries of the vector $U_1 H^{-1} A$. Terminate the algorithm.

Proof. First note that by the binomial theorem, if x and y are in $\mathbb{Z}_K$ we have $(x+y)^\ell \equiv x^\ell + y^\ell \pmod{\ell\mathbb{Z}_K}$. Thus when $k \leq e(\mathfrak{p}/\ell) = v_\mathfrak{p}(\ell\mathbb{Z}_K)$, this is true also modulo $\mathfrak{p}^k$. It follows that if we write $x = \sum_{1\leq j\leq n} x_j\omega_j$ with $x_j \in \mathbb{Z}$, then

$$x^\ell \equiv \sum_{1\leq j\leq n} x_j^\ell \omega_j^\ell \equiv \sum_{1\leq j\leq n} x_j \omega_j^\ell \pmod{\mathfrak{p}^k}$$

by Fermat's little theorem. If $\alpha = \sum_{1\leq i\leq n} a_i\omega_i$, and if $(\beta_j)_{1\leq j\leq n}$ is a $\mathbb{Z}$-basis of $\mathfrak{p}^k$, it follows that the congruence $x^\ell \equiv \alpha \pmod{\mathfrak{p}^k}$ is soluble if and only if there exist x_j and y_j in $\mathbb{Z}$ such that

$$\sum_{1\le j\le n} x_j\omega_j^{\ell} + \sum_{1\le j\le n} y_j\beta_j = \sum_{1\le i\le n} a_i\omega_i \ .$$

This means that the column vector A of the a_i belongs to the image of the matrix $(M|P)$, where M is the matrix of the ω_j^{ℓ} and P the matrix of the β_j.

On the other hand, the HNF matrix of $\mathfrak{p}^k$ is computed as explained in Proposition 2.3.15. Since $k \le e(\mathfrak{p}/\ell)$, all computations can of course be done modulo ℓ. This proves the algorithm's validity. □

Remark. If we are interested only in the *existence* of a solution and not in the solution itself, in step 3 we can simply compute the image of the matrix $(M|P)$ by Gaussian elimination in the field $\mathbb{F}_\ell$, ignoring completely any HNF computation. This will probably be faster.

10.3 Dirichlet Series with Functional Equation

10.3.1 Computing L-Functions Using Rapidly Convergent Series

In this section, we will state and prove a generalization of Theorem 6.1.4, which is essentially due to Lavrik (see [Lav]). I thank E. Friedman for help in writing this section.

In the sequel, using a traditional notation, we will usually write a complex number s as $s = \sigma + iT$, where σ is the real part of s and T the imaginary part. We start with the following definition.

Definition 10.3.1. *Let a_i be a finite sequence of* positive *real numbers, let b_i be a sequence of complex numbers, and let D be a positive real number. If $n \ge 1$, the function $\gamma(s)$ defined by*

$$\gamma(s) = D^{s/2} \prod_{i=1}^{n} \Gamma(a_i s + b_i)$$

will be called a gamma product.

The main properties of gamma products are summarized in the following proposition.

Proposition 10.3.2. *Let $\gamma(s) = D^{s/2} \prod_{i=1}^{n} \Gamma(a_i s + b_i)$ be a gamma product. Then we have the following results.*

(1) *The function $\gamma(s)$ is a meromorphic function of s.*
(2) *There exists $\sigma_0 \in \mathbb{R}$ such that $\gamma(s)$ is holomorphic for* $\mathrm{Re}(s) > \sigma_0$.
(3) *For any real numbers σ_1 and σ_2 such that $\sigma_1 < \sigma_2$, the function $\gamma(s)$ has only a finite number of poles s such that* $\sigma_1 < \mathrm{Re}(s) < \sigma_2$.

(4) *Set* $N = 2(\sum_i a_i)$, $P = D^{1/2}\prod_i a_i^{a_i}$, $B_r = \sum_i \mathrm{Re}(b_i)$, *and* $B_i = \sum_i \mathrm{Im}(b_i)$. *For any fixed real number* σ, *as* $T \to \pm\infty$ *we have*

$$|\gamma(\sigma + iT)| \sim C_\pm P^\sigma\, |T|^{N\sigma/2+B_r-n/2}\, e^{-\pi|T|N/4} \ ,$$

with

$$C_\pm = (2\pi)^{n/2} \prod_i a_i^{\mathrm{Re}(b_i)-1/2} e^{\mp\pi B_i/2} \ .$$

Proof. It is clear from the definition that $\gamma(s)$ is a meromorphic function whose poles are the complex numbers s such that $a_i s + b_i = -k$ for some nonnegative integer k — in other words, the numbers $s = -(b_i+k)/a_i$. Since the a_i are positive real numbers, we have $\mathrm{Re}(s) = -(\mathrm{Re}(b_i)+k)/a_i$, so we can take $\sigma_0 = \max_i(-\mathrm{Re}(b_i)/a_i)$. In addition, for each i there are only a finite number of values of the integer k such that $\sigma_1 < -(\mathrm{Re}(b_i)+k)/a_i < \sigma_2$, so this proves (1), (2), and (3).

For (4), we recall the complex Stirling formula for the gamma function, which implies that for fixed real σ, as $|T| \to \infty$ we have

$$|\Gamma(\sigma + iT)| \sim \sqrt{2\pi}\, |T|^{\sigma-1/2}\, e^{-\pi|T|/2} \ .$$

A short computation gives the result of the proposition. □

The number $N = 2\sum_{1\le i\le n} a_i$ is the most important number associated with the gamma product $\gamma(s)$ and will be called the *degree* of the function $\gamma(s)$.

Example. Let $\gamma(s)$ be the gamma factor associated with the Dedekind zeta function of a number field K of signature (r_1, r_2) and degree $n = r_1 + 2r_2$. Then we know that

$$\gamma(s) = D^{s/2}\Gamma\left(\frac{s}{2}\right)^{r_1+r_2}\Gamma\left(\frac{s+1}{2}\right)^{r_2}$$

with $D = |d(K)|\,\pi^{-n}$ (see, for example, [Coh0, Theorem 4.9.12]). Thus we have $N = r_1 + 2r_2 = n$, $P = |d(K)|^{1/2}/(2\pi)^{n/2}$, $B_r = r_2/2$, and $B_i = 0$; hence, in particular, the degree of $\gamma(s)$ is equal to the degree of the number field K, whence the name.

The following lemma gives the essential properties that we need about gamma products.

Lemma 10.3.3. *Let* $\gamma(s)$ *be a gamma product. Denote by* $W(t)$ *the inverse Mellin transform of* $\gamma(s)$, *in other words,*

$$W(t) = \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} t^{-z}\gamma(z)\,dz \ ,$$

and let σ_0 *be the real part of the rightmost pole of* γ.

(1) *The integral defining* $W(t)$ *converges absolutely and is independent of* $\delta > \sigma_0$.

(2) *When* t *tends to* ∞, *the function* $W(t)$ *tends to* 0 *faster than any power of* t. *More precisely, using the notation of Proposition 10.3.2, as* t *tends to* ∞ *we have*

$$W(t) \leq At^{(2B_r-n+1)/N} e^{-(\pi N/4)(t/P)^{2/N}}$$

for some explicit constant $A > 0$.

(3) *For* $\mathrm{Re}(s) > \sigma_0$ *we have*

$$\gamma(s) = \int_0^\infty t^s W(t) \frac{dt}{t} \ .$$

Proof. (1). Since $\gamma(z)$ decreases exponentially when $|\mathrm{Im}(z)|$ tends to infinity with $\mathrm{Re}(z)$ fixed, it is clear that the integral defining $W(t)$ converges absolutely. For the same reason, by integrating over the rectangle $[\delta_1, \delta_2] \times [-T, T]$ and letting T tend to ∞, it is clear that the integral is independent of $\delta > \sigma_0$, proving (1).

(2). For any $\delta > \sigma_0$, we clearly have

$$|W(t)| \leq \frac{1}{2\pi} t^{-\delta} \int_{-\infty}^{\infty} |\gamma(\delta + iT)| \, dT \ ,$$

and this last integral is convergent, so $W(t)$ tends to 0 faster than any power of t since δ can be chosen arbitrarily large. More precisely, if for simplicity we set $C_2 = B_r - n/2$ and if C_1 is suitably large, Proposition 10.3.2 implies that

$$|\gamma(\delta + iT)| \leq C_1 P^\delta \, |T|^{N\delta/2 + C_2} \, e^{-(\pi/4)N|T|} \ ,$$

from which it follows as above that

$$W(t) \leq \frac{C_1 (P/t)^\delta}{\pi} \frac{\Gamma(N\delta/2 + C_2 + 1)}{(\pi N/4)^{N\delta/2 + C_2 + 1}} \ .$$

We choose δ close to the smallest possible value of the right-hand side, for example, $\delta = (\pi/2)(t/P)^{2/N}$, which will indeed be larger than σ_0 for t sufficiently large. A small computation gives (2) with an explicit constant A (see Exercise 18).

(3). This is simply Mellin's inversion formula, which is applicable here since $\gamma(s)$ and $W(t)$ are rapidly decreasing functions as $|\mathrm{Im}(s)| \to \infty$ and $t \to \infty$, respectively. □

Remark. If we use the method of steepest descent instead of simple inequalities, it is not difficult to obtain an asymptotic formula for $W(t)$ instead of an upper bound (see [Bra]), but we will not need this. See also [Tol] for precise inequalities.

The main result that we want to prove is the following. (Recall that a meromorphic function f having a finite number of poles is of *finite type* $\alpha \geq 0$ if for all $\varepsilon > 0$ and all sufficiently large $|z|$ we have $|f(z)| = O\big(\exp\big(|z|^{\alpha+\varepsilon}\big)\big)$.)

Theorem 10.3.4. *For $i = 1$ and $i = 2$, let $L_i(s) = \sum_{n\geq 1} a_i(n)n^{-s}$ be Dirichlet series such that the $a_i(n)$ have at most polynomial growth (or, equivalently, the series $L_i(s)$ converge in some right half-plane $\mathrm{Re}(s) \geq \sigma_0$). For $i = 1$ and $i = 2$, let $\gamma_i(s)$ be functions such that the following conditions hold.*

(1) *The functions $\gamma_i(s)$ are gamma products with the same degree N.*
(2) *The functions $\Lambda_i(s) = \gamma_i(s)L_i(s)$ extend analytically to the whole complex plane into meromorphic functions of finite type having a finite number of poles.*
(3) *There exists a functional equation*

$$\Lambda_1(k-s) = w \cdot \Lambda_2(s)$$

for some constant $w \in \mathbb{C}^*$, some real number k, valid for all s different from the poles of $\Lambda_2(s)$.

Define the functions $F_i(s,x)$ by

$$F_i(s,x) = \frac{x^s}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} \frac{x^{-z}\gamma_i(z)}{z-s}\,dz = \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} \frac{x^{-z}\gamma_i(z+s)}{z}\,dz \ ,$$

where δ is sufficiently large, and the functions $p_i(s,x)$ by

$$p_i(s,x) = \sum_{a\neq s} \mathrm{Res}_{z=a}\left(\frac{x^{z-s}\Lambda_i(z)}{s-z}\right) \ ,$$

where the sum is over all poles a of $\Lambda_i(z)$ different from s.

Then for all $t_0 > 0$, we have

$$\Lambda_1(s) = \sum_{n\geq 1} \frac{a_1(n)}{n^s} F_1(s, nt_0) + w\sum_{n\geq 1} \frac{a_2(n)}{n^{k-s}} F_2\Big(k-s, \frac{n}{t_0}\Big) + p_1\Big(s, \frac{1}{t_0}\Big)$$

and symmetrically

$$\Lambda_2(s) = \sum_{n\geq 1} \frac{a_2(n)}{n^s} F_2\Big(s, \frac{n}{t_0}\Big) + w^{-1}\sum_{n\geq 1} \frac{a_1(n)}{n^{k-s}} F_1(k-s, nt_0) + p_2(s,t_0) \ .$$

If, in addition, $t_0 = 1$, we have $p_i(s,1) = \phi_i(s)$, where $\phi_i(s)$ is the polar part of $\Lambda_i(s)$: in other words, the unique rational function such that $\Lambda_i(s) - \phi_i(s)$ is an entire function and $\phi_i(s)$ tends to 0 as $|s|$ tends to ∞.

Proof. Before beginning the proof itself, we will make a few remarks about convergence. First, it is clear that $L_i(s)$ converges in some right half-plane if

and only if $a_i(n)$ has at most polynomial growth. Since the functions $\Lambda_i(s)$ have only a finite number of poles, we may choose a real number σ_0 large enough so that for $i = 1$ and $i = 2$, the functions $L_i(s)$ converge absolutely for $\mathrm{Re}(s) > \sigma_0$ and all the poles of $\Lambda_i(s)$ are in the strip $k - \sigma_0 < \mathrm{Re}(s) < \sigma_0$.

For $i = 1$ and $i = 2$, denote by $W_i(t)$ the inverse Mellin transform of $\gamma_i(s)$ and by σ_i the real part of the rightmost pole of $\gamma_i(s)$. The definition and main properties of this function are given in Lemma 10.3.3. In addition, for $\mathrm{Re}(s) > \sigma_i$ we have

$$\begin{aligned}\int_x^\infty W_i(t)t^s\,\frac{dt}{t} &= \frac{1}{2i\pi}\int_x^\infty \int_{\delta-i\infty}^{\delta+i\infty} t^{-z}\gamma_i(z)\,dz\; t^s\,\frac{dt}{t}\\ &= \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty}\gamma_i(z)\int_x^\infty t^{s-z-1}\,dtdz\\ &= \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty}\gamma_i(z)\frac{x^{s-z}}{z-s}\,dz = F_i(s,x)\ ,\end{aligned}$$

where the exchange of integral signs is justified by the fact that the functions $\gamma_i(z)$ are rapidly decreasing as $|z| \to \infty$. Thus, for $\mathrm{Re}(s) > \sigma_i$ we have

$$F_i(s,x) = \int_x^\infty W_i(t)t^s\,\frac{dt}{t}\ .$$

Set $\sigma = \max(\sigma_0, \sigma_1, \sigma_2)$, and choose s such that $\mathrm{Re}(s) > \sigma$. We have

$$\begin{aligned}\Lambda_i(s) &= \sum_{n\ge1}\int_0^\infty \frac{a_i(n)}{n^s}W_i(t)t^s\,\frac{dt}{t} = \sum_{n\ge1}\int_0^\infty a_i(n)W_i(nu)u^s\,\frac{du}{u}\\ &= \int_0^\infty \sum_{n\ge1} a_i(n)W_i(nu)u^s\,\frac{du}{u} = \int_0^\infty \theta_i(u)u^s\,\frac{du}{u}\end{aligned}$$

with

$$\theta_i(u) = \sum_{n\ge1} a_i(n)W_i(nu)\ .$$

In the above derivation, we can justify the exchange of summation and integration as follows. We break up the integral into an integral from 0 to 1 plus an integral from 1 to ∞. The exchange in the integral from 0 to 1 is justified since the interval is compact and the series converges uniformly. The exchange in the integral from 1 to ∞ is justified by the fact that by Lemma 10.3.3, $|W_i(u)|$ decreases more rapidly than any power of u as $u \to \infty$.

It follows from the Mellin inversion formula that for $\delta > \sigma$,

$$\theta_i(u) = \frac{1}{2i\pi}\sum_{n\ge1} a_i(n)\int_{\delta-i\infty}^{\delta+i\infty}(nu)^{-s}\gamma_i(s)\,ds = \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty}\Lambda_i(s)u^{-s}\,ds\ ,$$

where the interchange of summation and integration follows immediately from the fact that $\gamma_i(s)$ is a gamma product.

We now want to shift the line of integration to $\mathrm{Re}(s) = k - \delta$. To justify this, we need the following lemma

Lemma 10.3.5. *Keep all the above notation.*

(1) *For all fixed $r \in \mathbb{R}$, there exists $e(r)$ such that as $|T| \to \infty$, we have*

$$L_i(r + iT) = O\left(|T|^{e(r)}\right) .$$

(2) *For all $r_1 < r_2$, we have*

$$\lim_{|T|\to\infty} \int_{r_1+iT}^{r_2+iT} \Lambda_i(s) u^{-s}\, ds = 0 .$$

Proof. We prove (1) for L_1, the result following by symmetry. For $r > \sigma$, we have $L_1(r+iT) = O(1)$ since the series converges absolutely. For $r < k-\sigma$, we apply the functional equation, which gives us

$$L_1(r + iT) = wL_2(k - r - iT)\frac{\gamma_2(k - r - iT)}{\gamma_1(r + iT)} .$$

Since γ_1 and γ_2 are gamma products with the same degree N and $L_2(k-r-iT) = O(1)$ for $r < k - \sigma$, it follows that for $r < k - \sigma$ we have

$$L_1(r + iT) = O\left(|T|^{c(r)}\right)$$

for some constant $c(r)$ depending only on r. Since $\Lambda_1(s)$ is a function of finite type and $\gamma_1(s)$ is a gamma product, the function $L_1(s)$ is of finite type in any strip $\sigma_1 \le \mathrm{Re}(s) \le \sigma_2$. Thus we may apply the Phragmén–Lindelöf convexity theorem on the strip $[k - \sigma, \sigma]$ (see, for example, [Lan3]), which implies that $L_1(r + iT) = O\big(|T|^{e(r)}\big)$ for any r in the strip $k - \sigma < r < \sigma$, with an explicit value of the exponent $e(r)$ (which we do not need), proving (1). The result of (2) is an immediate consequence of (1) and of the fact that $\gamma_i(s)$ is a gamma product. □

Resuming the proof of the theorem, thanks to this lemma we may shift the line of integration to $\mathrm{Re}(s) = k - \delta$, and we simply catch all the residues at the poles of $\Lambda_i(s)$. Changing s into $k - s$ gives

$$\theta_i(u) = \sum_a \mathrm{Res}_{z=a}(\Lambda_i(z)u^{-z}) + \frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty} \Lambda_i(k - s)u^{s-k}\, ds ,$$

where the sum is over all poles a of $\Lambda_i(s)$. Applying the functional equation, we obtain

$$\theta_1(u) = \sum_a \mathrm{Res}_{z=a}(\Lambda_i(z)u^{-z}) + wu^{-k}\frac{1}{2i\pi}\int_{\delta-i\infty}^{\delta+i\infty}\Lambda_2(s)\Big(\frac{1}{u}\Big)^{-s}\,ds$$
$$= \sum_a \mathrm{Res}_{z=a}(\Lambda_i(z)u^{-z}) + wu^{-k}\theta_2\Big(\frac{1}{u}\Big)\ .$$

Changing u into $1/u$ gives the functional equation

$$\theta_1\Big(\frac{1}{u}\Big) = wu^k\theta_2(u) + \sum_a \mathrm{Res}_{z=a}(\Lambda_i(z)u^z)\ .$$

Conversely, by computing the Mellin transform it is easily seen that this functional equation is equivalent to the functional equation linking $\Lambda_1(s)$ with $\Lambda_2(s)$.

Coming back to our integral representation of $\Lambda_1(s)$, for any $t_0 > 0$, we can write $\Lambda_1(s) = I_1 + I_2$ with

$$I_1 = \int_0^{t_0}\theta_1(u)u^s\,\frac{du}{u} \quad \text{and} \quad I_2 = \int_{t_0}^{\infty}\theta_1(u)u^s\,\frac{du}{u}\ .$$

Consider first the integral I_2. We have

$$I_2 = \sum_{n\geq 1} a_1(n)\int_{t_0}^{\infty} W_1(nu)u^s\,\frac{du}{u} = \sum_{n\geq 1}\frac{a_1(n)}{n^s}\int_{nt_0}^{\infty} W_1(t)t^s\,\frac{dt}{t}$$
$$= \sum_{n\geq 1}\frac{a_1(n)}{n^s}F_1(s, nt_0)\ .$$

Consider now the integral I_1. Changing u into $1/u$ and using the functional equation for $\theta_i(u)$ that we just derived, we obtain

$$I_1 = \int_{1/t_0}^{\infty}\theta_1\Big(\frac{1}{u}\Big)u^{-s}\,\frac{du}{u} = wI_1' + \sum_a I_3(a)$$

with

$$I_1' = \int_{1/t_0}^{\infty}\theta_2(u)u^{k-s}\,\frac{du}{u}$$

and

$$I_3(a) = \int_{1/t_0}^{\infty}\mathrm{Res}_{z=a}(u^z\Lambda_1(z))u^{-1-s}\,du\ .$$

The integral I_1' is of the same form as I_2 with θ_1 replaced by θ_2, s by $k-s$, and t_0 by $1/t_0$; hence we obtain

$$I_1' = \sum_{n\geq 1}\frac{a_2(n)}{n^{k-s}}F_2\Big(k-s, \frac{n}{t_0}\Big)\ .$$

But since $\mathrm{Re}(s) > \sigma$, we have

$$I_3(a) = \mathrm{Res}_{z=a}\left(\Lambda_1(z)\int_{1/t_0}^{\infty} u^{z-1-s}\,du\right) = \mathrm{Res}_{z=a}\left(\frac{\Lambda_1(z)t_0^{s-z}}{s-z}\right) .$$

Putting everything together gives the desired formula for $\Lambda_1(s)$ for $\mathrm{Re}(s) > \sigma$.

Since $W_i(t)$ tends to 0 faster than any power of t as $t \to \infty$, it is easily seen that because s is fixed, $F_i(s,x)$ tends to 0 faster than any power of x as $x \to \infty$ (see Exercise 19). It follows that the right-hand side of the formula that we just proved for $\Lambda_1(s)$ defines an analytic continuation of $\Lambda_1(s)$ to the whole complex plane outside the poles of $p_1(s, 1/t_0)$. By uniqueness of analytic continuation, it follows that the identity is valid for every s that is not a pole of $\Lambda_1(s)$, and the formula for $\Lambda_2(s)$ follows by symmetry.

The last statement of the theorem follows from the next lemma.

Lemma 10.3.6. *For any rational function $\phi(z)$ that tends to 0 when $|z|$ tends to ∞, we have the identity*

$$\sum_{a \neq s} \mathrm{Res}_{z=a}\left(\frac{\phi(z)}{s-z}\right) = \phi(s) ,$$

where the sum is over all the poles a of $\phi(z)$ different from s.

Proof. Let R be a real positive number larger than the modulus of the poles of $\phi(z)$ and of $|s|$, and let C_R be the circle of radius R centered at the origin. By the residue theorem we have

$$\sum_{a \neq s} \mathrm{Res}_{z=a}\left(\frac{\phi(z)}{s-z}\right) - \phi(s) = \frac{1}{2i\pi}\int_{C_R} \frac{\phi(z)}{s-z}\,dz .$$

Since $\phi(z)$ is a rational function that tends to 0 when $|z|$ tends to infinity, there exists B such that $|\phi(z)| \leq B/|z|$ for $|z|$ sufficiently large. It follows that

$$\left|\int_{C_R} \frac{\phi(z)}{s-z}\,dz\right| \leq 2\pi R \frac{B}{R}\frac{1}{R-|s|} = \frac{2\pi B}{R-|s|} ,$$

and this tends to zero when R tends to infinity, proving the lemma and hence the theorem. □

The main point of the above theorem is that it allows us to compute the value of $\Lambda_i(s)$ in the whole complex plane, which is not possible using the Dirichlet series directly, and also to compute it using a rapidly convergent series, since the functions $F_i(s,x)$ decrease faster than any power of x to 0 as $x \to \infty$, as we have seen above. In fact, we know from Lemma 10.3.3 (2) and Exercise 19 that they even decrease exponentially fast.

We have kept the arbitrary positive parameter t_0 so that we can *check* the validity of an implementation, since the result must be independent of t_0. In practice, to ensure fastest convergence, if $\mathrm{Im}(s)$ is small, one should choose t_0 close to 1, or equal to 1 if the implementation has been sufficiently checked. On the other hand, if $\mathrm{Im}(s)$ is large, one should choose a *complex* value for t_0 (see Exercise 20). Another use for having a variable t_0 is to obtain *approximate functional equations*; see [Lav] for details.

Remark. The formula given above for $\Lambda_i(s)$ is one among an infinite family of formulas. Indeed, if $g(s)$ is an entire function of finite type and if we set $\gamma_1'(s) = g(s)\gamma_1(s)$, $\gamma_2'(s) = g(k-s)\gamma_2(s)$, and $\Lambda_i'(s) = \gamma_i'(s)L_i(s)$, then we still have a functional equation $\Lambda_1'(k-s) = w\Lambda_2'(s)$. Although the functions $\gamma_i'(s)$ are strictly speaking not gamma products in general, if $g(s)$ satisfies very mild conditions, it is easy to see that the proof of Theorem 10.3.4 goes through without change for this type of function (see Exercise 21). The possibility of using an arbitrary complex value of t_0 is a special case of this construction. The advantage of having an auxiliary function $g(s)$ is that we can tailor it so that the computation of $L_i(s)$ is as fast and accurate as possible. In practice, the method sketched in Exercise 20 using complex values of t_0 is usually sufficient.

10.3.2 Computation of $F_i(s,x)$

The computation of $p_i(s,x)$ being trivial, to be able to compute $\Lambda_i(s)$ we must give an algorithm for computing the functions $F_i(s,x)$. From now on, we drop the index i, and write $W(t)$ instead of $W_i(t)$, $\gamma(s)$ instead of $\gamma_i(s)$, and $F(s,x)$ instead of $F_i(s,x)$.

As can be seen for the formulas, we will need $F(s,x)$ for complex s and positive real x. We can write

$$F(s,x) = \gamma(s) - \int_0^x W(t)t^s \frac{dt}{t} .$$

We are going to expand first $W(t)$ and then $F(s,x)$ into a (generalized) power series. By definition,

$$W(t) = \frac{1}{2i\pi} \int_{\delta-i\infty}^{\delta+i\infty} t^{-z}\gamma(z)\,dz .$$

Here, δ is in particular larger than the real part of all the poles of the function $\gamma(z)$. If we push the line of integration to the left and stop at an abscissa δ' that is not the real part of a pole of $\gamma(z)$, we will have "caught" the residues of the function $t^{-z}\gamma(z)$ with real part between δ' and δ. Since $\gamma(z)$ is a gamma product, the integral of $t^{-z}\gamma(z)$ on the line $\mathrm{Re}(z) = \delta'$ tends to zero when δ' tends to $-\infty$, assuming we choose δ' such that the distance of δ' to the real parts of the poles of $\gamma(z)$ is always larger than η for some fixed $\eta > 0$. By definition of a gamma product, it is easy to see that this is indeed possible.

It follows that

$$W(t) = \sum_a \mathrm{Res}_{z=a}\left(t^{-z}\gamma(z)\right) \ ,$$

where the sum is over all the poles of the function $\gamma(z)$. If

$$\gamma(z) = \sum_{1\le k\le m} \frac{u_{k,a}}{(z-a)^k} + \text{holomorphic at } a$$

is the polar expansion of $\gamma(z)$ in the neighborhood of $z = a$, we have

$$\mathrm{Res}_{z=a}\left(t^{-z}\gamma(z)\right) = t^{-a}\sum_{k=0}^{m-1}\frac{(-1)^k}{k!}u_{k+1,a}\log^k t \ .$$

A little computation shows that if s is not a pole of $\gamma(s)$, we have

$$F(s,x) = \gamma(s) - \sum_a x^{s-a}\sum_{j=0}^{m-1}(-1)^j\frac{\log^j x}{j!}\sum_{k=j+1}^{m}\frac{u_{k,a}}{(s-a)^{k-j}} \ .$$

This is the desired generalized power series expansion of $F(s,x)$.

When s is a pole a_0 of $\gamma(s)$, a similar computation shows that

$$\begin{aligned} F(s,x) = &\sum_{j=0}^{m} u_{j,a_0}(-1)^j\frac{\log^j x}{j!} \\ &- \sum_{a\neq a_0} x^{s-a}\sum_{j=0}^{m-1}(-1)^j\frac{\log^j x}{j!}\sum_{k=j+1}^{m}\frac{u_{k,a}}{(s-a)^{k-j}} \ , \end{aligned}$$

where u_{0,a_0} is the next term (in other words, the constant term) in the Laurent series expansion of $\gamma(z)$ around $z = a_0$ (see Exercise 22).

It is easy to show that this power series converges for all x (see Exercise 23) but suffers from bad cancellation problems when x is large. More precisely, although some of the terms in the above series for $F(s,x)$ are quite large, we know from Exercise 19 that the final result will be exponentially small when x is large, hence we should avoid using the above generalized power series when $N(x/P)^{2/N} > 20$, say.

Consider, for example, one of the simplest cases, $\gamma(z) = \Gamma(z)$. The function $F(s,x)$ is the incomplete gamma function, and the above expansion reads in this case

$$F(s,x) = \Gamma(s) - \sum_{n=0}^{\infty}\frac{(-1)^n x^{s+n}}{n!(s+n)} \ .$$

When $s = 1$ for instance, this is exactly the expansion of e^{-x}, which should evidently *not* be used to compute e^{-x} if x is large.

In this special case, for x large we can use a *continued fraction expansion* for the incomplete gamma function, of which [Coh0, Propositions 5.3.15 and

5.6.12] are special cases. In the general case, however, one would need to generalize this continued fraction by considering higher-order linear recursions. This can indeed be done (see [Del]).

Meanwhile, there are two possible methods. One is to use the generalized power series expansion above for all x. By closely analyzing its convergence, we can quite easily see in any given situation how much accuracy is needed, and E. Tollis (see [Tol]) has shown that it is sufficient to compute with approximately *twice* the required accuracy. For example, to obtain the value of $F(s,x)$ to 28 decimal digits, in reasonable ranges of s it is enough to perform the computations using 56 decimal digits. This result generalizes the well-known fact that for $x > 0$, the largest summand in the expansion of e^{-x} is of the order of e^x (more precisely, $e^x/\sqrt{2\pi x}$).

A complementary method is to choose a value of t_0 that depends on the value of s at which we want to evaluate $\Lambda_i(s)$ (see Exercise 20). A suitable choice can avoid many of the cancellation problems.

10.4 Exercises for Chapter 10

1. With the notation of Section 10.1.2, let $\mathfrak{P}$ be a prime ideal of L above $\mathfrak{p}$, and set $\mathfrak{P}_I = \mathfrak{P} \cap L^I$ and $\mathfrak{P}_D = \mathfrak{P} \cap L^D = \mathfrak{P}_I \cap L^D$. Prove the following results.
   a) We have $\mathfrak{P} = \mathfrak{P}_I^e$; in other words, $\mathfrak{P}_I$ is totally ramified of degree e in the extension L/L^I.
   b) We have $\mathfrak{P}_I = \mathfrak{P}_D \mathbb{Z}_{LI}$; in other words, $\mathfrak{P}_D$ is inert of degree f in the extension L^I/L^D.
   c) We have $f(\mathfrak{P}_D/\mathfrak{p}) = 1$; in other words, $\mathfrak{P}_D$ is of degree 1 over $\mathfrak{p}$.
   d) If $\mathrm{Gal}(L/K)$ is Abelian, then $f(\mathfrak{P}'_D/\mathfrak{p}) = 1$ for all prime ideals $\mathfrak{P}'_D$ of L^D above $\mathfrak{p}$.
   e) Give an explicit example where there exists another prime ideal $\mathfrak{P}'_D$ of L^D above $\mathfrak{p}$ such that $f(\mathfrak{P}'_D/\mathfrak{p}) > 1$.

2. Let D_n be the dihedral group with $2n$ elements, generated by two elements σ and τ such that $\sigma^n = \tau^2 = 1_{D_n}$ and $\tau\sigma\tau^{-1} = \sigma^{-1}$.
   a) Give the complete list of subgroups of D_n.
   b) Find among these subgroups those that are normal.
   c) Find among the normal subgroups H those for which D_n/H is cyclic of prime power order.
   d) Deduce that, as claimed in the text, if n is odd the only normal subgroup H of D_n of prime power order such that D_n/H is cyclic is the group C_n.

3. Let N/K be a Galois extension of number fields with Galois group G, and let L/K be a subextension of N/K. Let $\mathfrak{p}$ be a prime ideal of K, let $\mathfrak{P}_L$ be a prime ideal of L above $\mathfrak{p}$, and let $\mathfrak{P}_N$ be a prime ideal of N above $\mathfrak{P}_L$.
   a) Show that for all k,

$$G_k(\mathfrak{P}_N/\mathfrak{P}_L) = G_k(\mathfrak{P}_N/\mathfrak{p}) \cap \mathrm{Gal}(N/L) \ .$$

b) Assume that L/K is Galois, so that $L = N^H$ for a normal subgroup H of G, and let ϕ be the canonical surjection from $G = \mathrm{Gal}(N/K)$ to $\mathrm{Gal}(L/K) \simeq G/H$. Show that for $-1 \leq k \leq 1$,

$$G_k(\mathfrak{P}_L/\mathfrak{p}) = \phi(G_k(\mathfrak{P}_N/\mathfrak{p})) \ .$$

(Hint: for $k = 0$ and $k = 1$, show that there is an inclusion and then use a counting argument.)

Note that this last result is *not* true in general for $k \geq 2$, but it is true if we replace G_k by the *upper ramification groups* G^k (which are the same groups numbered differently); see [Ser].

4. Using similar techniques as those used in the proofs of Proposition 10.1.16 and Corollary 10.1.17, show (in the given order) the following additional results on ramification groups. Let $\sigma \in G_j$ and $\tau \in G_k$ for j and $k \geq 1$.

a) We have $\sigma\tau\sigma^{-1}\tau^{-1} \in G_{j+k}$ and

$$\theta_{j+k}(\sigma\tau\sigma^{-1}\tau^{-1}) = (k-j)\theta_j(\sigma)\theta_k(\tau) \ .$$

b) The integers $k \geq 1$ such that $G_k \neq G_{k+1}$ are in the same residue class modulo p, where p is the prime number below $\mathfrak{p}$.

c) We have $\sigma\tau\sigma^{-1}\tau^{-1} \in G_{j+k+1}$ (which is of course stronger than the first statement of a)).

5. Using the usual techniques of rational function decomposition, prove that if $f(X) = \prod_{1 \leq k \leq n}(X - \alpha_k)$ with distinct roots α_k, then, as claimed in the text

$$\frac{1}{f(X)} = \sum_{k=1}^{n} \frac{1}{f'(\alpha_k)(X - \alpha_k)} \ .$$

6. Let L/K be an Abelian extension of K of prime power degree ℓ^r, where ℓ is prime, and let $\mathfrak{f}$ be the conductor of L/K. Generalizing Corollary 10.1.24, prove that if $\mathfrak{p}$ is a prime ideal of K above ℓ, then $v_{\mathfrak{p}}(\mathfrak{f}) \leq \lfloor r\ell e(\mathfrak{p}/\ell)/(\ell-1) \rfloor + 1$ (use Exercise 15 of Chapter 3). Deduce from this a slightly weaker version of Hasse's Theorem 5.4.6, at least in the case of extensions of prime-power degree.

7. By giving explicit examples with $\ell = 3$ and $K = \mathbb{Q}$, show that all the decomposition types not excluded by Proposition 10.1.26 can actually occur (case (8), which is the most difficult to attain, occurs for $L = \mathbb{Q}(\theta)$ with θ root of $X^3 - X^2 - 9X + 8$ and $p = 7$).

8. Prove statements (3) and (5) of Theorem 9.2.6 in the case $n = \ell$ prime.

9. Let L/K be a cyclic extension of Galois group G generated by σ.

a) Show the following additive version of Hilbert's Theorem 90 (Lemma 10.2.4). If $\alpha \in L$ is an element of relative trace equal to 0, there exists $\beta \in L$ such that $\alpha = \beta - \sigma(\beta)$. Give an explicit formula for β.

b) Show the following ideal-theoretic version of Hilbert's Theorem 90. If I is a fractional ideal of L of relative norm equal to $\mathbb{Z}_K$, there exists a fractional ideal J of L such that $I = J\sigma(J)^{-1}$.

10. Extend as many results of Sections 10.2.1 and 10.2.2 as you can to the case of noncyclic, but still Abelian, extensions. In particular, generalize Hilbert's theorem 90, Corollary 10.2.7, and the notion of Kummer-equivalence.

11. Let A be a finite Abelian group of exponent dividing n. Show that $\mathrm{Hom}(A, \mu_n) \simeq A$ noncanonically by first proving it for cyclic groups and then using the elementary divisor theorem giving the structure of finite Abelian groups.
12. Let K be a number field, let P be a monic polynomial with coefficients in $\mathbb{Z}_K$, let $\mathfrak{p}$ be a prime ideal of K, and let $\alpha \in K$. Assume that the $\mathfrak{p}$-adic valuations of all the coefficients of P are nonnegative. Show that we also have $v_{\mathfrak{p}}(\alpha) \geq 0$.
13. Let ℓ be a prime number, and let K be a number field containing a primitive ℓth root of unity ζ_ℓ.
    a) Show that for all j coprime to ℓ, the number $(1-\zeta_\ell^j)/(1-\zeta_\ell)$ is a unit in K.
    b) Deduce from this the formula $\ell\mathbb{Z}_K = (1-\zeta_\ell)^{\ell-1}\mathbb{Z}_K$, hence that, as claimed in the text, for any prime ideal $\mathfrak{p}$ of K,
$$v_{\mathfrak{p}}(1-\zeta_\ell) = \frac{v_{\mathfrak{p}}(\ell)}{\ell-1} = \frac{e(\mathfrak{p}/\ell)}{\ell-1} \ .$$
14. Let $\mathfrak{P}$ be a prime ideal above a prime ideal $\mathfrak{p}$ in an extension L/K of number fields. Show that
$$\mathfrak{P}^k \cap \mathbb{Z}_K = \mathfrak{p}^a \quad \text{with} \quad a = \left\lceil \frac{k}{e(\mathfrak{P}/\mathfrak{p})} \right\rceil \ .$$
    In particular, if $\mathfrak{P}$ is unramified, then $\mathfrak{P}^k \cap \mathbb{Z}_K = \mathfrak{p}^k$ (see Proposition 2.3.15).
15. Let K be a number field, let ℓ is a prime number, and set $d = [K(\zeta_\ell) : K]$ and $m = (\ell-1)/d$.
    a) Using Hecke's theorem or otherwise, show that
$$\mathfrak{f}(K(\zeta_3)/K) = \mathfrak{d}(K(\zeta_3)/K) = \prod_{\mathfrak{p}|3,\ 2\nmid e(\mathfrak{p}/3)} \mathfrak{p} \ .$$
    b) More generally, show that for any prime ideal $\mathfrak{p}$ of K above ℓ, we have $m \mid e(\mathfrak{p}/\ell)$, and that
$$\mathfrak{d}(K(\zeta_\ell)/K) = \prod_{\mathfrak{p}|\ell} \mathfrak{p}^{d-(d,e(\mathfrak{p}/\ell)/m)} = \prod_{\mathfrak{p}|\ell} \mathfrak{p}^{(\ell-1-(\ell-1,e(\mathfrak{p}/\ell)))/m} \ .$$
    c) Show that
$$\mathfrak{f}(K(\zeta_\ell)/K) = \prod_{\mathfrak{p}|\ell,\ (\ell-1)\nmid e(\mathfrak{p}/\ell)} \mathfrak{p} \ .$$
16. Let $\mathfrak{p}$ be a prime ideal of K above $\ell \in \mathbb{Z}$.
    a) Show that the map $x \mapsto x^\ell$ from $(\mathbb{Z}_K/\mathfrak{p})^*$ to itself is injective.
    b) Deduce from this that the field $\mathbb{Z}_K/\mathfrak{p}$ is a *perfect field*: in other words, show that the map $x \mapsto x^\ell$ is surjective (the same proof works of course for any finite field).
    c) Let $\overline{\alpha} \in (\mathbb{Z}_K/\mathfrak{p})^*$. Give a formula for the element $x \in \mathbb{Z}_K/\mathfrak{p}$ such that $x^\ell = \overline{\alpha}$, and deduce from this that x can be found in polynomial time.
17. Prove the validity of Algorithm 10.2.14.
18. Fill in the missing details of the proof of Lemma 10.3.3.
19. Show that, as claimed in the text, when s is fixed the functions $F_i(s,x)$ defined in Theorem 10.3.4 tend to 0 faster than any power of x as $x \to \infty$. More precisely, find an upper bound for $|F_i(s,x)|$ of the same type as the upper bound for $|W_i(t)|$ given in Lemma 10.3.3.

20.

a) Show that with a suitable definition of the functions $F_i(s,x)$ for *complex* values of x, the formulas for $\Lambda_i(s)$ given by Theorem 10.3.4 are valid for t_0 belonging to a certain domain of the complex plane.

b) For a fixed value of s, estimate the (complex) value of t_0 that gives the best result for the evaluation of $\Lambda_i(s)$ using Theorem 10.3.4.

21. Give sufficient conditions on the entire function $g(s)$ so that Theorem 10.3.4 remains valid if the gamma products $\gamma_i(s)$ are replaced by the functions $\gamma_1'(s) = g(s)\gamma_1(s)$ and $\gamma_2'(s) = g(k-s)\gamma_2(s)$.

22. Prove the generalized power series expansion of $F(s,x)$ given in the text for s not a pole of $\gamma(s)$. Then by letting s tend to a pole a_0 of $\gamma(s)$, prove the given expansion for $s = a_0$.

23. Show that, as claimed in the text, the series for $F(s,x)$ given in Section 10.3.2 converges for all x.

24. Let E be an elliptic curve defined over $\mathbb{Q}$, let N be its conductor, and let $L(E,s) = \sum_{n\geq 1} a_n n^{-s}$ be its associated L-series (see [Coh0, Definition 7.3.3]). Now that the complete Taniyama–Weil conjecture has been proved, we know that

$$\Lambda(E,s) = N^{s/2}(2\pi)^{-s}\Gamma(s)L(E,s)$$

can be analytically continued to the whole complex plane into a holomorphic function satisfying the functional equation $\Lambda(E,2-s) = w\Lambda(E,s)$ for $w = \pm 1$.

a) Using Theorem 10.3.4, find a rapidly convergent series for $L(E,s)$.

b) Compute $L(E,1)$ when $w = 1$ and $L'(E,1)$ when $w = -1$.

c) More generally, assume that $L^{(k)}(E,1) = 0$ for $0 \leq k \leq r-1$ and that $w = (-1)^r$. Compute $L^{(r)}(E,1)$ using special functions, which you will need to define.

11. Appendix B: Electronic Information

In [Coh0], I have given a rather extensive list of packages more or less related to number theory. While that list is not out of date, it is perhaps worthwhile to insist here on the really useful packages and also on other types of electronically available information. I also mention programs or data not related to the subject matter of this book but important for number theory and for related subjects. Since pricing policies change very rapidly (and free packages sometimes become not free), I have decided to include no pricing information, but only pointers to the relevant web pages or e-mail addresses.

11.1 General Computer Algebra Systems

The packages listed here have the advantage of being able to perform a wide range of symbolic operations, but the applications to number theory almost always suffer from the excessive slowness of these packages (a factor of 100 in speed is not uncommon). Most of these packages are expensive, but there is usually an inexpensive student edition that already has many features. Also, site licenses are often available.

Axiom

This is a huge system containing a wealth of structures and algorithms. However, the author does not advise it for number theory. All information can be obtained from the URL

`http://www.nag.com`

Macsyma

A very robust package, which has undergone enormous improvements in recent years. All information can be obtained from the URL

`http://www.macsyma.com`

There also exist free (although licensed) versions, for instance **Maxima**, maintained by W. Schelter. For more information, ftp the file

`ftp://rascal.ics.utexas.edu/README.MAXIMA`

Maple

Probably the most popular package in the academic community. All information can be obtained from the URL

http://www.maplesoft.com

Mathematica

The most aggressively advertised package but probably not the best, in the author's opinion (who is perhaps biased by the aggressiveness). All information can be obtained from the URL

http://www.wolfram.com

Note. The large commercial packages such as the three "Ms" above can usually be obtained from your local computer software shop or specialized bookstore.

MuPaD

Until very recently, this was the only free general computer algebra system (at least for nonprofit purposes). Unfortunately, it is not completely free anymore, although a free version (MuPaD-lite) is available for nonprofit organizations and students. In any case, it is a very nice system with a Maple-like syntax.

For detailed information about the system, pricing and license conditions, see the URLs

http://www.mupad.de

and

http://www.sciface.com

11.2 Semi-general Computer Algebra Systems

There are two systems in this category.

Magma

This is a huge system, specialized in algebraic structures and morphisms between them and containing an incredible number of algorithms, usually the fastest existing ones. The environment is mathematically rigorous and hence very pleasing to mathematicians. It is very fast, is programmed in a high-level programming language, and has extensive on-line as well as written documentation. Many people developing algorithms and implementations in the fields covered by Magma (which expand frequently) have been asked to contribute their code. It is developed in Sydney by a group directed by J. Cannon. It is not free, but it is *strongly* recommended to acquire at least one license. As for general-purpose computer algebra systems, a student PC version is available for a reasonable price.

A description is given in [Bo-Ca-Pl]. For complete details, consult the home page at

http://www.maths.usyd.edu.au:8000/u/magma/

or send e-mail to

magma@maths.usyd.edu.au

Pari/GP

Developed by the author and collaborators in Bordeaux, this system is copyrighted but free. Contains the implementation of almost all the algorithms presented in [Coh0] and in the present book, plus many other types of algorithms such as multiple precision evaluation of many transcendental functions, elliptic curve computations, and so forth. Complete source code is available. A large part of Pari is included in Magma, and a smaller part is included in MuPaD.

Available by anonymous ftp from the URL

`ftp://megrez.math.u-bordeaux.fr/pub/pari`

The website, maintained by G. Niklasch, is at the URL

`http://hasse.mathematik.tu-muenchen.de/ntsw/pari/`

11.3 More Specialized Packages and Programs

All these packages are free, although usually copyrighted.

Kant/Kash

Developed by M. Pohst and collaborators in Berlin, this is for the moment the only package with Pari/GP that is able to perform heavy-duty work in algebraic number theory. It also contains algorithms for lattices. Almost all of Kant is included in Magma. This package is free, but the source code is not available.

A description is given in [DFKPRSW]. Available by anonymous ftp from the URL

`ftp://ftp.math.tu-berlin.de/pub/algebra/Kant/Kash`

LiDIA

Developed by J. Buchmann and collaborators in Darmstadt. This young package will certainly become a third package able to perform heavy-duty work in algebraic number theory. It is a C++ library that provides multi-precision arithmetic for basic domains, finite fields and classes for integer factorization, factorization of polynomials over finite fields, lattices, matrices, quadratic orders, number fields of larger degree, and elliptic curves.

Like Pari, complete source code is available. Since the developers of LiDIA take among other things Pari as a benchmark for their code, you can expect that the common functions will run faster with LiDIA than with Pari.

Free for noncommercial use, and available by anonymous ftp from the URL

`ftp://ftp.informatik.tu-darmstadt.de/pub/TI/systems/LiDIA`

or from the URL

`http://www.informatik.tu-darmstadt.de/TI/LiDIA`

Simath

Developed by H. Zimmer and collaborators in Saarbrücken. This package is more oriented towards elliptic curve computations.

All information is given at the Simath home page at

`http://emmy.math.uni-sb.de/~simath`

The package can be ftp'ed from the URL

`ftp://ftp.math.uni-sb.de/pub/simath`

or from the mirror site

`ftp://ftp.math.orst.edu/pub/simath`

Ubasic

This is a quite old but very nice and fast Basic-like language for multi-precision computations written by Y. Kida, which works only on DOS-based PCs. Many people love it, although it is quite limited in scope. Many scripts are given with the package, including sophisticated factoring and primality proving algorithms.

Available by anonymous ftp from the URL

`ftp://ftp.math.ohio-state.edu/pub/msdos/ubasic/`

Calc

This package, written (in C) and maintained by K. Matthews, contains refined methods for extended GCD, HNF, SNF, LLL, and MLLL computations which avoid coefficient explosion, as well as a crude MPQS/ECM factoring program. Complete source code is available from the URL

`ftp://www.maths.uq.edu.au/pub/krm/calc/`

and the web page is at the URL

`http://www.maths.uq.edu.au/~krm/krm_calc.html`

(see also the author's home page at the same location).

Miscellaneous

J. Guardia, using an algorithm of J. Montes, has written a program that computes the decomposition type of primes in very high-degree number fields. It is available from the URL

`ftp://drac.mat.ub.es/pub/FiniteFields`

11.4 Specific Packages for Curves

In addition to Magma, Pari, and Simath, which contain many programs for elliptic curves in particular, the following packages are specific to curves.

Apecs

Apecs is a Maple program written by I. Connell for the arithmetic of (plane) elliptic curves. An elliptic curve E can be introduced in various ways (not only by a Weierstrass equation) and essentially over any field. Apecs maintains a

catalog of E defined over $\mathbb{Q}$, and for such E all isogenies over $\mathbb{Q}$ and all the ingredients of the Birch–Swinnerton-Dyer conjecture can be calculated. Together with `mwrank`, this package is very popular among people working on elliptic curves.

Get the `README` file from the URL

`ftp://math.mcgill.ca/pub/apecs/README`

and then get the necessary files.

Mwrank

This standalone C++ program written by J. Cremona is the best existing program for computing the rank and generators of an elliptic curve defined over $\mathbb{Q}$. Several other useful elliptic curve related programs are in the distribution. Like most other standalone programs, `mwrank` can be called directly from a program like Pari/GP. Be sure to obtain the very latest version.

Available from the URL

`http://www.maths.nott.ac.uk/personal/jec/ftp/progs`

More General Curves

A few programs are available for curves of higher genus or for other types of computations on curves.

- V. Flynn's ftp site contains Maple routines for computing with curves of genus 2, available from the URL

  `ftp://ftp.liv.ac.uk/pub/genus2/`

- M. van Hoeij has a Maple package for computing the genus of plane curves, rational parameterizations of genus 0 curves with a point, and Weierstrass forms for genus 1 curves with a point. Available from the URL

  `http://klein.math.fsu.edu/~hoeij/compalg/IntBasis/`

- Q. Liu has a Pari program that computes the local reduction of genus 2 curves (for the moment outside the prime 2) analogous to Tate's algorithm for elliptic curves. It is available from the URL

  `ftp://megrez.math.u-bordeaux.fr/pub/liu/genus2reduction.gz`

11.5 Databases and Servers

Two large databases exist for algebraic number theory: the Berlin and Bordeaux databases. Both contain all the important information about number fields of small degree (about one million number fields of degree up to 7).

The Berlin Database

The first such database from the `Kant/Kash` group is not available by ftp but is directly accessible using `kash`, which opens for you a TCP/IP connection to Berlin (type `?Db` at the kash prompt).

The Bordeaux Database

The second database cannot be queried but is available by anonymous ftp from the URL

`ftp://megrez.math.u-bordeaux.fr/pub/numberfields`

Recover first the file `README.tex` (or `README.dvi`) for a complete description of the contents of the database.

In addition, we know of at least two databases of a different kind related to number fields.

The Munich Database

This database, maintained by G. Niklasch, is much smaller than the previous two but presents a different selection of fields (of larger degrees and small absolute discriminants). It is available at the URL

`http://hasse.mathematik.tu-muenchen.de/nfdb/`

The Arizona Database

This database, maintained by J. Jones (`jj@asu.edu`), lists number fields in an ordering not related to the size of the discriminant, but to their ramification behavior. This is a natural ordering in certain contexts, for example in the number fields occurring in the theory of coverings of the projective line minus 3 points (A. Grothendieck's theory of "dessins d'enfants"). It is available from the URL

`http://math.la.asu.edu/~jj/numberfields`

It contains complete lists of number fields of a given degree which are unramified outside of a small set of small primes. The site provides lists of defining polynomials and many field invariants, with an emphasis on decompositions of associated local algebras.

In addition to databases related to number fields, there are databases related to algebro-geometric objects, essentially curves.

The Nottingham Database

This database, maintained by J. Cremona (`jec@maths.nott.ac.uk`), is the best available database for elliptic curves. It is available by ftp from the URL

`http://www.maths.nott.ac.uk/personal/jec/ftp/data`

It contains all the important arithmetical and analytical data for elliptic curves of conductor up to 6000.

The Liverpool Database

Maintained by V. Flynn, it contains data about curves of genus 2. For the moment it is still very fragmentary. As with Flynn's programs, it is available from the URL

`ftp://ftp.liv.ac.uk/pub/genus2/`

In most of the sites mentioned above, one can also find miscellaneous information. For example, at the Bordeaux site, one can find errata sheets

for the author's books, and at the Nottingham site, one can find errata sheets for Cremona's book.

11.6 Mailing Lists, Websites, and Newsgroups

Several mailing lists, websites, and newsgroups are useful for number theory.

The Number Theory Mailing List

The most important mailing list is the number theory mailing list maintained by V. Miller. To subscribe or send a message to the list, connect to

`http://listserv.nodak.edu/archives/nmbrthry.html`

and follow the links on the top of the page.

The Pari Mailing Lists

There are three mailing lists concerning the Pari package, maintained by D. Bernstein (and archived on the Pari website; see above). The first one (`pari-users`) is a general forum for discussing the Pari package itself and difficulties users may have in installing or using the package. The second one (`pari-dev`) is much more technical and is a forum devoted to exchanges between people involved in the development of the package, including of course bug reports and so on. A user should post in this mailing list only if he or she already has some knowledge of the inner workings of the package. The third one (`pari-announce`) can and should be subscribed to, but is reserved for announcements of new versions of the package and related information, posted by the Pari developers or close collaborators.

To subscribe, send an e-mail to

`pari-xxx@list.cr.yp.to`

where `xxx` should be replaced appropriately by `announce`, `users`, or `dev`. The same e-mail address is used for posting a message to the list.

The Munich Website

This website is a server maintained by G. Niklasch which comprises the Pari website, the Munich number field database and related auxiliary data, and links to further number-theoretical online information. The URL is

`http://hasse.mathematik.tu-muenchen.de`

The Number Theory Website

Maintained by K. Matthews (`krm@maths.uq.edu.au`), in addition to scientific material this site contains information about people, conferences, book announcements, and so forth, related to number theory. The URL is

`http://www.math.uga.edu/~ntheory/web.html`

The L-function Website

Maintained by B. Conrey and D. Farmer, this contains information on L-functions, mainly from the viewpoint of Selberg's attempt to characterize functions that should satisfy the Riemann hypothesis. The URL is

`http://www.math.okstate.edu/~loriw/start/start.html`

11.7 Packages Not Directly Related to Number Theory

Of course, a large number of such packages exist. We mention here a few that may be useful from time to time in number-theoretical work. Of course, the general or semi-general computer algebra systems contain a lot of programs useful in fields other than number theory. For example, Magma, which is certainly the most advanced in its class, contains a huge amount of programs and algorithms in group theory, finite geometries, combinatorics, geometry of numbers, and so on, and is expanding to include algebraic and arithmetic geometry. In addition, there are the following more specialized systems.

Plouffe and Sloane's Programs and Databases

If you have a sequence of natural numbers that you want to identify, your best bet is to look in the book of S. Plouffe and N. Sloane [Plo-Slo], which completely supersedes a much older book by Sloane with the same title. You can also use an online search by connecting to

`http://www.research.att.com/~njas/sequences`

This page also includes links to two automated e-mail servers, called "sequences" and "superseeker":

`sequences,superseeker@research.att.com`

Both servers search a database containing a more complete and updated Plouffe–Sloane book. However, "superseeker" applies additionally clever algorithms which try to transform your sequence into one it can recognize.

On the other hand, if you want to recognize a real number, you may use an LLL algorithm, as explained in [Coh0, Algorithm 2.7.4], if you think the number is algebraic, but you can also use the huge and powerful database of S. Plouffe (the Inverter) by using the server available at

`http://www.lacim.uqam.ca/pi/`

This server contains a database of almost a hundred million mathematical constants, specialized programs to recognize a real number, and information about records on the computation of mathematical constants.

Shoup's NTL package

This package, developed by V. Shoup, is the most efficient package for factoring polynomials over finite fields. Complete details are available at the URL

`http://www.cs.wisc.edu/~shoup/ntl`

Macaulay 2

This is a package developed by D. Grayson and M. Stillman that is designed to support mathematics research in commutative algebra and algebraic geometry. It features fast algorithms for computing Gröbner bases and projective resolutions, together with an object-oriented interpreted user language that supports high-level mathematical concepts.

For information, send e-mail to `Macaulay2@math.uiuc.edu`, and see the URL at

`http://www.math.uiuc.edu/Macaulay2`

CoCoa

This package is aimed at computations in commutative algebra. It is available by ftp at the URL

`ftp://ideal.dima.unige.it/cocoa`

with the U.S. mirror at

`ftp://ftp.reed.edu/mirrors/cocoa`

The web page is at the URL

`http://cocoa.dima.unige.it/`

with the U.S. mirror at

`http://ftp.reed.edu/cocoa`

Faugère's FGb (Fast Gb)

A professional and orders of magnitude faster version of J-C. Faugère's program Gb for computing Gröbner bases (which is incorporated in some of the larger systems), this is by far the best package for Gröbner bases computations. All information can be found on the author's home page at the URL

`http://posso.lip6.fr/~jcf/`

12. Appendix C: Tables

In this appendix, we group a number of tables related to the subject matter of this book.

12.1 Hilbert Class Fields of Quadratic Fields

Let K be a quadratic field and $L = K(1)$ its Hilbert class field. Recall that Theorem 6.2.8 tells us that $L = KL_K$ for a number field L_K linearly disjoint from K. Thus, to give a relative defining polynomial for L/K we may simply give an absolute defining polynomial for $L_K/\mathbb{Q}$ which will be a relative defining polynomial for L/K with coefficients in $\mathbb{Z}$ and not simply in $\mathbb{Z}_K$.

12.1.1 Hilbert Class Fields of Real Quadratic Fields

In Section 6.1 we saw how to use Stark's conjectures and Stark units to compute Hilbert and ray class fields of totally real fields. Here, we give results obtained in the simplest nontrivial case, that of Hilbert class fields of real quadratic fields. The tables are taken from [Coh-Rob].

In the case of real quadratic fields, the class number is usually rather small and the same number fields L_K occur frequently. Thus, instead of giving a table ordered by discriminant, it is more reasonable to give a table roughly ordered by the complexity of the field L_K, and this is what we will do.

We give the table of Hilbert class fields of all real quadratic fields of discriminant less than or equal to 2000. They have been obtained as follows. We first apply the algorithms described in Section 6.1 to obtain a preliminary relative defining polynomial. We then check that this polynomial is correct (to remove the dependence on GRH), and we compute a defining polynomial for L_K. Finally, we use a strong polynomial reduction algorithm of the Polred type to obtain a polynomial which is as simple as possible.

For each of the 607 real quadratic fields K of discriminant less than 2000, we give a polynomial defining a field L_K over $\mathbb{Q}$ as in Proposition 6.2.8. For the sake of completeness, we recall the list of the 319 fields K with class number equal to 1, for which trivially $L_K = \mathbb{Q}$.

Discriminant of the fields such that $L_K = \mathbb{Q}$										
5	8	12	13	17	21	24	28	29	33	37
41	44	53	56	57	61	69	73	76	77	88
89	92	93	97	101	109	113	124	129	133	137
141	149	152	157	161	172	173	177	181	184	188
193	197	201	209	213	217	233	236	237	241	248
249	253	268	269	277	281	284	293	301	309	313
317	329	332	337	341	344	349	353	373	376	381
389	393	397	409	412	413	417	421	428	433	437
449	453	457	461	472	489	497	501	508	509	517
521	524	536	537	541	553	556	557	569	573	581
589	593	597	601	604	613	617	632	633	641	649
652	653	661	664	668	669	673	677	681	701	709
713	716	717	721	737	749	753	757	764	769	773
781	789	796	797	809	813	821	824	829	844	849
853	856	857	869	877	881	889	893	908	913	917
921	929	933	937	941	953	956	973	977	989	997
1004	1013	1021	1033	1041	1048	1049	1052	1057	1061	1069
1077	1081	1084	1097	1109	1112	1117	1121	1132	1133	1137
1141	1149	1153	1169	1177	1181	1193	1201	1208	1213	1217
1228	1237	1244	1249	1253	1273	1277	1289	1293	1301	1317
1321	1324	1329	1333	1336	1337	1349	1357	1361	1381	1388
1389	1397	1401	1409	1432	1433	1437	1441	1453	1457	1461
1468	1473	1477	1481	1493	1497	1501	1516	1528	1529	1532
1541	1549	1553	1561	1569	1577	1589	1592	1597	1609	1613
1621	1633	1637	1657	1661	1669	1673	1676	1688	1689	1693
1697	1709	1713	1721	1724	1733	1741	1753	1757	1777	1784
1789	1793	1797	1801	1816	1817	1821	1829	1837	1841	1852
1857	1861	1868	1873	1877	1889	1893	1909	1912	1913	1916
1933	1941	1948	1949	1964	1969	1973	1977	1981	1993	1997

There are 194 fields with class number 2. We give a table for each possible value of the discriminant $d(L_K)$ of L_K.

First, there are 70 real quadratic fields K of discriminant less than 2000 with class number 2 and such that $L_K = Q(\sqrt{5})$.

Discriminant of the fields K such that $L_K = \mathbb{Q}(\sqrt{5})$										
40	60	65	85	105	120	140	165	185	205	220
265	280	285	305	345	365	380	385	440	460	465
485	545	565	620	645	665	685	705	745	760	805
860	865	885	920	965	1005	1065	1085	1165	1180	1185
1205	1240	1245	1265	1285	1340	1385	1405	1420	1465	1505
1545	1565	1580	1585	1605	1645	1660	1685	1720	1865	1880
1905	1945	1965	1985							

There are 34 real quadratic fields K of discriminant less than 2000 with class number 2 and such that $L_K = \mathbb{Q}(\sqrt{2})$.

Discriminant of the fields K such that $L_K = \mathbb{Q}(\sqrt{2})$										
104	136	168	232	264	296	424	456	488	552	584
616	712	744	776	808	872	1032	1064	1128	1192	1256
1416	1448	1544	1576	1608	1672	1704	1832	1864	1896	1928
1992										

There are 14 real quadratic fields K of discriminant less than 2000 with class number 2 and such that $L_K = \mathbb{Q}(\sqrt{3})$.

Discriminant of the fields K such that $L_K = \mathbb{Q}(\sqrt{3})$										
156	204	348	444	492	636	732	1068	1212	1308	1356
1644	1788	1884								

There are 26 real quadratic fields K of discriminant less than 2000 with class number 2 and such that $L_K = \mathbb{Q}(\sqrt{13})$.

Discriminant of the fields K such that $L_K = \mathbb{Q}(\sqrt{13})$										
221	273	312	364	377	429	481	533	572	728	741
949	988	1001	1144	1157	1196	1209	1261	1417	1469	1612
1729	1781	1833	1976							

There are 21 real quadratic fields K of discriminant less than 2000 with class number 2 and such that $L_K = \mathbb{Q}(\sqrt{17})$.

Discriminant of the fields K such that $L_K = \mathbb{Q}(\sqrt{17})$										
357	408	476	493	561	629	748	952	969	1037	1173
1241	1309	1496	1513	1564	1581	1649	1717	1853	1921	

There remain 29 fields K with class number 2 and such that $d(L_K) > 17$. We give them in a single table containing first the discriminant of K and then $d(L_K)$, ordered by increasing value of $d(L_K)$.

Discriminant of the fields K and $d(L_K)$ for $d(L_K) > 17$											
609	21	861	21	1113	21	1281	21	1533	21	1869	21
696	24	888	24	984	24	1272	24	1464	24	812	28
1036	28	1148	28	1484	28	957	29	1073	29	1189	29
1276	29	1537	29	1624	29	1653	29	1769	29	1353	33
1749	33	1517	37	1628	37	1961	37	1804	41		

There are 24 real quadratic fields with class number equal to 3 and discriminant less than 2000. In the following table, we give their discriminants together with a polynomial defining the field L_K.

Discriminants of the fields K such that $h_K = 3$ and polynomials for L_K			
229	$X^3 - 4X - 1$	257	$X^3 - X^2 - 4X + 3$
316	$X^3 - X^2 - 4X + 2$	321	$X^3 - X^2 - 4X + 1$
469	$X^3 - X^2 - 5X + 4$	473	$X^3 - 5X - 1$
568	$X^3 - X^2 - 6X - 2$	733	$X^3 - X^2 - 7X + 8$
761	$X^3 - X^2 - 6X - 1$	892	$X^3 - X^2 - 8X + 10$
993	$X^3 - X^2 - 6X + 3$	1016	$X^3 - X^2 - 6X + 2$
1101	$X^3 - X^2 - 9X + 12$	1229	$X^3 - X^2 - 7X + 6$
1257	$X^3 - X^2 - 8X + 9$	1304	$X^3 - 11X - 2$
1373	$X^3 - 8X - 5$	1436	$X^3 - 11X - 12$
1489	$X^3 - X^2 - 10X - 7$	1509	$X^3 - X^2 - 7X + 4$
1772	$X^3 - X^2 - 12X + 8$	1901	$X^3 - X^2 - 9X - 4$
1929	$X^3 - X^2 - 10X + 13$	1957	$X^3 - X^2 - 9X + 10$

There are 41 real quadratic fields with class number equal to 4 and discriminant less than 2000. In the following table, we give their discriminants together with a polynomial defining the field L_K.

Discriminants of the fields K such that $h_K = 4$ and polynomials for L_K			
145	$X^4 - X^3 - 3X^2 + X + 1$	328	$X^4 - 2X^3 - 3X^2 + 2X + 1$
445	$X^4 - X^3 - 5X^2 + 2X + 4$	505	$X^4 - 2X^3 - 4X^2 + 5X + 5$
520	$X^4 - 6X^2 + 4$	680	$X^4 - 6X^2 + 4$
689	$X^4 - X^3 - 5X^2 + X + 1$	777	$X^4 - 2X^3 - 4X^2 + 5X + 1$
780	$X^4 - 2X^3 - 7X^2 + 8X + 1$	793	$X^4 - X^3 - 6X^2 + 8X - 1$
840	$X^4 - 6X^2 + 4$	876	$X^4 - 7X^2 - 6X + 1$
897	$X^4 - 2X^3 - 4X^2 + 5X + 3$	901	$X^4 - 2X^3 - 4X^2 + 5X + 2$
905	$X^4 - X^3 - 7X^2 + 3X + 9$	924	$X^4 - 5X^2 + 1$
1020	$X^4 - 2X^3 - 7X^2 + 8X + 1$	1045	$X^4 - X^3 - 8X^2 + X + 11$
1096	$X^4 - 2X^3 - 5X^2 + 6X + 7$	1105	$X^4 - 9X^2 + 4$
1145	$X^4 - X^3 - 8X^2 + 6X + 11$	1160	$X^4 - 6X^2 + 4$
1164	$X^4 - 2X^3 - 7X^2 + 8X + 4$	1221	$X^4 - X^3 - 10X^2 + X + 1$
1288	$X^4 - 2X^3 - 7X^2 + 8X + 8$	1292	$X^4 - X^3 - 11X^2 + 12X + 8$
1313	$X^4 - X^3 - 8X^2 - 4X + 3$	1320	$X^4 - 6X^2 + 4$
1365	$X^4 - 9X^2 + 4$	1480	$X^4 - 6X^2 + 4$
1560	$X^4 - 9X^2 + 4$	1640	$X^4 - 6X^2 + 4$
1677	$X^4 - X^3 - 7X^2 + 2X + 4$	1736	$X^4 - 2X^3 - 7X^2 + 6X + 9$
1740	$X^4 - 2X^3 - 7X^2 + 8X + 1$	1745	$X^4 - X^3 - 10X^2 + 2X + 19$
1752	$X^4 - 2X^3 - 5X^2 + 6X + 3$	1820	$X^4 - 9X^2 + 4$
1848	$X^4 - 10X^2 + 4$	1885	$X^4 - 9X^2 + 4$
1932	$X^4 - 5X^2 + 1$		

Finally, there are 29 real quadratic fields with class number greater than or equal to 5 and discriminant less than 2000. In the following table, we give their discriminants together with a polynomial defining the field L_K.

Discriminants of the fields K such that $h_K \geq 5$ and polynomials for L_K	
401	$X^5 - X^4 - 5X^3 + 4X^2 + 3X - 1$
577	$X^7 - 2X^6 - 7X^5 + 10X^4 + 13X^3 - 10X^2 - X + 1$
697	$X^6 - 3X^5 - 3X^4 + 11X^3 - X^2 - 5X + 1$
785	$X^6 - X^5 - 8X^4 + 6X^3 + 16X^2 - 10X - 5$
817	$X^5 - X^4 - 6X^3 + 5X^2 + 3X - 1$
904	$X^8 - 2X^7 - 9X^6 + 10X^5 + 22X^4 - 14X^3 - 15X^2 + 2X + 1$
940	$X^6 - 3X^5 - 5X^4 + 14X^3 + 9X^2 - 15X - 5$
985	$X^6 - 3X^5 - 4X^4 + 13X^3 + 3X^2 - 10X + 1$
1009	$X^7 - X^6 - 9X^5 + 2X^4 + 21X^3 + X^2 - 13X - 1$
1093	$X^5 - 8X^3 - 3X^2 + 10X + 4$
1129	$X^9 - 3X^8 - 10X^7 + 38X^6 + 5X^5 - 107X^4 + 58X^3 + 78X^2 - 60X - 1$
1297	$X^{11} - 5X^{10} - 4X^9 + 54X^8 - 53X^7 - 127X^6 + 208X^5 + 69X^4 - 222X^3 + 29X^2 + 56X - 5$
1345	$X^6 - 3X^5 - 8X^4 + 16X^3 + 24X^2 - 5$
1384	$X^6 - 2X^5 - 7X^4 + 14X^3 + 3X^2 - 12X + 4$
1393	$X^5 - X^4 - 7X^3 + 6X^2 + 3X - 1$
1429	$X^5 - X^4 - 13X^3 + 23X^2 + 9X - 23$
1596	$X^8 - 2X^7 - 13X^6 + 16X^5 + 43X^4 - 10X^3 - 34X^2 - 4X + 4$
1601	$X^7 - 2X^6 - 14X^5 + 34X^4 + 4X^3 - 38X^2 + 7X + 1$
1641	$X^5 - X^4 - 10X^3 + X^2 + 21X + 9$
1705	$X^8 - X^7 - 14X^6 + 9X^5 + 62X^4 - 23X^3 - 84X^2 + 20X - 1$
1708	$X^6 - 3X^5 - 8X^4 + 21X^3 - 6X^2 - 5X + 1$
1756	$X^5 - 2X^4 - 10X^3 + 14X^2 + 21X - 16$
1761	$X^7 - 2X^6 - 14X^5 + 14X^4 + 50X^3 - 22X^2 - 51X - 3$
1765	$X^6 - 3X^5 - 6X^4 + 17X^3 + 5X^2 - 14X + 4$
1768	$X^8 - 4X^7 - 6X^6 + 32X^5 - 5X^4 - 48X^3 + 14X^2 + 16X - 4$
1785	$X^8 - 2X^7 - 13X^6 + 17X^5 + 48X^4 - 23X^3 - 33X^2 + 3X + 1$
1897	$X^5 - X^4 - 13X^3 + 8X^2 + 27X + 1$
1937	$X^6 - 10X^4 + 25X^2 - 13$
1996	$X^5 - 9X^3 - 4X^2 + 10X + 4$

12.1.2 Hilbert Class Fields of Imaginary Quadratic Fields

In Section 6.3 we saw how to use complex multiplication and modular functions to compute Hilbert and ray class fields of imaginary quadratic fields. Here, we give some results obtained in the simplest nontrivial case, that of Hilbert class fields.

Contrary to the case of real quadratic fields where the class number is usually rather small, in the case of imaginary quadratic fields the class number is roughly of the order of the square root of the discriminant, hence grows quite rapidly. This does not prevent the complex multiplication methods to be very efficient (even more so than the methods using Stark's conjectures for the real case), but the results that we obtain become large quite rapidly, even after applying polynomial reduction algorithms. Thus we give tables only for discriminants in absolute value less than or equal to 451 (corresponding to four pages of the present book). The polynomials have been obtained as follows. We use Schertz's method described in Section 6.3 with several choices of the primes p and q, and we choose the one giving the smallest polynomial in some sense (in fact, for the T_2-norm always used in number field normalizations; see [Coh0, Section 4.4.2]). We then use an absolute polynomial reduction algorithm to obtain one of the polynomials with smallest T_2-norm. This is feasible in practice up to degree 30, but in the range of our table this is not a problem since the largest degree is equal to 21.

Discriminants and Hilbert class fields of imaginary quadratic fields	
-3	X
-4	X
-7	X
-8	X
-11	X
-15	$X^2 - X - 1$
-19	X
-20	$X^2 - X - 1$
-23	$X^3 - X^2 + 1$
-24	$X^2 - 2$
-31	$X^3 + X - 1$
-35	$X^2 - X - 1$
-39	$X^4 - X^3 - X^2 + X + 1$
-40	$X^2 - X - 1$
-43	X
-47	$X^5 - 2X^4 + 2X^3 - X^2 + 1$
-51	$X^2 - X - 4$
-52	$X^2 - X - 3$
-55	$X^4 - X^3 + X^2 + X + 1$
-56	$X^4 - X^3 + X + 1$
-59	$X^3 + 2X - 1$
-67	X
-68	$X^4 + X^2 - 2X + 1$
-71	$X^7 - X^6 - X^5 + X^4 - X^3 - X^2 + 2X + 1$
-79	$X^5 - X^4 + X^3 - 2X^2 + 3X - 1$
-83	$X^3 - X^2 + X - 2$
-84	$X^4 - X^2 + 1$
-87	$X^6 - X^5 + 4X^4 - 4X^3 + 5X^2 - 3X + 1$
-88	$X^2 - 2$
-91	$X^2 - X - 3$
-95	$X^8 - X^7 + X^5 - 2X^4 - X^3 + 2X^2 + 2X - 1$
-103	$X^5 - 2X^4 + 3X^3 - 3X^2 + X + 1$
-104	$X^6 - X^5 + 2X^4 + X^3 - 2X^2 - X - 1$
-107	$X^3 - X^2 + 3X - 2$
-111	$X^8 - 3X^7 + 3X^6 - 3X^5 + 5X^4 - 6X^3 + 6X^2 - 3X + 1$
-115	$X^2 - X - 1$
-116	$X^6 - 2X^3 + X^2 + 2X + 2$
-119	$X^{10} - X^9 + 2X^8 - 4X^7 + 5X^6 - 7X^5 + 9X^4 - 8X^3 + 5X^2 - 4X + 1$
-120	$X^4 - X^3 + 2X^2 + X + 1$
-123	$X^2 - X - 10$
-127	$X^5 - X^4 - 2X^3 + X^2 + 3X - 1$
-131	$X^5 - X^4 + 2X^3 - X^2 + X + 2$
-132	$X^4 - X^2 + 1$

Discriminants and Hilbert class fields of imaginary quadratic fields	
-136	$X^4 - 2X^3 + X^2 + 2X + 1$
-139	$X^3 - X^2 + X + 2$
-143	$X^{10} - 3X^9 + 6X^8 - 6X^7 + 3X^6 + 3X^5 - 9X^4 + 13X^3 - 12X^2 + 6X - 1$
-148	$X^2 - X - 9$
-151	$X^7 - X^6 + X^5 + 3X^3 - X^2 + 3X + 1$
-152	$X^6 + 2X^4 - 2X^3 - 2X^2 - 1$
-155	$X^4 - 2X^3 + 2X^2 - X + 8$
-159	$X^{10} - X^9 + 3X^8 - 4X^7 + 2X^6 - 2X^5 + X^4 + 6X^3 + 4X^2 + 6X + 3$
-163	X
-164	$X^8 + 3X^6 - 2X^3 - X^2 + 2X + 1$
-167	$X^{11} - X^{10} + 5X^9 - 4X^8 + 10X^7 - 6X^6 + 11X^5 - 7X^4 + 9X^3 - 4X^2 + 2X + 1$
-168	$X^4 - X^3 - X^2 - 2X + 4$
-179	$X^5 - X^4 + 3X^2 - X + 2$
-183	$X^8 + 5X^4 - 9X^3 + 6X^2 - 3X + 1$
-184	$X^4 - X^3 + 4X^2 + X + 1$
-187	$X^2 - X - 4$
-191	$X^{13} - 2X^{12} + 4X^{10} - 5X^9 + X^8 + 5X^7 - 11X^6 + 19X^5 - 22X^4 + 16X^3 - 10X^2 + 6X - 1$
-195	$X^4 - X^3 + 2X^2 + X + 1$
-199	$X^9 - X^8 - 3X^6 + 3X^3 + 3X^2 + 5X + 1$
-203	$X^4 - X^3 - 3X^2 - 2X + 4$
-211	$X^3 - 2X - 3$
-212	$X^6 - 2X^5 - 2X^4 + 6X^3 - 2X^2 - 4X + 5$
-215	$X^{14} - 2X^{13} + 6X^{11} - 3X^{10} - 8X^9 + 13X^8 + 4X^7 - 16X^6 + 7X^5 + 13X^4 - 11X^3 - 4X^2 + 6X - 1$
-219	$X^4 - X^3 + 5X^2 - 2X + 4$
-223	$X^7 + X^5 - 4X^4 - X^3 + 5X + 1$
-227	$X^5 - X^3 - 2X^2 + 3X + 4$
-228	$X^4 - X^2 + 1$
-231	$X^{12} - X^{11} - 2X^{10} - 5X^9 + 7X^8 + 8X^7 - 7X^6 + 8X^5 + 7X^4 - 5X^3 - 2X^2 - X + 1$
-232	$X^2 - X - 7$
-235	$X^2 - X - 1$
-239	$X^{15} - 4X^{14} + 4X^{13} + 4X^{12} - 5X^{11} - 13X^{10} + 20X^9 + 4X^8 - 15X^7 - 13X^6 + 27X^5 - 4X^4 - 8X^3 - 2X^2 + 6X - 1$
-244	$X^6 - 2X^5 - X^4 + 4X^3 + 3X^2 - 6X + 2$
-247	$X^6 - 3X^5 + 6X^4 - 7X^3 + 7X^2 - 4X - 1$
-248	$X^8 - 3X^7 + 3X^6 - 2X^5 + 2X^4 + 2X^3 + 3X^2 + 3X + 1$
-251	$X^7 + 4X^5 - 2X^4 + 2X^3 - 3X + 2$

Discriminants and Hilbert class fields of imaginary quadratic fields	
-255	$X^{12} - X^{11} - 4X^{10} + 10X^9 - 3X^8 - 14X^7 + 25X^6 - 25X^5 + 18X^4 - 7X^3 + 2X^2 - 2X + 1$
-259	$X^4 - 2X^3 + X + 2$
-260	$X^8 - 2X^7 + 3X^6 + 2X^5 - 4X^3 + 7X^2 - 4X + 1$
-263	$X^{13} - 6X^{12} + 15X^{11} - 21X^{10} + 19X^9 - 13X^8 + 12X^7 - 22X^6 + 36X^5 - 38X^4 + 27X^3 - 16X^2 + 8X - 1$
-264	$X^8 - 3X^7 + 3X^6 + 2X^4 + 3X^2 - 3X + 1$
-267	$X^2 - X - 22$
-271	$X^{11} + X^9 - X^8 + 3X^7 + 3X^6 - 6X^5 + 3X^4 + 5X^3 - 6X^2 + 5X + 1$
-276	$X^8 - 4X^7 + 11X^6 - 18X^5 + 23X^4 - 18X^3 + 14X^2 - 4X + 4$
-280	$X^4 - 6X^2 + 4$
-283	$X^3 + 4X - 1$
-287	$X^{14} - X^{13} + 3X^{12} + X^{11} - X^{10} + 2X^8 - 6X^7 - 8X^6 + 7X^5 - 5X^4 - 6X^3 + 9X^2 + 8X + 1$
-291	$X^4 - X^3 - 4X^2 + X + 7$
-292	$X^4 - X^2 - 4X + 5$
-295	$X^8 - 4X^7 + 5X^6 - X^5 + 11X^4 - 25X^3 + 14X^2 - X + 1$
-296	$X^{10} - 2X^9 - 3X^8 + 4X^7 + 5X^6 - 2X^5 + 5X^4 + 4X^3 - 3X^2 - 2X + 1$
-299	$X^8 - X^7 - 5X^6 + 6X^5 + 2X^4 - 10X^3 - 7X^2 - X - 1$
-303	$X^{10} - 3X^8 - 2X^7 + 6X^6 - 8X^4 + 21X^3 + 12X^2 - 9X + 9$
-307	$X^3 - X^2 + 3X + 2$
-308	$X^8 - 4X^7 + 12X^6 - 22X^5 + 25X^4 - 18X^3 + 8X^2 - 2X + 5$
-311	$X^{19} - X^{18} + 2X^{17} - 5X^{16} + 8X^{15} - 14X^{14} + 13X^{13} - 10X^{12} - X^{11} + 9X^{10} - 18X^9 + 25X^8 - 10X^7 - 4X^6 + 38X^5 - 42X^4 + 37X^3 - 16X^2 + 4X + 1$
-312	$X^4 - X^3 + 4X^2 + 3X + 9$
-319	$X^{10} - 5X^9 + 11X^8 - 14X^7 + 10X^6 - 2X^5 + X^4 - 5X^3 + 9X^2 - 6X - 1$
-323	$X^4 - 2X^3 - 2X^2 + 3X - 2$
-327	$X^{12} - 2X^{11} + 8X^{10} - 16X^9 + 24X^8 - 28X^7 + 22X^6 - 5X^5 - 9X^4 + 4X^3 + 5X^2 - 4X + 1$
-328	$X^4 - 2X^3 + 3X^2 - 2X + 3$
-331	$X^3 - X^2 + 3X - 4$
-335	$X^{18} + 3X^{16} - 8X^{15} + 14X^{14} - 20X^{13} + 37X^{12} + 10X^{11} + 13X^{10} + 41X^9 - 48X^8 + 16X^7 - 8X^6 + 4X^5 + 45X^4 + 4X^3 + 15X^2 + X + 1$
-339	$X^6 - 4X^4 + 4X^2 + 3$
-340	$X^4 + 3X^2 + 1$
-344	$X^{10} + 4X^8 - 4X^7 + 5X^6 - 8X^5 + 10X^4 - 8X^3 - X^2 - 8X + 2$
-347	$X^5 - X^4 + 4X^3 - X^2 + 5X - 4$
-355	$X^4 + 3X^2 + 20$

Discriminants and Hilbert class fields of imaginary quadratic fields	
-356	$X^{12} - 4X^{10} - 2X^9 + 12X^8 + 20X^7 + 18X^6 + 16X^5 + 20X^4 + 18X^3 + 12X^2 + 4X + 1$
-359	$X^{19} - 2X^{18} + 2X^{17} - 2X^{16} - 3X^{15} + 14X^{14} - 7X^{13} - 22X^{12} + 30X^{11} - 9X^{10} + 5X^9 - 2X^8 - 51X^7 + 90X^6 - 19X^5 - 91X^4 + 113X^3 - 59X^2 + 14X - 1$
-367	$X^9 - X^8 - 3X^7 - X^6 + 6X^5 - X^4 + 6X^3 - 7X^2 + 2X - 3$
-371	$X^8 - 4X^7 + 8X^6 - 10X^5 - 6X^4 + 24X^3 - 27X^2 + 14X - 4$
-372	$X^4 - X^2 + 1$
-376	$X^8 - 5X^6 + 13X^4 - 11X^2 + 18$
-379	$X^3 - X^2 + X - 4$
-383	$X^{17} - X^{16} - X^{15} - X^{14} + X^{12} + 13X^{11} + 7X^{10} + 11X^9 + 4X^8 + X^7 + 7X^6 + 23X^5 + 31X^4 + 42X^3 + 24X^2 + 6X - 1$
-388	$X^4 - 2X^3 - 3X^2 + 4X + 5$
-391	$X^{14} - 5X^{12} - 5X^{11} + 7X^{10} + 15X^9 + 6X^8 - 8X^7 - 3X^6 + 4X^5 - 12X^3 - 12X^2 - 8X + 1$
-395	$X^8 - 4X^7 + X^6 + 11X^5 + 4X^4 - 31X^3 + 38X^2 - 20X + 16$
-399	$X^{16} - 3X^{15} + 6X^{14} - 3X^{13} + 7X^{12} - 12X^{11} - 3X^{10} - 3X^9 + 30X^8 - 42X^7 + 21X^6 + 12X^5 - 11X^4 - 6X^3 + 12X^2 - 6X + 1$
-403	$X^2 - X - 3$
-404	$X^{14} - 2X^{13} + 8X^{12} - 18X^{11} + 33X^{10} - 66X^9 + 99X^8 - 136X^7 + 188X^6 - 192X^5 + 181X^4 - 150X^3 + 86X^2 - 28X + 4$
-407	$X^{16} - X^{15} + 2X^{14} - X^{13} + 9X^{12} + 2X^{11} + 15X^{10} + 12X^8 + 4X^6 - 19X^5 - 17X^4 - 33X^3 - 4X^2 - 10X - 1$
-408	$X^4 + 2X^2 + 4$
-411	$X^6 + 5X^4 + 13X^2 + 12$
-415	$X^{10} + 6X^8 - 2X^7 + 14X^6 - 2X^5 + 11X^4 + 7X^3 + X^2 + 10X - 1$
-419	$X^9 - 2X^7 - 4X^6 + 4X^5 + 10X^4 - X^3 - 8X^2 + 8X + 8$
-420	$X^8 - 3X^6 + 8X^4 - 3X^2 + 1$
-424	$X^6 - 3X^5 + 5X^3 - 2X^2 - X - 7$
-427	$X^2 - X - 15$
-431	$X^{21} - 3X^{20} + 6X^{19} - 9X^{18} + 9X^{17} + 4X^{16} - 10X^{15} + 36X^{14} - 30X^{13} + 14X^{12} - 2X^{11} - 66X^{10} + 41X^9 - 83X^8 + 44X^7 - 10X^6 + 21X^5 + 40X^4 + 16X^3 + 15X^2 + 12X + 1$
-435	$X^4 - X^3 + 2X^2 + X + 1$
-436	$X^6 + X^4 - 6X^3 + 13X^2 - 12X + 4$
-439	$X^{15} - 5X^{14} + 11X^{13} - 9X^{12} - 7X^{11} + 17X^{10} - X^9 - 29X^8 + 38X^7 - 13X^6 - 20X^5 + 24X^4 + 7X^3 - 23X^2 + 13X - 1$
-440	$X^{12} - X^{11} + 6X^{10} + X^9 - 2X^8 + 7X^7 + 10X^6 - 7X^5 - 2X^4 - X^3 + 6X^2 + X + 1$
-443	$X^5 - X^4 - X^3 + 5X^2 + 3X - 5$
-447	$X^{14} - 2X^{13} - 3X^{12} + 7X^{11} + 9X^{10} - 34X^9 + 60X^8 - 118X^7 + 174X^6 - 164X^5 + 135X^4 - 113X^3 + 58X^2 - 12X + 3$
-451	$X^6 - 3X^5 + 3X^4 - X^3 + 7X^2 - 7X - 4$

12.2 Small Discriminants

In this section, we give information on the best results about number fields of small discriminant known to the author at the date of writing.

12.2.1 Lower Bounds for Root Discriminants

Assuming the GRH, it is possible to give very good lower bounds for the root discriminant of an algebraic number field, depending on its signature (r_1, r_2). Many people, and especially H. Stark, A. Odlyzko, J.-P. Serre, G. Poitou, and F. Diaz y Diaz, have contributed to this subject (see, for example, [Odl] for an up-to-date account).

It is also possible to give bounds if we do *not* assume the validity of GRH, but these are so far from the smallest known values that it does not seem reasonable to use these bounds.

The GRH bounds that we use are now called (for short) Odlyzko bounds. It is widely believed that these bounds are very close to the actual truth, in other words that it is not possible to substantially improve these bounds, at least in full generality. The knowledgeable reader will understand that the bounds can of course be improved if we have specific knowledge on the splitting of small primes or on the height of the first zeros of the Dedekind zeta function.

We give a table of such bounds up to degree 26 for every signature. A complete table up to degree 100 can be obtained by `ftp` at
`ftp://megrez.math.u-bordeaux.fr/pub/numberfields/odlyzkobounds`
(take also the file `README.odlyzko` from the same directory).

Since they are lower bounds, the values are given rounded down to three decimals after the decimal point.

Odlyzko GRH bounds for $n \leq 26$ and $0 \leq r_2 \leq 6$							
$n \backslash r_2$	0	1	2	3	4	5	6
1	0.997						
2	2.228	1.722					
3	3.633	2.821					
4	5.127	4.038	3.266				
5	6.644	5.326	4.348				
6	8.148	6.643	5.487	4.595			
7	9.617	7.964	6.657	5.622			
8	11.042	9.271	7.838	6.678	5.737		
9	12.418	10.553	9.016	7.749	6.703		
10	13.744	11.805	10.182	8.824	7.685	6.730	
11	15.021	13.023	11.330	9.893	8.673	7.637	
12	16.250	14.206	12.454	10.951	9.661	8.554	7.602
13	17.432	15.353	13.554	11.995	10.644	9.473	8.457
14	18.571	16.465	14.627	13.020	11.617	10.389	9.316
15	19.668	17.542	15.672	14.026	12.577	11.300	10.176
16	20.726	18.586	16.691	15.011	13.523	12.203	11.032
17	21.747	19.597	17.682	15.975	14.452	13.094	11.883
18	22.733	20.578	18.647	16.916	15.365	13.974	12.726
19	23.686	21.528	19.586	17.836	16.260	14.840	13.560
20	24.608	22.450	20.499	18.735	17.138	15.692	14.384
21	25.500	23.346	21.389	19.612	17.997	16.530	15.196
22	26.365	24.215	22.255	20.468	18.839	17.353	15.997
23	27.204	25.060	23.099	21.305	19.663	18.160	16.785
24	28.017	25.881	23.920	22.121	20.469	18.953	17.560
25	28.807	26.680	24.721	22.918	21.259	19.730	18.323
26	29.575	27.457	25.502	23.697	22.031	20.493	19.072

Odlyzko GRH bounds for $n \leq 26$ and $7 \leq r_2 \leq 13$							
$n \backslash r_2$	7	8	9	10	11	12	13
14	8.377						
15	9.184						
16	9.994	9.073					
17	10.802	9.838					
18	11.607	10.603	9.702				
19	12.407	11.367	10.429				
20	13.200	12.127	11.156	10.276			
21	13.984	12.882	11.880	10.969			
22	14.760	13.631	12.601	11.661	10.802		
23	15.526	14.373	13.317	12.350	11.464		
24	16.281	15.107	14.027	13.036	12.125	11.288	
25	17.026	15.832	14.731	13.717	12.783	11.922	
26	17.760	16.548	15.428	14.394	13.438	12.555	11.739

To test whether the Odlyzko bounds are good, we can compare them with actual data coming from known number fields. As already explained in Chapter 9, the number fields are obtained by a variety of methods:

- complete enumeration using the geometry of numbers for degrees up to 8;
- the class field method described in Chapter 5;
- searches for *polynomials* of small discriminant, using resultant or similar methods (see [Sim2]).

We refer the interested reader to [Sim2] for a table of such polynomials. The conclusion is that first of all the Odlyzko bounds are indeed excellent, and second that we do have good methods to construct number fields of small root discriminant since in most cases we approach the GRH bound by less than 2 or 3 percent. An exception is the smallest totally real field in degree 7, which is known, and whose root discriminant is almost 15% above the Odlyzko bound.

12.2.2 Totally Complex Number Fields of Smallest Discriminant

In the totally complex case, we can explicitly give a table of the smallest known discriminants, first because the number of signatures is linear instead of quadratic, second because the discriminants are smaller, and third because the computations are simpler.

Thus, in this section, we give a table of totally complex number fields of degree less than or equal to 80, having the smallest known root discriminant. They have all been obtained by using the class field method described in Section 9.2.1, with the exception of the field in degree 26, found by D. Simon by more elementary methods (see [Sim2]). The table is taken from joint work with F. Diaz y Diaz and M. Olivier (see [Co-Di-Ol4]).

In each case, the congruence subgroup is equal to $P_{\mathfrak{m}}$, and $\mathfrak{m}_\infty$ is the set of all real places of the base field K (if this was not the case, either L would not be totally complex or $\mathfrak{m}$ would not be its conductor).

Two tables are given. The first table lists the absolute degree $[L : \mathbb{Q}]$, the base field K, the modulus $\mathfrak{m}_0$ as a product of prime ideals (written $\mathfrak{P}_p$ to indicate a prime ideal of degree 1 above p and $\mathfrak{p}_p$ a prime ideal of degree 2 above p), the discriminant in factored form, the root discriminant, and the percentage above the Odlyzko bound that we have computed, except for degree 26, where K and $\mathfrak{m}_0$ are irrelevant.

The second table lists the absolute defining polynomials for the above fields L up to degree 36, since above degree 36 (and even the largest degrees below) the polynomials become unwieldy. Almost all of the missing polynomials can be found in [Co-Di-Ol4].

With the exception of degree 26, all the polynomials have been obtained using the methods of Kummer theory described in Chapter 5, and each gives an exercise on the use of Kummer theory for finding such polynomials, as

we saw in the two examples of Section 5.6 (see, for example, Exercise 16 of Chapter 5).

Smallest known totally complex discriminants

n	K	$\mathfrak{m}_0$	$d(L)$	$\sqrt[n]{\lvert d(L)\rvert}$	%
2	X	$\mathfrak{P}_3$	-3	1.732	0.558%
4	X^2-X+1	$\mathfrak{P}_{13}$	$3^2\cdot 13$	3.289	0.685%
6	X^2-X+1	$\mathfrak{P}_{19}$	$-3^3\cdot 19^2$	4.622	0.576%
8	X^2+1	$\mathfrak{P}_{17}$	$2^8\cdot 17^3$	5.787	0.856%
10	X^5-X^2+1	$\mathfrak{P}_{23}$	$-7^2\cdot 23\cdot 431^2$	6.793	0.939%
12	$X^6-X^5+2X^3-2X^2+1$	$\mathfrak{P}_{41}$	$37^2\cdot 41\cdot 857^2$	7.666	0.843%
14	X^2-X+18	(1)	-71^7	8.426	0.581%
16	X^4-X-1	$\mathfrak{P}_{17}\mathfrak{P}_{37}$	$17^2\cdot 37^2\cdot 283^4$	9.179	1.164%
18	$X^6-2X^5+3X^4+X^2+3X+1$	(2)	$-2^{12}\cdot 23^6\cdot 107^3$	9.836	1.378%
20	X^4+1	$\mathfrak{p}_{11}$	$2^{40}\cdot 11^8$	10.438	1.573%
22	X^2-X+2	$\mathfrak{P}_{23}$	$-7^{11}\cdot 23^{10}$	11.003	1.854%
24	$X^4-X^3-X^2+X+1$	$\mathfrak{P}_{397}$	$3^{12}\cdot 13^6\cdot 397^5$	11.441	1.349%
26	***	$-239\cdot 1535761^2\cdot 7036903^2$		12.419	5.788%
28	X^4+2X^2-2X+1	$\mathfrak{P}_{71}$	$2^{28}\cdot 37^7\cdot 71^6$	12.296	1.135%
30	X^5-X-1	$\mathfrak{P}_{307}$	$-19^6\cdot 151^6\cdot 307^5$	12.766	1.721%
32	X^4-X^3+2X+1	$\mathfrak{P}_3\mathfrak{p}_{13}$	$3^{28}\cdot 7^8\cdot 13^{14}$	13.065	1.135%
36	$X^4-X^3+31X^2-24X+252$	(1)	$3^{18}\cdot 4057^9$	13.823	1.709%
40	X^2+2	$\mathfrak{P}_3\mathfrak{P}_3'\mathfrak{P}_{11}$	$2^{60}\cdot 3^{20}\cdot 11^{18}$	14.412	1.543%
44	X^4-X^3+2X+1	$\mathfrak{P}_{463}$	$3^{33}\cdot 7^{11}\cdot 463^{10}$	14.960	1.511%
48	$X^4-X^3+4X^2+3X+9$	$\mathfrak{p}_2\mathfrak{p}_5$	$2^{16}\cdot 3^{24}\cdot 5^{20}\cdot 13^{24}$	15.386	1.006%
52	$X^4-2X^3+21X^2-20X+68$	(1)	$2^{78}\cdot 1009^{13}$	15.941	1.626%
56	X^4-X^3-2X+8	$\mathfrak{P}_2^3$	$2^{49}\cdot 3^{42}\cdot 241^{14}$	16.472	2.283%
60	$X^4-X^3-2X^2+3$	$\mathfrak{p}_{19}$	$3^{30}\cdot 19^{28}\cdot 37^{15}$	16.880	2.352%
64	X^4-2X^3-2X+5	$\mathfrak{P}_2^3\mathfrak{p}_3\mathfrak{p}_3'$	$2^{128}\cdot 3^{48}\cdot 13^{16}$	17.314	2.738%
68	X^4-X+1	$\mathfrak{P}_{647}$	$229^{17}\cdot 647^{16}$	17.838	3.777%
72	X^4+1	$\mathfrak{P}_{577}$	$2^{144}\cdot 577^{17}$	17.948	2.531%
76	$X^4-2X^3+21X^2-20X+32$	(1)	$17^{38}\cdot 433^{19}$	18.808	5.656%
80	X^4+X^2-X+1	$\mathfrak{P}_{641}$	$257^{20}\cdot 641^{19}$	18.583	2.774%

Note that the field in degree 26 obtained with the methods of Simon can also be obtained with the class field method by choosing K defined by a root of the polynomial

$$X^{13}+X^{12}-10X^{11}-8X^{10}+38X^9+22X^8 \\ -69X^7-24X^6+62X^5+7X^4-26X^3+2X^2+4X-1$$

and a modulus $\mathfrak{m}$ such that $\mathfrak{m}_0=\mathfrak{P}_{239}$.

Absolute defining polynomials of L for $2 \le n \le 36$

$x^2 - x + 1$

$x^4 - x^3 - x^2 + x + 1$

$x^6 - x^5 + x^4 - 2x^3 + 4x^2 - 3x + 1$

$x^8 - 2x^7 + 4x^5 - 4x^4 + 3x^2 - 2x + 1$

$x^{10} - 3x^9 + 7x^8 - 11x^7 + 13x^6 - 12x^5 + 9x^4 - 5x^3 + 3x^2 - 2x + 1$

$x^{12} - 2x^{11} + 2x^{10} - x^9 + 2x^8 - 5x^7 + 8x^6 - 7x^5 + 4x^4 - 3x^3 + 4x^2 - 3x + 1$

$x^{14} - 7x^{13} + 25x^{12} - 59x^{11} + 103x^{10} - 141x^9 + 159x^8 - 153x^7 + 129x^6 - 95x^5 + 58x^4 - 27x^3 + 10x^2 - 3x + 1$

$x^{16} + 2x^{14} - x^{13} + 3x^{12} - 4x^{11} + 4x^{10} - 7x^9 + 5x^8 - 7x^7 + 4x^6 - 4x^5 + 3x^4 - x^3 + 2x^2 + 1$

$x^{18} - x^{17} + 3x^{16} + 2x^{15} - x^{14} + 11x^{13} + 3x^{12} + 3x^{11} + 28x^{10} - 18x^9 + 47x^8 - 27x^7 + 45x^6 - 23x^5 + 27x^4 - 11x^3 + 9x^2 - 2x + 1$

$x^{20} - 4x^{19} + 8x^{18} - 8x^{17} - x^{16} + 12x^{15} - 8x^{14} - 16x^{13} + 43x^{12} - 44x^{11} + 24x^{10} - 12x^9 + 24x^8 - 44x^7 + 48x^6 - 36x^5 + 21x^4 - 12x^3 + 8x^2 - 4x + 1$

$x^{22} - 5x^{21} + 13x^{20} - 26x^{19} + 48x^{18} - 82x^{17} + 127x^{16} - 179x^{15} + 238x^{14} - 309x^{13} + 391x^{12} - 475x^{11} + 560x^{10} - 644x^9 + 703x^8 - 690x^7 + 578x^6 - 398x^5 + 220x^4 - 95x^3 + 31x^2 - 7x + 1$

$x^{24} - 6x^{23} + 22x^{22} - 62x^{21} + 146x^{20} - 295x^{19} + 522x^{18} - 829x^{17} + 1191x^{16} - 1559x^{15} + 1874x^{14} - 2078x^{13} + 2127x^{12} - 2007x^{11} + 1752x^{10} - 1403x^9 + 1023x^8 - 683x^7 + 407x^6 - 216x^5 + 103x^4 - 41x^3 + 15x^2 - 4x + 1$

$x^{26} - x^{25} + 3x^{24} - 4x^{23} + 6x^{22} - 8x^{21} + 9x^{20} - 12x^{19} + 12x^{18} - 14x^{17} + 14x^{16} - 14x^{15} + 15x^{14} - 13x^{13} + 15x^{12} - 14x^{11} + 14x^{10} - 14x^9 + 12x^8 - 12x^7 + 9x^6 - 8x^5 + 6x^4 - 4x^3 + 3x^2 - x + 1$

$x^{28} - 6x^{27} + 14x^{26} - 12x^{25} - 15x^{24} + 64x^{23} - 94x^{22} + 38x^{21} + 106x^{20} - 230x^{19} + 198x^{18} + 20x^{17} - 268x^{16} + 324x^{15} - 128x^{14} - 132x^{13} + 241x^{12} - 164x^{11} + 6x^{10} + 82x^9 - 68x^8 + 28x^7 - 2x^6 - 10x^5 + 9x^4 - 2x^3 + 1$

$x^{30} - 5x^{29} + 13x^{28} - 20x^{27} + 22x^{26} - 36x^{25} + 77x^{24} - 141x^{23} + 211x^{22} - 237x^{21} + 247x^{20} - 329x^{19} + 456x^{18} - 543x^{17} + 580x^{16} - 538x^{15} + 327x^{14} - 54x^{13} - 34x^{12} - 85x^{11} + 176x^{10} - 109x^9 + 16x^8 + x^7 + 13x^6 - 9x^5 - 4x^3 + 9x^2 - 5x + 1$

$x^{32} - 5x^{31} + 17x^{30} - 40x^{29} + 77x^{28} - 131x^{27} + 200x^{26} - 295x^{25} + 385x^{24} - 496x^{23} + 575x^{22} - 647x^{21} + 669x^{20} - 585x^{19} + 561x^{18} - 292x^{17} + 323x^{16} + 52x^{15} + 162x^{14} + 183x^{13} + 111x^{12} + 146x^{11} + 92x^{10} + 67x^9 + 31x^8 + 22x^7 + 11x^6 + 11x^5 + 11x^4 + 7x^3 + 8x^2 + 5x + 1$

$x^{36} + 2x^{35} - x^{34} - 6x^{33} - 10x^{32} - 7x^{31} + 6x^{30} + 16x^{29} + 64x^{28} + 18x^{27} - 72x^{26} - 119x^{25} + 140x^{24} + 20x^{23} + 96x^{22} - 528x^{21} + 429x^{20} - 237x^{19} + 613x^{18} - 533x^{17} + 1151x^{16} - 484x^{15} + 664x^{14} - 464x^{13} + 161x^{12} + 1006x^{11} - 1324x^{10} + 716x^9 - 36x^8 - 239x^7 + 245x^6 - 197x^5 + 121x^4 - 55x^3 + 17x^2 - 4x + 1$

The reason we have not given any examples in degrees 34, 38, and so forth, is that they are too far away from the Odlyzko bounds. In fact, the examples given in degrees 26, 68, or 76 are already not very good, but they are the best known to us.

In any case, there is no reason to believe, especially in large degrees, that small discriminants will correspond to Abelian extensions of subfields. Already in degree 26 the best-known example has not been obtained in this way, although the field L is an Abelian extension of a subfield of degree 13. In fact, it is plausible that the Galois group of the fields with smallest discriminants will tend to be the complete symmetric group S_n, which prevents the fields from having nontrivial subfields.

Remark. Although we have never considered this aspect of the question in this book, it should be noted that the search for number fields with small root discriminant is in many ways analogous to the search for curves of small genus having many rational points over a finite field. We refer the interested reader to [Nie] for this very interesting subject.

Bibliography

[Adl-Hua] L. Adleman and M.-D. Huang, editors, *Algorithmic Number Theory Symposium ANTS-I*, Lecture Notes in Comp. Sci. **877**, Springer-Verlag (1994).

[Ami] Y. Amice, *Les nombres p-adiques*, SUP/Le Mathématicien **14**, Presses Universitaires de France (1975).

[Art-Tat] E. Artin and J. Tate, *Class field theory*, Benjamin, New York (1967).

[Bac] G. Bachman, *Introduction to p-adic numbers and valuation theory*, Academic paperbacks, Acad. Press (1964).

[Bac-Sha1] E. Bach and J. Shallit, *Factor refinement*, J. Algorithms **15** (1993), 199–222.

[Bac-Sha2] E. Bach and J. Shallit, *Algorithmic number theory. Volume 1: Efficient algorithms*, MIT Press, Cambridge, MA (1996).

[Bac-Sor] E. Bach, J. Sorenson, *Explicit bounds for primes in residue classes*, Math. Comp. **65** (1996), 1717–1735.

[Bai] A. Baily, *On the density of discriminants of quartic fields*, J. reine angew. Math. **315** (1980), 190–210.

[BBBCO] C. Batut, K. Belabas, D. Bernardi, H. Cohen, and M. Olivier, *User's guide to Pari-GP version 2.x.x*, available by anonymous ftp.

[Bec-Wei] T. Becker and V. Weispfenning, *Gröbner bases, a computational approach to commutative algebra*, Graduate Texts in Math. **141**, Springer-Verlag (1993).

[Bel1] K. Belabas, *A fast algorithm to compute cubic fields*, Math. Comp. **66** (1997), 1213–1237.

[Bel2] K. Belabas, *On the mean 3-rank of quadratic fields*, Compositio Math., to appear.

[Bel3] K. Belabas, *Variations sur un thème de Davenport et Heilbronn*, Thesis, Université Bordeaux I (1996).

[Ber-Mar] A.-M. Bergé and J. Martinet, *Notions relatives de régulateurs et de hauteurs*, Acta Arith. **54** (1989), 155–170.

[Be-Ma-Ol] A.-M. Bergé, J. Martinet, and M. Olivier, *The computation of sextic fields with a quadratic subfield*, Math. Comp. **54** (1990), 869–884.

[Bir] G. Birkhoff, *Subgroups of Abelian groups*, Proc. Lond. Math. Soc. (2) **38** (1934-5), 385–401.

[Bo-Ca-Pl] W. Bosma, J. Cannon and C. Playoust, *The Magma algebra system I: The user language*, J. Symb. Comput. **24** (1997), 235–265.

[Bos-Poh] W. Bosma and M. Pohst, *Computations with finitely generated modules over Dedekind rings*, Proceedings ISSAC'91 (1991), 151–156.

[Bou1] N. Bourbaki, *Algèbre Commutative*, Chapitre VII, Hermann, Paris.

[Bou2] N. Bourbaki, *Algèbre*, Chapitre VIII, Hermann, Paris.

[Bra] B. Braaksma, *Asymptotic expansions and analytic continuations for a class of Barnes integrals*, Compos. Math. **15** (1964), 239–341.

[Buc-For] J. Buchmann and D. Ford, *On the computation of totally real quartic fields of small discriminant*, Math. Comp. **52** (1989), 161–174.

[Bu-Fo-Po] J. Buchmann, D. Ford, and M. Pohst, *Enumeration of quartic fields of small discriminant*, Math. Comp. **61** (1993), 873–879.

[Buh] J. Buhler, editor, *Algorithmic Number Theory Symposium ANTS-III*, Lecture Notes in Comp. Sci. **1423**, Springer-Verlag (1998).

[But] L. Butler, *Subgroup lattices and symmetric functions*, Memoirs of the A.M.S. **539** (1994).

[Cas-Frö] J. Cassels and A. Fröhlich, *Algebraic Number Theory*, Academic Press, London, New York (1967).

[Cas-Jeh] P. Cassou-Noguès and A. Jehanne, *Parité du nombre de classes des S_4-extensions de $\mathbb{Q}$ et courbes elliptiques*, J. Number Theory **57** (1996), 366–384.

[Cav-Lem] S. Cavallar and F. Lemmermeyer, *The Euclidean algorithm in cubic number fields*, Number Theory (Eger, 1996), de Gruyter, Berlin (1998), 123–146.

[Che] C. Chevalley, *Sur la théorie du corps de classe dans les corps finis et les corps locaux*, J. Fac. Sci. Tokyo **2** (1933), 365–475.

[Coh0] H. Cohen, *A Course in Computational Algebraic Number Theory (3rd corrected printing)*, Graduate Texts in Math. **138**, Springer-Verlag (1996).

[Coh1] H. Cohen, *Hermite and Smith normal form algorithms over Dedekind domains*, Math. Comp. **65** (1996), 1681–1699.

[Coh2] H. Cohen, editor, *Algorithmic Number Theory Symposium ANTS-II*, Lecture Notes in Comp. Sci. **1122**, Springer-Verlag (1996).

[Coh-Dia] H. Cohen and F. Diaz y Diaz, *A polynomial reduction algorithm*, Sém. de Théorie des Nombres Bordeaux (Série 2) **3** (1991), 351–360.

[Co-Di-Ol1] H. Cohen, F. Diaz y Diaz, and M. Olivier, *Subexponential algorithms for class group and unit computations*, J. Symb. Comput. **24** (1997), 433–441.

[Co-Di-Ol2] H. Cohen, F. Diaz y Diaz, and M. Olivier, *Computing ray class groups, conductors and discriminants*, Math. Comp. **67** (1998), 773–795.

[Co-Di-Ol3] H. Cohen, F. Diaz y Diaz, and M. Olivier, *Imprimitive octic fields with small discriminant*, Algorithmic Number Theory Symposium ANTS-III (J. Buhler, ed.), Lecture Notes in Comp. Sci. **1423**, Springer-Verlag (1998), 372–380.

[Co-Di-Ol4] H. Cohen, F. Diaz y Diaz, and M. Olivier, *A table of totally complex number fields of small discriminants*, Algorithmic Number Theory Symposium ANTS-III (J. Buhler, ed.), Lecture Notes in Comp. Sci. **1423**, Springer-Verlag (1998), 381–391.

[Co-Di-Ol5] H. Cohen, F. Diaz y Diaz, and M. Olivier, *Computation of relative quadratic class groups*, Algorithmic Number Theory Symposium ANTS-III (J. Buhler, ed.), Lecture Notes in Comp. Sci. **1423**, Springer-Verlag (1998), 433–440.

[Co-Di-Ol6] H. Cohen, F. Diaz y Diaz, and M. Olivier, *Tables of octic fields containing a quartic subfield*, Math. Comp., to appear.

[Co-Di-Ol7] H. Cohen, F. Diaz y Diaz, and M. Olivier, *Algorithmic Methods for Finitely Generated Abelian Groups*, Proc. 2nd Magma Conference, J. Symb. Comput., to appear.

[Coh-Rob] H. Cohen and X.-F. Roblot, *Computing the Hilbert class field of real quadratic fields*, Math. Comp., to appear.

[Con-Slo] J.-H. Conway and N. Sloane, *Sphere packings, lattices and groups (3rd ed.)*, Grundlehren der math. Wiss. **290**, Springer-Verlag, New York (1999).

[Cor-Ros] G. Cornell and M. Rosen, *A note on the splitting of the Hilbert class fields*, J. Number Theory **11** (1988), 152–158.

[Co-Li-Os] D. Cox, J. Little, and D. O'Shea, *Ideals, varieties and algorithms. An introduction to computational algebraic geometry and commutative algebra*, Undergraduate Texts in Math., Springer-Verlag, New York (1992).

[Cre] J. Cremona, *Reduction of cubic and quartic forms*, LMS Journal of Computation and Math. **2** (1999), 62–92.

[Dab1] M. Daberkow, *Bestimmung relativer Ganzheitsbasen in relativquadratischen Zahlkörpern*, Diplomarbeit, Universität Düsseldorf (1993).

[Dab2] M. Daberkow, *Über die Bestimmung der ganzen Elemente in Radikalerweiterungen algebraischer Zahlkörper*, Thesis, Technische Universität Berlin (1995).

[DFKPRSW] M. Daberkow, C. Fieker, J. Klüners, M. Pohst, K. Roegner, M. Schörnig, and K. Wildanger, *KANT V4*, J. Symb. Comput. **24** (1997), 267–283.

[Dat-Wri] B. Datskowsky and D. J. Wright, *Density of discriminants of cubic extensions*, J. reine angew. Math. **386** (1988), 116–138.

[Dav-Hei1] H. Davenport and H. Heilbronn, *On the density of discriminants of cubic fields (I)*, Bull. London Math. Soc. **1** (1969), 345–348.

[Dav-Hei2] H. Davenport and H. Heilbronn, *On the density of discriminants of cubic fields (II)*, Proc. Roy. Soc. London **322** (1971), 405–420.

[Del] C. Delaunay, work in progress.

[Diaz] F. Diaz y Diaz, *A table of totally real quintic number fields*, Math. Comp. **56** (1991), 801–808.

[Dia-Oli] F. Diaz y Diaz and M. Olivier, *Algorithmique algébrique dans les corps de nombres*, Etats de la Recherche en Algorithmique Arithmétique, Bordeaux, Soc. Math. France (1995).

[Dum-Tan] D. Dummit and B. Tangedal, *Computing the leading term of an Abelian L-function*, Algorithmic Number Theory Symposium ANTS-III (J. Buhler, ed.), Lecture Notes in Comp. Sci. **1423**, Springer-Verlag (1998), 400–411.

[Eic-Oli] Y. Eichenlaub and M. Olivier, *Computation of Galois groups for polynomials with degree up to eleven*, submitted.

[Fie] C. Fieker, *Computing class fields via the Artin map*, J. Symb. Comput., submitted.

[Fie-Poh] C. Fieker and M. Pohst, *On lattices over number fields*, Algorithmic Number Theory Symposium ANTS-II (H. Cohen, ed.), Lecture Notes in Comp. Sci. **1122**, Springer-Verlag (1996), 133–139.

[For1] D. Ford, *Enumeration of totally complex quartic fields of small discriminant*, Computational Number Theory (1989) (A. Pethö, M. Pohst, H. C. Williams, and H. Zimmer, eds.), de Gruyter, Berlin and New York (1991), 129–138.

[For-Let] D. Ford and P. Letard, *Implementing the Round Four maximal order algorithm*, J. Théorie des Nombres Bordeaux **6** (1994), 39–80.

[Fri] E. Friedmann, *Hecke's integral formula*, Sém. de Théorie des Nombres de Bordeaux, Exposé No. **5** (1987-1988).

[Frö-Tay] A. Fröhlich and M. Taylor, *Algebraic number theory*, Cambridge Studies in Adv. Math. **27**, Cambridge Univ. Press (1991).

[GCL] K. Geddes, S. Czapor, and G. Labahn, *Algorithms for Computer Algebra*, Kluwer Academic Publishers, Boston, Dordrecht, London (1992).

[Gei] K. Geissler, *Zur Berechnung von Galoisgruppe*, Diplomarbeit, Technische Universität Berlin (1997).

[Gor] P. Gordan, *Beweis, dass jede Covariante und Invariante einer binären Form eine ganze Funktion mit numerischen Coefficienten einer endlichen Anzahl solcher Formen ist*, J. reine angew. Math. **69** (1868), 323–354.

[Gras] G. Gras, *Théorie du corps de classes global*, Faculté des Sciences de Besançon (1979–1980).

[Haf-McC] J. Hafner and K. McCurley, *Asymptotically fast triangularization of matrices over rings*, SIAM J. Comput. **20** (1991), 1068–1083.

[Haj-Mai] F. Hajir and C. Maire, *Tamely ramified towers and discriminant bounds for number fields*, preprint.

[Har] D. Harbater, *Galois groups with prescribed ramification*, Arithmetic Geometry (N. Childress and J. Jones, eds.), Contemp. Math. **174**, American Math. Soc. (1994), 35–60.

[Has1] H. Hasse, *Bericht über neuere Untersuchungen und Probleme aus der Theorie der algebraischen Zahlkörper, Teil 1: Klassenkörpertheorie, Teil 1a: Beweise zu Teil 1, Teil 2: Reziprozitätsgesetz*, Physica-Verlag (1965).

[Has2] H. Hasse, *Zahlentheorie*, Akademie-Verlag GmbH (1949).

[Hav-Maj1] G. Havas and B. Majewski, *Integer matrix diagonalization*, J. Symb. Comput. **24** (1997), 399–408.

[Hav-Maj2] G. Havas, B. Majewski, and K. Matthews, *Extended GCD and Hermite normal form algorithms via lattice basis reduction*, Experiment. Math. **7** (1998), 125–136.

[Hec] E. Hecke, *Lectures on the theory of algebraic numbers*, Graduate Texts in Math. **77**, Springer-Verlag, Berlin, Heidelberg, New York (1981).

[Hop] A. Hoppe, *Normal forms over Dedekind domains, efficient implementation in the computer algebra system KANT*, Thesis, Technische Universität Berlin (1998).

[Jan] G. Janusz, *Algebraic number fields (2nd ed.)*, Graduate Studies in Math. **7**, American Math. Soc. (1996).

[Klu] J. Klüners, *On computing subfields – A detailed description of the algorithm*, J. Théorie des Nombres Bordeaux **10** (1998), 243–271.

[Klu-Poh] J. Klüners and M. Pohst, *On computing subfields*, J. Symb. Comput. **24** (1997), 385–397.

[Kob] N. Koblitz, *p-adic numbers, p-adic analysis, and zeta-functions (2nd edition)*, Graduate Texts in Math. **58**, Springer-Verlag, Berlin, Heidelberg, New York (1984).

[Lan1] S. Lang, *Elliptic functions*, Addison-Wesley, Reading, MA (1973).

[Lan2] S. Lang, *Introduction to modular forms*, Springer-Verlag, Berlin, Heidelberg, New York (1976).

[Lan3] S. Lang, *Algebraic number theory (2nd ed.)*, Graduate Texts in Math. **110**, Springer-Verlag, Berlin, Heidelberg, New York (1994).

[Lav] A. F. Lavrik, *On functional equations of Dirichlet functions*, Math. USSR–Izvestija **1** (1967), 421–432.

[Lem] F. Lemmermeyer, *The Euclidean algorithm in algebraic number fields*, Expo. Math. **13** (1995), 385–416.

[Let] P. Letard, *Construction de corps de nombres de degré 7 et 9*, Thesis, Université Bordeaux I (1995).

[Leu] A. Leutbecher, *Euclidean fields having a large Lenstra constant*, Ann. Inst. Fourier **35** (1985), 83–106.

[Leu-Nik] A. Leutbecher and G. Niklasch, *On cliques of exceptional units and Lenstra's construction of Euclidean fields*, Number Theory, Proceedings Journées Arithmétiques, Ulm 1987 (H. Schlickewei and E. Wirsing, eds.), Lecture Notes in Math. **1380**, Springer-Verlag (1989), 150–178.

[Marc] D. A. Marcus, *Number fields*, Springer-Verlag, New York (1977).

[Mart1] J. Martinet, *Character theory and Artin L-functions*, Algebraic number fields (A. Fröhlich, ed.), Academic Press, London (1977), 1–87.

[Mart2] J. Martinet, *Petits discriminants des corps de nombres*, Journées arithmétiques 1980 (J. V. Armitage, ed.), London Math. Soc. Lecture Notes Ser. **56** (1982), 151–193.

[Mart3] J. Martinet, *Méthodes géométriques dans la recherche des petits discriminants*, Prog. Math. **59**, Birkhäuser, Boston (1985), 147–179.

[Mart4] J. Martinet, *Une introduction à la théorie du corps de classes* (notes de M. Olivier), Ecole doctorale de mathématiques de Bordeaux (1991).

[Mart5] J. Martinet, *Les réseaux parfaits des espaces euclidiens*, Masson, Paris (1996).

[Mar-Pay] J. Martinet et J.-J. Payan, *Sur les extensions cubiques non galoisiennes des rationnels et leur clôture galoisienne*, Jour. reine angew. Math. **228** (1967), 15–37.

[Mon] P. Montgomery, *Partial LLL reduction*, personal communication.

[Nak1] J. Nakagawa, *On the relations among the class numbers of binary cubic forms*, Invent. Math. **134** (1998), 101–138.

[Nak2] N. Nakagoshi, *The structure of the multiplicative group of residue classes modulo* $\mathfrak{p}^{N+1}$, Nagoya Math. J. **73** (1979), 41–60.

[Neu] J. Neukirch, *Class field theory*, Grundlehren der math. Wiss. **280**, Springer-Verlag (1986).

[Nie] H. Niederreiter and C. Xing, *Algebraic curves over finite fields with many rational points*, Number Theory (Eger, 1996), de Gruyter, Berlin (1998), 423–443.

[Odl] A. Odlyzko, *Bounds for discriminants and related estimates for class numbers, regulators, and zeros of zeta functions: A survey of recent results*, Sém. de Théorie des Nombres Bordeaux (Série 2) **2** (1990), 119–141.

[Oli1] M. Olivier, *The computation of sextic fields with a cubic subfield and no quadratic subfield*, Math. Comp. **58** (1992), 419–432.

[Oli2] M. Olivier, *Corps sextiques primitifs*, Ann. Inst. Fourier **40** (1990), 757–767.

[Poh] M. Pohst, *On the computation of number fields of small discriminants including the minimum discriminants of sixth degree fields*, J. Number Theory **14** (1982), 99-117.

[Poh-Zas] M. Pohst and H. Zassenhaus, *Algorithmic algebraic number theory (3rd ed.)*, Cambridge Univ. Press, Cambridge (1993).

[Que] R. Quême, *A computer algorithm for finding new Euclidean number fields*, J. Théorie des Nombres Bordeaux **10** (1998), 33–48.

[Rob] D. Roberts, *Density of cubic field discriminants*, preprint.

[Rob1] X.-F. Roblot, *Algorithmes de factorisation dans les extensions relatives et applications de la conjecture de Stark à la construction des corps de classes de rayon*, Thesis, Université Bordeaux I (1997).

[Rob2] X.-F. Roblot, *Unités de Stark et corps de classes de Hilbert*, C. R. Acad. Sci. Paris **323** (1996), 1165–1168.

[Rob3] X.-F. Roblot, *Stark's Conjectures and Hilbert's Twelfth Problem*, Experiment. Math. , to appear.

[Sch1] R. Schertz, *Zur expliciten Berechnung von Ganzheitbasen in Strahlklassenkörpern über einem imaginär-quadratischen Zahlkörper*, J. Number Theory **34** (1990), 41–53.

[Sch2] R. Schertz, *Galoismodulstruktur und elliptische Funktionen*, J. Number Theory **39** (1991), 285–326.

[Sch3] R. Schertz, *Problèmes de Construction en Multiplication Complexe*, Sém. de Théorie des Nombres Bordeaux (Série 2), **4** (1992), 239–262.

[Sch4] R. Schertz, *Construction of ray class fields by elliptic units*, J. Théorie des Nombres Bordeaux **9** (1997), 383–394.

[Sch5] R. Schertz, *Lower powers of elliptic units*, preprint (1998).

[Sc-Po-Di] A. Schwarz, M. Pohst, and F. Diaz y Diaz, *A table of quintic number fields*, Math. Comp. **63** (1994), 361–376.

[Ser] J.-P. Serre, *Corps locaux (2nd ed.)*, Hermann, Paris (1968). English translation: Graduate Texts in Math. **67**, Springer-Verlag (1979).

[Shim] G. Shimura, *Abelian varieties with complex multiplication and modular functions*, Princeton Univ. Press, Princeton, NJ (1998).

[Shin] T. Shintani, *On zeta-functions associated with the vector space of quadratic forms*, J. Fac. Sci. Univ. Tokyo, Sec. 1a, **22** (1975), 25–66.

[Sie] C.-L. Siegel, *The trace of totally positive and real algebraic integers*, Ann. of Math. **46** (1945), 302–312.

[Sil1] J. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Math. **106**, Springer-Verlag (1986).

[Sil2] J. Silverman, *Advanced Topics in the arithmetic of elliptic curves*, Graduate Texts in Math. **151**, Springer-Verlag (1994).

[Sim1] D. Simon, *Solving relative norm equations in number fields using S-units*, Math. Comp., submitted.

[Sim2] D. Simon, *Construction de polynômes de petit discriminant*, C. R. Acad. Sci. Paris, to appear.

[Sim3] D. Simon, *Equations dans les corps de nombres et discriminants minimaux*, Thesis, Université Bordeaux I (1998).

[Smy1] C. Smyth, *Totally positive algebraic integers of small trace*, Ann. Inst. Fourier **33** (1984), 1–28.

[Smy2] C. Smyth, *The mean value of totally real algebraic integers*, Math. Comp. **42** (1984), 663–681.

[Smy3] C. Smyth, *An inequality for polynomials*, Proceedings of the CTNA Ottawa conference, to appear.

[Suz] H. Suzuki, *A generalization of Hilbert's Theorem 94*, Nagoya Math. J. **121** (1991), 161–169.

[Tat] J. Tate, *Les conjectures de Stark sur les fonctions L d'Artin en $s = 0$*, Progress in Math. **47**, Birkhaüser, Boston (1984).

[Tol] E. Tollis, *Zeros of Dedekind zeta functions in the critical strip*, Math. Comp. **66** (1997), 1295–1321.

[Wal] P. Walker, *Elliptic functions, a constructive approach*, John Wiley and Sons, New York (1996).

[Wri1] D. Wright, *Distribution of discriminants of Abelian extensions*, Proc. London Math. Soc. **58** (1989), 17–50.

[Wri2] D. Wright, personal communication.

[Wri-Yuk] D. Wright and A. Yukie, *Prehomogeneous vector spaces and field extensions*, Invent. math. **110** (1992), 283–314.

[Yam] K. Yamamura, *Maximal unramified extensions of imaginary quadratic number fields of small conductors*, J. Théorie des nombres Bordeaux **9** (1997), 405–448.

[Zan] H. Zantema, *Class numbers and units*, Computational methods in number theory II (H. W. Lenstra and R. Tijdeman, eds.), Math. Centrum tracts **155** (1982), 213–234.

[Zim] H. Zimmer, *Computational Problems, Methods, and Results in Algebraic Number Theory*, Lecture Notes in Math. **262**, Springer-Verlag (1972).

Index of Notation

Symbols

$1_{\mathcal{A}}$	unit element of an Abelian group $\mathcal{A}$, 164
$\mathbf{1}_{\mathcal{A}}$	row vector of unit elements of an Abelian group $\mathcal{A}$, 164
$(1+\mathfrak{a})/(1+\mathfrak{b})$	group useful for computing $(\mathbb{Z}_K/\mathfrak{m})^*$, 188

A

$\lfloor A \rceil$	entries of the matrix A rounded to the nearest integer, 211
$[a, b[$	semi-open interval, also denoted by $[a, b)$ in certain countries, 32
$(A\|B)$	horizontal concatenation of matrices A and B, 164
$\left(\frac{A}{B}\right)$	vertical concatenation of matrices A and B, 164
$A(\chi)$	$:= \prod_{\mathfrak{p}\|\mathfrak{f},\mathfrak{p}\nmid\mathfrak{f}(\chi)}(1-\chi(\mathfrak{p}))$, 300
(A, D_A)	Smith normal form of Abelian group $\mathcal{A}$, 165
(A, I)	pseudo-matrix, 29
(A, I, J)	integral pseudo-matrix for SNF, 43
a_k	kth elementary symmetric functions, 451
$\alpha \equiv 1 \pmod{{}^*\mathfrak{m}}$	multiplicative congruence, see definition, 136
$((\alpha, \mathfrak{a}), (\beta, \mathfrak{b}))$	pseudo–two-element representation of an ideal, 87
$\alpha \equiv \beta \pmod{{}^*\mathfrak{m}}$	multiplicative congruence, see definition, 136
α_j	jth conjugate of α, 451
$\overline{\alpha}$	class of α in $(\mathbb{Z}_K/\mathfrak{m})^*$, 136
$A_{\mathfrak{m}}(L/K)$	Artin group for the modulus $\mathfrak{m}$, kernel of Art, 153
$\mathrm{Art}(\mathfrak{a})$	$:= \left(\frac{L/K}{\mathfrak{a}}\right)$, Artin reciprocity map, 153

B

(B, D_B)	Smith normal form of Abelian group $\mathcal{B}$, 165

C

C_α	characteristic polynomial of α, 70
$C(\chi)$	$:= (\pi^{-m} d(K) \mathcal{N}(\mathfrak{f}(\chi)))^{1/2}$, 300
(C, D_C)	Smith normal form of Abelian group $\mathcal{C}$, 165

$Cl_i(K)$	abbreviation for $Cl_{i,L}(K)$, 348
$Cl_i(L/K)$	relative pseudo-class group, 348
$Cl_{i,L}(K)$	capitulating subgroup of the extension L/K, 348
$Cl_{i,S}(L/K)$	S-relative class group, 380
$Cl(K)$	class group of K, 134
$Cl_{\mathfrak{m}}$	$:= Cl_{\mathfrak{m}}(K)$, 138
$Cl_{\mathfrak{m}}(K)$	ray class group for the modulus $\mathfrak{m}$, 136
$Cl_N(K)$	abbreviation for $Cl_{N,L}(K)$, 349
$Cl_N(L/K)$	relative norm-class group, 349
$Cl_{N,L}(K)$	norm-default quotient of the extension L/K, 349
$Cl_S(K)$	S-class group of K, 373
$\overline{C}$	$:= C/P_{\mathfrak{m}}$, 138
$\mathcal{C}_p$	p-Sylow subgroup of $\mathcal{C}$, 177

D

D_1	set of indices such that $a_{i,i} = 1$ in HNF of a prime ideal, 105
D_p	set of indices such that $a_{i,i} = p$ in HNF of a prime ideal, 105
$\deg(T)$	degree of polynomial T, 51
$\Delta(\mathfrak{a})$	value of the discriminant form on the ideal $\mathfrak{a}$, 318
$\delta(x)$	1 if $x = 0$, 0 if $x \neq 0$, 157
$\mathfrak{d}(L/K)$	relative discriminant ideal, 79
$D(\mathfrak{P}/\mathfrak{p})$	decomposition group at $\mathfrak{P}$ (also $G_{-1}(\mathfrak{P}/\mathfrak{p})$), 151
$\mathrm{diag}((b_i)_i)$	diagonal matrix whose diagonal entries are the b_i, 166
$\mathfrak{d}_i$	elementary divisors of a torsion module, 6
$\mathrm{disc}(L/K)$	$:= (\mathfrak{d}(L/K), \overline{d(L/K)})$, 79
$\mathrm{disc}_T(M)$	$:= (\mathfrak{d}_T(M), d_T(M))$, 28
$D(L/K)$	discriminant of relative Abelian extension together with part at infinity, 158
$\overline{\mathfrak{D}(L/K)}$	relative different of L/K, 96
$\overline{d(L/K)}$	relative discriminant in K^*/K^{*2}, 79
$\overline{\mathfrak{d}_T(M)}$	discriminant ideal of M, 28
$\overline{d_T(M)}$	discriminant of M in K^*/K^{*2}, 28

E

e_i	$:= e(\mathfrak{P}_i/\mathfrak{p})$, 83
ℓ	a prime number, often degree of a field extension, 227
$e(\mathfrak{P}/\mathfrak{p})$	ramification index of $\mathfrak{P}$ above $\mathfrak{p}$, 83
$\eta(\tau)$	Dedekind eta function at τ, 316
$\exp_a(x)$	Artin–Hasse exponential of x, 202, 221
$\exp_{\mathfrak{p}}(x)$	$\mathfrak{p}$-adic exponential of x, 190

F

$\mathfrak{f}$	can be the conductor of a congruence subgroup, 144
$\mathfrak{f}$	can be the index-ideal of a suborder, 79
$F_{\mathcal{B}}(x,y)$	cubic form associated to an integral basis $\mathcal{B}$, 395
$\mathfrak{f}(\chi)$	conductor of the character χ, 146
f_H	content of the Hessian, 408
f_i	$:= f(\mathfrak{P}_i/\mathfrak{p})$, 83
$F_K(x,y)$	cubic form associated to the cubic field K, 396
$\mathfrak{f}(L/K)$	conductor of the extension L/K (usually Abelian), 151
$f(\mathfrak{P}/\mathfrak{p})$	residual degree of $\mathfrak{P}$ above $\mathfrak{p}$, 83

G

G_3	subset of $\mathbb{R}^n$ for Lagrange multiplier method, 456
G_4	subset of $\mathbb{R}^n$ for Lagrange multiplier method, 456
$\mathrm{Gal}(L/K)$	Galois group of L/K, 74
$\Gamma^0(pq)$	upper Cartan congruence subgroup of level pq of $\mathrm{PSL}_2(\mathbb{Z})$, 318
$\widehat{G}$	group of characters of Abelian group G, 157
$\mathfrak{g}_j(A)$	ideal generated by minor-ideals in the last $n+1-j$ rows, 31
$G_k(\mathfrak{P}/\mathfrak{p})$	kth ramification group of $\mathfrak{P}$, 502
$g_{p,q}(\tau)$	$:= \eta(\tau/p)\eta(\tau/q)/(\eta(\tau/pq)\eta(\tau))$, 318
$g_{p,q,e}(\mathfrak{a})$	$:= g_{p,q}(\omega_1/\omega_2)^e$ for a $\mathbb{Z}$-basis of $\mathfrak{a}$, 319
$g_{\mathfrak{p},\mathfrak{q},e}(\mathfrak{a})$	$:= g_{p,q,e}(\mathfrak{a})$ for all suitable $\mathbb{Z}$-bases of $\mathfrak{a}$, 320
$G_{\mathfrak{p}}$	$:= (1+\mathfrak{p})/(1+\mathfrak{p}^k)$, p-part of group $G = (\mathbb{Z}_K/\mathfrak{p}^k)^*$, 189

H

$H_3^-(X)$	number of complex cubic forms of discriminant down to $-X$, 421
$H_3^+(X)$	number of real cubic forms up to discriminant X, 417
H_A	HNF matrix representing subgroup of $\mathcal{A}$, 172
H_B	HNF matrix representing subgroup of $\mathcal{B}$, 172
H_C	HNF matrix representing subgroup of $\mathcal{C}$, 172
$H(F)$	Hessian of F, 393
H_F	$:= -H(F)/4$, 394
$h_i(L/K)$	relative pseudo-class number, 348
$h(K)$	class number of K, 134
$h_{\mathfrak{m},C}$	$:= \lvert Cl_{\mathfrak{m}}/\overline{C}\rvert = \lvert I_{\mathfrak{m}}/C\rvert = [K(\mathfrak{m})^{\overline{C}} : K]$, 142
$h_{\mathfrak{m}}(K)$	$:= \lvert Cl_{\mathfrak{m}}(K)\rvert$, ray class number for the modulus $\mathfrak{m}$, 137
HNF	Hermite normal form, 1
$h_N(L/K)$	relative norm-class number, 349

I

id_P identity map on P, 8
$I(\mathfrak{P}/\mathfrak{p})$ inertia group at $\mathfrak{P}$ (also $G_0(\mathfrak{P}/\mathfrak{p})$), 151
$i_{L/K}$ natural map from K to an extension L, 348
$I_{\mathfrak{m}}$ $:= I_{\mathfrak{m}}(K)$, 135
$I_{\mathfrak{m}}(K)$ group of ideals of K coprime to $\mathfrak{m}$, 135
$I_{\mathfrak{m},L}$ $:= I_{\mathfrak{m}}\mathbf{z}_L(L)$, 153
I_n $n \times n$ identity matrix, 11
$I_{\mathfrak{p}}$ the $\mathfrak{p}$-radical of an order, 102

J

$j(\mathfrak{a})$ value of modular invariant j on the ideal $\mathfrak{a}$, 315
$J(F)$ Jacobian covariant of F, 393

K

$K(1)$ Hilbert class field of K, ray class field for the trivial modulus 1, 134
K_2 quadratic subfield of the Galois closure of a dihedral, usually cubic, extension, 440
K number field or field of fractions of a Dedekind domain, 2
K_j kernel of the map $x \mapsto x^{p^j}$ from $G_{\mathfrak{p}}$ into itself, 197
KM $:= K \otimes_R M$, 6
$K(\mathfrak{m})$ ray class field for the modulus $\mathfrak{m}$, 138
$K^*_{\mathfrak{m}}$ group of $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}}$, 136
K_z Kummer extension $K(\zeta_n)$ of base field K, 227

L

L_2 can be the Galois closure of a dihedral, usually cubic, extension, 440
L^{ab} maximal Abelian subextension of L/K, 216
λ simpler lift in characteristic zero of a multiple of idempotent e_1, 250
λ_0 lift in characteristic zero of a multiple of idempotent e_1, 248
$\Lambda(s,\chi)$ Hecke L-function with gamma and exponential factors, 300
$L(s,\chi)$ pure Hecke L-function, 300
$[\ell]$ ℓth power map, 232
L_K splitting field of the Hilbert class field of K, 312
$[L:K]$ degree of relative extension L/K, 50
$\left(\frac{L/K}{\mathfrak{a}}\right)$ $:= \mathrm{Art}(\mathfrak{a})$, Artin reciprocity map, 153
L^{*K} extended multiplicative group of L by K, 352
LLL Lenstra–Lenstra–Lovász algorithm, 23

$L_{\mathfrak{m}}(\alpha)$	short discrete logarithm of α, 238
$\log_a(x)$	Artin–Hasse logarithm of x, 202, 221
$\log_{\mathfrak{p}}(x)$	$\mathfrak{p}$-adic logarithm of x, 190
$L_S(s,\chi)$	Hecke L-function outside S, 299
L_z	Kummer extension of $L(\zeta_n)$ of extension field L, 227

M

$\mathfrak{m}$	modulus, 135
$\mathfrak{m}_0$	finite part of the modulus $\mathfrak{m}$, 135
$(\mathfrak{m}_1, C_1) \sim (\mathfrak{m}_2, C_2)$	equivalence of congruence subgroups, 140
$M^{\mathbf{adj}}$	adjoint matrix of M, 51
$(\mathfrak{m}, C)$	congruence subgroup C modulo $\mathfrak{m}$, 138
$\mathfrak{m}_\infty$	infinite part of the modulus $\mathfrak{m}$, 135
$\mathfrak{m} \mid \mathfrak{n}$	divisibility relation between moduli, 135
$M/\mathfrak{p}M$	reduction modulo $\mathfrak{p}$ of a module M, 37
M^*	dual of a $\mathbb{Z}_K$-module for the trace, 484
M^t	transpose matrix of M, 51
$M_{\mathbf{tors}}$	torsion submodule of M, 7
μ_b	a useful element of $\mathbb{Z}[G]$ (see text), 250
$\boldsymbol{\mu}(K)$	group of roots of unity in K, 232
$\boldsymbol{\mu}_n = \boldsymbol{\mu}_n(K)$	subgroup of nth roots of unity in K, 495

N

$N_3^-(X)$	number of complex cubic fields of discriminant down to $-X$, 421
$N_3^+(X)$	number of real cubic fields up to discriminant X, 417
$[n]_G$	raising to the nth power in the group G, 348
$\mathcal{N}_{L/K}$	relative norm map from an extension L to K, 348
$\mathcal{N}_{L/K}(\alpha)$	relative norm of α, 76
$\mathcal{N}(\alpha)$	absolute norm of α, 3

O

$\mathcal{O}$	an order, 102
$(\omega, \mathfrak{a})$	pseudo-element, 26
$(\omega_i, \mathfrak{a}_i)$	pseudo-generating set or pseudo-basis, 26

P

$\mathfrak{P}$	a prime ideal above $\mathfrak{p}$, 83
$\mathfrak{p}$	prime ideal of a base field K, 3
$\phi_{C\Phi}$	Davenport–Heilbronn map from cubic fields to cubic forms, 397
$\phi_h(f_1, f_2)$	hth higher covariant made from f_1 and f_2, 392
$\phi(\mathfrak{m})$	Euler ϕ-function for modulus or ideal $\mathfrak{m}$, 137
$\boldsymbol{\Phi}_n(K)$	space of binary forms of degree n on K, 389

$\phi_{\Phi C}$	Davenport–Heilbronn map from cubic forms to cubic fields, 397
π	often a uniformizer at the prime ideal $\mathfrak{p}$, 104
PID	Principal ideal domain, 1
(p^∞, m)	$:= p^{v_p(m)}$, 177
$P(K)$	group of principal ideals of K, 134
$P_{\mathfrak{m}}$	$:= P_{\mathfrak{m}}(K)$, 136
$P_{\mathfrak{m}}(K)$	group of principal ideals $\alpha\mathbb{Z}_K$ with $\alpha \equiv 1 \pmod{{}^*\mathfrak{m}}$, 136
$[p]$	multiplication by p in an Abelian group, 179
$P_{\mathfrak{p},\mathfrak{q},e}(X)$	characteristic polynomial of $g_{\mathfrak{p},\mathfrak{q},e}(\mathbb{Z}_K)$, 321
$\psi_{i,N}$	natural map from $Cl_i(L/K)$ to $Cl_N(L/K)$, 350
$\psi_{N,i}$	natural map from $Cl_N(L/K)$ to $Cl_i(L/K)$, 350
$\mathcal{P}^*$	group of pseudo-principal ideals, 348

Q

$[Q : P]$	index-ideal of P into Q, 15

R

R	often a Dedekind domain, 2
$\mathcal{R}$	resultant, 51
r_1	number of real embeddings of K, 3
r_2	one half of the number of nonreal embeddings of K, 3
r_a	$:= g^a \bmod \ell$, 250
r_c	ℓ-rank of Abelian group C, 181
$R_i(L/K)$	relative regulator associated to $i_{L/K}$, 358
$R_N(L/K)$	relative regulator associated to $\mathcal{N}_{L/K}$, 359
$R/\mathfrak{p}$	residue field at $\mathfrak{p}$, 37
r_u	unit rank of K, 231
r_v	ℓ-rank of the ℓ-Selmer group of K, 231
R'_X	partial derivative of R with respect to X, 53
$\mathcal{R}_Y$	resultant with respect to the variable Y, 53
R'_Z	partial derivative of R with respect to Z, 53

S

S	usually a finite set of places of a number field, 2
S_0	usually a set of prime ideals (or finite places) in S, 4
S_∞	set of embeddings in S, 4
$s(\alpha)$	vector of signs of embeddings of α, 186
S_ℓ	set of $\mathfrak{p} \nmid \mathfrak{m}$, $\mathfrak{p} \mid \ell$, 227
$S_\varnothing$	set of $\mathfrak{p} \nmid \mathfrak{m}$, $\mathfrak{p} \nmid \ell$, 227
σ_i	one of the real or complex embeddings of a number field, 3
$\sigma_{i,K}$	relative K-linear embeddings, 73

$\sigma_{\mathfrak{P}}(x)$ Frobenius homomorphism for $\mathfrak{P}$, 152
$\sigma_{\mathfrak{p}}(x)$ common Frobenius homomorphism for all $\mathfrak{P}$ above $\mathfrak{p}$, 152
$\sigma(z, L)$ Weierstrass σ-function of a lattice L at z, 328, 334
$s_k = s_k(\alpha)$ kth power sum of conjugates of α, 451
$S_{\mathfrak{m}}$ set of $\mathfrak{p} \mid \mathfrak{m}$, $\mathfrak{p} \nmid \ell$, 227
$s_{\mathfrak{m}_1,\mathfrak{m}_2}$ canonical surjective map from $I_{\mathfrak{m}_1}$ to $I_{\mathfrak{m}_2}$ if $\mathfrak{m}_2 \mid \mathfrak{m}_1$, 146
$S_{\mathfrak{m},\ell,1}$ set of $\mathfrak{p} \mid (\mathfrak{m}, \ell)$ such that $v_{\mathfrak{p}}(\mathfrak{m}) = z(\mathfrak{p}, \ell)$, 227
$S_{\mathfrak{m},\ell,2}$ set of $\mathfrak{p} \mid (\mathfrak{m}, \ell)$ such that $v_{\mathfrak{p}}(\mathfrak{m}) < z(\mathfrak{p}, \ell)$, 227
$S_{\mathfrak{m},\ell,3}$ set of $\mathfrak{p} \mid (\mathfrak{m}, \ell)$ such that $v_{\mathfrak{p}}(\mathfrak{m}) > z(\mathfrak{p}, \ell)$, 227
SNF Smith normal form, 42
$\mathrm{St}(M)$ Steinitz class of M, 11

T

$T_2 = T_2(\alpha)$ $\sum_{1 \le j \le n} |\alpha_j|^2$, 451
t_2 upper bound for T_2 usually given by Hunter's theorem, 451
$T_k = T_k(\alpha)$ $\sum_{1 \le j \le n} |\alpha_j|^k$, 452
$T_{\mathfrak{m}}(L/K)$ Takagi (or norm) group for the modulus $\mathfrak{m}$, 153
tors torsion, 7
$\mathrm{Tr}_{L/K}(\alpha)$ relative trace of α, 76
T^{σ} conjugate by σ of polynomial T, 72

U

U_a part of unimodular matrix necessary for computing discrete logarithms in Abelian groups, 165
$U_i(K)$ an abbreviation for $U_{i,L}(K)$, 358
$U_i(L/K)$ group of relative pseudo-units, 353
$U_{i,L}(K)$ in fact, the trivial group, 358
$U(K)$ unit group of K, 136
U usually a unimodular transformation matrix, 17
$U_{\mathfrak{m}}(K)$ $:= U(K) \cap K^*_{\mathfrak{m}}$, 136
$U_{N,0}(L/K)$ subgroup of $U_N(L/K)$, 358
$U_N(K)$ abbreviation for $U_{N,L}(K)$, 358
$U_N(L/K)$ group of relative norm-units, 358
$U_{N,L}(K)$ $U(K)/(\mu(K) \cdot N_{L/K}(U(L)))$, 358
$U_S(K)$ group of S-units of K, 371

V

$v_{\eta}(\gamma)$ multiplier system for the modular form η, 317
$V_{\ell}(K)$ group of ℓ-virtual units of K, 231
$v_{\mathfrak{p}}(x)$ $\mathfrak{p}$-adic valuation of x, 3

W

$W(\chi)$ Artin root number of χ, 300
$w(L/K)$ $w(L)/w(K)$, 355
$\wp(z, L)$ Weierstrass $\wp$-function of a lattice L at z, 325
$w(z, L)$ Weber functions for ray class field computations, 333

X

$\lfloor x \rceil$ floor of $x + 1/2$, 33
$|x|_{\mathfrak{p}}$ $\mathfrak{p}$-adic norm of x, 3
$\binom{X}{Y}$ vertical concatenation of matrices or column vectors X and Y, 164

Z

$\zeta_{K,S}(s, \sigma)$ partial Dedekind zeta function, 298
ζ_n primitive nth root of unity, specifically $e^{2i\pi/n}$, 227
$\zeta(z, L)$ Weierstrass ζ-function of a lattice L at z, 326
$\mathbb{Z}_K$ the ring of integers of K, 2
$(\mathbb{Z}_K/\mathfrak{m})^*$ $:= (\mathbb{Z}_K/\mathfrak{m}_0)^* \times \mathbb{F}_2^{\mathfrak{m}_\infty}$, 135
$\mathbb{Z}_{K,S}$ ring of S-integers of K, 371
$z(\mathfrak{p}, \ell)$ $:= \ell e(\mathfrak{p}/\ell)/(\ell - 1) + 1$, 498

Index of Algorithms

A

Absolute defining polynomial (from relative), 63
Absolute to relative defining polynomial, 66
$ad - bc = 1$ algorithm, 25
Addition of ideals, 94
Artin map on Kummer extensions, 272
Artin root number $W(\chi)$, 308

B

Buchmann–Lenstra algorithm, 111

C

Capitulation group, 361
Chinese remainder algorithm for ideals, nonrecursive, 188
Class group and capitulation group, 361
Norm class group and norm-default quotient group, 361
S-class group, 373
$Cl_i(L/K)$ and $Cl_i(K) = Cl_{i,L}(K)$, 361
$Cl_N(L/K)$ and $Cl_N(K) = Cl_{N,L}(K)$, 361
Compositum of two number fields using $\theta_1\theta_2$, 60
Compositum of two number fields using $k\theta_1 + \theta_2$, 57
Conductor of a character, 219
Conductor of a congruence subgroup, 214
Conductor of an Abelian extension, 216
Congruence subgroups dividing a given one, 214
Congruence subgroups of index ℓ, 182
Coprime representative computation, 207
Cubic extensions list
– relative cyclic, 438
– relative noncyclic, 441
Is a cubic form the image of a cubic field (version 1), 408
Is a cubic form the image of a cubic field, 422

D

Dedekind η-function on $\mathcal{H}$, 317
Dedekind criterion, 106
Discrete logarithm in $(1+\mathfrak{p})/(1+\mathfrak{p}^k)$, 201
Discrete logarithm in $(\mathbb{Z}_K/\mathfrak{m})^*$, 208
Discrete logarithm in $(\mathbb{Z}_K/\mathfrak{p}^k)^*$, 202
Discrete logarithm in $Cl(K)/Cl(K)$ 236
Discrete logarithm in the ℓ-Selmer group, 255
Discrete logarithm in the unit group 254

E

Element down, 65
element norm, 76
Euclidean algorithm in Dedekind domains, 17
Exceptional set S_0 for Galois extensions, 382
Exceptional set S_0 for non-Galois extensions, 382
Extended Euclidean algorithm in Dedekind domains, 17
Extension Abelian or not, 217
Extension of an Abelian group by another, 170

F

Factorization of an ideal, 99

H

Hermite normal form algorithm in Dedekind domains, 30
- interleaved version, 40
- modular version, 39
- transformation matrix, 41

Hilbert class field
- compute $P(X) \in \mathbb{Z}(X)$, 312
- suitability of $P(X)$, 312

Hilbert class fields of imaginary quadratic fields, 321
HNF reduction of an element modulo an ideal, 32

I

Ideal addition, 94
Ideal factorization, 99
Ideal inversion, 98
Ideal inversion using different, 98
Ideal list up to n
- conductor at ℓ, 101
- general, 100
- squarefree, 100

Ideal multiplication, 95
ideal norm, 116
Ideal powering, 95
Ideal product, 95
Ideal reduction (relative), 366
Ideal up in absolute HNF, 117
Image of a subgroup by a group homomorphism, 172
Integral pseudo-basis (driver algorithm), 107
Interleaved modular HNF algorithm in Dedekind domains, 40
Intersection and sum of two subgroups, 176
Intersection of two $\mathbf{Z}_K$-modules, 36
Intersection of two subgroups as a subgroup of one of them, 177
Inverse image of a subgroup by a group homomorphism, 173
Inverse of a prime ideal, 94
Inverse of an ideal, 98
Inverse of an ideal using different, 98

K

Kernel of a group homomorphism, 173
Kummer extension of prime degree when $\zeta_\ell \in K$ using Hecke, 238
Kummer extension of prime degree when $\zeta_\ell \notin K$ using Hecke, 265
Kummer extension when $\zeta_n \in K$ using Artin, 278
Kummer extension when $\zeta_n \notin K$ using Artin, 284

L

Left four-term exact sequence: computation of the second group, 175
Linear system in integers, 182
Linear system of congruences, 184

Linear system of congruences and equations, 185
List of reduced binary quadratic forms, 323
List of relative cyclic cubic extensions, 438
List of relative noncyclic cubic extensions, 441
List of relative quadratic extensions using class field theory, 436
List of relative quadratic extensions using squarefree ideals, 434
LLL-reduction of an element modulo an ideal, 33

M

Mixed linear system, 185
Modular HNF algorithm in Dedekind domains, 39
Modular HNF with transformation matrix, 41
Multiplication of ideals, 95

N

Norm (or Takagi) group of an Abelian extension, 215
Norm equations, 383
- integral, 385
Norm group of an extension that is not necessarily Abelian, 216
norm of an element, 76
Norm of an ideal, 116
Norm-default quotient group, 361
Numerical value belongs to $\mathbb{Z}_K$
- for K imaginary quadratic, 343
- for K real quadratic, 309, 344

O

One-element representation in $(\mathbb{Z}_K/\mathfrak{m})^*$, 205
Over-order computation, 104

P

$\mathfrak{p}^a/\mathfrak{p}^b$, 200
$(1+\mathfrak{p}^a)/(1+\mathfrak{p}^b)$, 200
$\mathfrak{p}$-maximal order, 108
Polynomial reduction in the relative case, 110
Powering of an ideal, 95
$\mathfrak{p}$-radical computation, 102
Prime ideal below a prime ideal, 116
Prime ideal decomposition, 111
Prime ideal factorization of an ideal, 99
Primitive representatives of ray class group, 342
Principal ideal algorithm in ray class groups, 210
Product of ideals, 95
Product of two orders, 108
Pseudo–two-element representation of a prime ideal, 91
Pseudo–two-element representation of an ideal, 89

Q

Quadratic extension list using class field theory, 436
Quadratic extension list using squarefree ideals, 434
Quasi-period computation, 327
Quotient of groups, 168

R

Raising an ideal to a power, 95
Random element in an ideal, 23
Ray class field of imaginary quadratic field using $\sigma(z, L)$, 341
Ray class group associated to a modulus, 209
Reduction modulo $\mathfrak{p}$ of a pseudo-basis, 37
Reduction modulo $\mathfrak{p}$ of an element, 106
Reduction of a relative ideal, 366
Reduction of an element modulo an ideal using HNF, 32

Reduction of an element modulo an ideal using LLL, 33
Reduction of an ideal in a fixed ray class, 212
Reduction of polynomials in the relative case, 110
Reduction of the representative of a ray ideal class, 213
Relative Buchmann–Lenstra algorithm, 111
Relative class group and capitulation group, 361
Relative defining polynomial (from absolute), 66
Relative ideal reduction, 366
Relative ideal reduction (naive), 366
Relative integral norm equations, 385
Relative integral pseudo-basis (driver algorithm), 107
Relative norm equations, 383
relative norm of an element, 76
Relative norm of an ideal, 116
Relative polynomial reduction, 110
Relative prime ideal decomposition, 111
Relative round 2 algorithm (driver algorithm), 107
Relative round 2 algorithm at $\mathfrak{p}$, 108
Relative to absolute defining polynomial, 63
Relative unit group for $i_{L/K}$, 363
Relative unit group for $\mathcal{N}_{L/K}$, 364
Representation of an element coprime to an ideal, 207
Representative of an ideal class coprime to an ideal, 24
Reversion of an algebraic number, 66
Right four-term exact sequences: compute third group, 171
Round 2 algorithm (driver algorithm), 107
Round 2 algorithm at $\mathfrak{p}$, 108

S

Weierstrass σ-function computation 331
Simple relative polynomial reduction, 110
Smith normal form algorithm in Dedekind domains, 44
Smith normal form of an Abelian group given by generators and relations, 165
Solving ℓth power congruences, 505
Solving ℓth power congruences when $k \leq e(\mathfrak{p}/\ell)$, 507
Splitting class field extensions, 224
Splitting of a prime ideal in a class field, 304
Stark units for cyclic, real ray class fields over real quadratic fields, 310
Stark units for real ray class fields over real quadratic fields, 310
Subgroup reconstruction from its p-Sylow subgroups, 180
Subgroups of index ℓ, 182
Sum of two subgroups, 176

T

Table of complex cubic fields up to a given absolute discriminant, 424
Table of real cubic fields up to a given discriminant, 423
θ_1 and θ_2 computation, 60
Pseudo–two-element representation of an ideal, 89
Two-element representation of an ideal, 24
Pseudo–two-element representation of a prime ideal, 91

U

$U(L)/i_{L/K}(U(K))$, 362
$U(L)/(\mu(L) \cdot i_{L/K}(U(K))$, 363
Unit group for $i_{L/K}$, 363
Unit group for $N_{L/K}$, 364
S-unit group, 376

V

Valuation at a prime ideal for a relative quadratic extension, 125
Valuation of an ideal at a prime ideal, 93

Z

$(\mathbb{Z}_K/\mathfrak{m})^*$, 206
$(\mathbb{Z}_K/\mathfrak{p}^k)^*$, 201

General Index

A

Abelian extension
- character, 157
- conductor, 151, 155, 158
- conductor computation, 216
- norm group computation, 215
- prime decomposition, 154
- signature, 157

Abelian group
- cokernel, 174, 220
- effective computation, 166
- effective homomorphism, 166
- exact functor, 178
- exact sequence, 178
- extension, 169, 199
- finite, 164
- finitely generated, 164
- generators and relations, 164
- image, 172
- inverse image, 173
- kernel, 174
- left four-term exact sequence, 175
- p-Sylow subgroup, 177
- presentation, 165
- quotient, 168
- right four-term exact sequence, 171
- six-term exact sequence, 354
- Smith normal form, 165
- subgroup, 166

absolute
- defining polynomial, 49
- discriminant, 158
- equation, 49
- extension, 49
- polynomial, 49

absolute and relative discriminant, 114
addition of ideals, 95
additive structure of $\mathbf{Z}_K/\mathfrak{p}^k$, 193
algebra (bigraded), 392
algebraic K-theory, 354, 359
algebraic number
- characteristic polynomial, 55
- norm, 55
- reversion, 65
- trace, 55

ambiguous class, 380
Amice, Y., 190
approximate functional equation, 516
approximation theorem, 20
- strong, 4
- weak, 2

Artin group, 153
Artin map in Kummer theory, 227
Artin reciprocity law, 153, 154
Artin reciprocity map, 153
Artin root number, 300
Artin, E., 133, 150, 153
Artin–Hasse exponential, 202, 221
Artin–Hasse logarithm, 202, 221

B

Bachmann, G, 190
base field, 49, 50
Belabas, K., iv, v, 389, 418, 430

Bergé, A.-M., 445
bigraded algebra, 392
Bilhan, M., vi
binary form, 389
Birkhoff, G., 180
Bosma, W., 1
bounds (Odlyzko), 431
Buchmann, J., 83, 111, 445
Butler, L., 180

C

capitulating subgroup, 348
capitulation, 134, 348, 352
character
- conductor, 146
- conductor computation, 218
- even, odd, 300
- of a congruence subgroup, 145
- of an Abelian extension, 157
- primitive, 146, 300
characteristic polynomial, 55, 76
- computation, 76
- transitivity, 130
Chevalley, C., 380
Chinese remainder theorem
- for ideals, 187
- for moduli, 143
class field
- Hilbert, 133, 134, 155
- ray, 139, 155
- real ray, 301
class field theory, 133
class field tower, 134, 349
class group
- ray, 135, 347
- relative, 348
clique of exceptional units, 466
closed (integrally), 2
cocycle condition, 493
codifferent (relative), 96
Cohen, H., 430, 448
cokernel, 174, 220
complex cubic form (reduced), 419
complex multiplication, 223, 314
composition of pseudo-quadratic for 126
compositum, 49
- discriminant, 71
- of étale algebras, 65
- of number fields, 56
conductor, 145
- necessary conditions, 148, 149
- of a character, 146, 218
- of a class of congruence subgroups 144
- of a congruence subgroup, 155, 214
- of an Abelian extension, 151, 155, 158, 216
conductor at ℓ, 101
congruence subgroup, 138
- character, 145
- conductor, 155
- conductor computation, 214
congruence subgroups
- conductor, 144
- equivalence relation, 140
- GCD, 144
congruences (linear system of), 184
content of an ideal, 86
Conway, J., 445
coprime ideal class, 21
Cornell–Rosen theorem, 312
Couveignes, J.-M., v
covariant
- Jacobian, 393
- of a binary form, 391
Cremona, J., 392, 419, 427
cubic form, 395
- reduced, 411, 419
cubic polynomial (reduced), 426
cubic resolvent, 461

D

Davenport, H., 405
Davenport–Heilbronn theorem, 405
decomposition (of an ideal), 236
decomposition group, 151

Dedekind criterion, 106
Dedekind domain, 2
defining polynomial, 54
degree (of a gamma product), 509
Delaunay, C., 518
determinantal ideal, 39, 95
Diaz y Diaz, F., v, 131, 290, 430, 443, 445, 543
different (relative), 96
discrete logarithm, 166, 237
– short, 238
discriminant, 28, 78
– absolute, 158
– absolute and relative, 114
– elementary divisor, 80
– factorization, 71
– ideal, 78
– of a form, 390
– of an étale algebra, 55
– of compositum, 71
– relative, 79, 158
– root, 430
discriminant ideal, 28
– relative, 151
domain (Dedekind), 2
dual of an ideal, 97

E

element
– reduction, 211
– relative norm, 82
– torsion, 7
elementary divisor, 15, 42
– discriminantal, 80
– theorem, 16
elementary transformation, 20, 30
elliptic function, 325
embedding
– extension, 72
– ramified, 73
– relative, 73
– unramified, 73
equivalence
– Kummer, 498
– of congruence subgroups, 140
étale algebra, 50, 51
– compositum, 65
– discriminant, 55
– Galois group, 56
Euclidean algorithm, 17
– in Dedekind domains, 18
Euler ϕ-function for moduli, 137
even character, 300
exceptional unit, 466, 467
exponent of an Abelian group, 377
exponential
– $\mathfrak{p}$-adic, 190
– Artin–Hasse, 202, 221
extended multiplicative group, 352
extension
– absolute, 49
– Galois, 74
– normal, 74
– relative, iii, 49
– tamely ramified, 483
– unramified, 133
extension of an embedding, 72
extension of groups, 169

F

factor refinement, 23
factorization
– ideal, 99
– of discriminant, 71
Fieker, C., 227, 270, 294
field
– global, 17
– Hilbert class, 155
– ray class, 155
field norm, 3
$\mathfrak{f}_\infty$-positive, 307
finite field polynomial factorization, 110
finite type (function of), 511
Ford, D., 445
form
– binary, 389
– covariant, 391

- cubic, 395
- discriminant, 390
- integral, 394
- invariant, 391
- irreducible, 394
- primitive, 394
- pseudo-quadratic, 122
- root of, 389
four-term exact sequence
- left, 175
- right, 171
fractional ideal, 2
freeness test, 10, 35
Friedman, E., 301, 386, 508
Frobenius homomorphism, 152, 478
function
- elliptic, 325
- of finite type, 511
- theta, 329
functional equation, 511
- approximate, 516

G

Galois
- extension, 74
- group, 74
Galois group
- of an étale algebra, 56
Galois representation, 445
gamma product, 508
Gauss sum, 307
Gauss–Bareiss, 39
GCD of congruence subgroups, 144
generalized Smith normal form, 355
generators and relations, 164
global field, 17
Golod, E., 134, 349
Gras, G., 133, 150
GRH, 431
group
- Artin, 153
- cokernel, 174, 220
- congruence, 138
- decomposition, 151
- extension, 169
- ideal, 138
- image, 172
- inertia, 151
- inverse image, 173
- kernel, 174
- left four-term exact sequence, 175
- norm, 153, 215
- quotient, 168
- ramification, 439
- ray, 136
- ray class, 135, 209, 347
- right four-term exact sequence, 171
- Selmer, 231
- subgroup, 166
- subgroups, 179
- Takagi, 153

H

Hasse, H., 133, 150, 157, 276
Havas, G., 38
Hecke L-function, 299
Hecke's theorem, 227, 498
Hecke, E., 299, 498
Heilbronn, H., 405
Hensel lift, 189, 500
HNF representation, 84
Hermite normal form in Dedekind domains, 30
Hermite's constant, 445
Hessian, 400
Hilbert class field, 133, 134, 155
Hilbert's Theorem 90, 494
Hilbert, D., 133, 494
HNF algorithm in Dedekind domains, 30
- modular, 39, 40
homomorphism
- Frobenius, 152, 478
Hoppe, A., 95
Hunter's theorem, 445

I

ideal
- Chinese remainder theorem, 187
- content, 86
- decomposition modulo $Cl(K)^\ell$, 236
- determinantal, 39
- discriminant, 28
- dual, 97
- Euler ϕ, 137
- factorization, 99
- fractional, 2
- index, 15
- inverse, 97
- minor, 31
- order, 15
- powering, 96
- primitive, 86, 319
- product, 95
- pseudo-principal, 348
- reduced, 365
- reduction, 211, 212
- relative norm, 80, 83
- relative reduction, 366
- representation, 83
- S-integral, 372
- sum, 95

ideal group, 138
image
- of a subgroup, 172
- pseudo-matrix, 29

index of a suborder, 79
index-ideal, 15, 79
inertia group, 151
inessential discriminantal divisor, 396
integral binary form, 394
integral pseudo-basis, 78
integral pseudo-matrix, 43
integrally closed, 2
invariant factor, 15
invariant of a binary form, 391
inverse image of a subgroup, 173
inverse Mellin transform, 511
inverse of a pseudo-quadratic form, 126
inverse of an ideal, 97
inverter, 530
irreducible binary form, 394

J

Jacobi quintuple-product identity, 345
Jacobi triple-product identity, 329
Jacobian covariant, 393
Janusz, G., 133, 150
Julia-reduced, 427

K

Kant/Kash, iii, 430, 525
kernel of a group map, 174
Klüners, J., 313
Koblitz, N., 190
Kummer theory, 223, 297, 545
- Artin map, 227
- using Artin when $\zeta_n \in K$, 270
- using Artin when $\zeta_n \notin K$, 280
- using Hecke when $\zeta_\ell \in K$, 227
- using Hecke when $\zeta_\ell \notin K$, 242

Kummer-equivalence, 498

L

Lagrange multiplier, 455
Lagrange resolvent, 227, 248
Langlands, R. P., 133
Lavrik, A. F., 508
left divisor of a matrix, 167
Lenstra, H. W., 83, 111, 466, 467
Leutbecher, A., 467
LiDIA, iii, 525
lift (Hensel), 189, 500
linear system
- of congruences, 184
- in integers, 182
- mixed, 184

LLL (partial), 34, 127
local norm, 151
logarithm

- $\mathfrak{p}$-adic, 190
- Artin–Hasse, 202, 221
- discrete, 166, 237
- short discrete, 238

M

Magma, iii, 524
Martinet, J., v, 2, 133, 349, 433, 439, 445, 448
matrix (left divisor), 167
maximal order, 17
Mellin inversion formula, 512
Mellin transform, 511
Minkowski, H., 430
minor-ideal, 31, 45
mixed linear system, 184
modular HNF algorithm, 39, 40
module
- finitely generated, 6
- projective, 8, 46
- pseudo-matrix, 29
- rank, 6

modulus, 135
- Chinese remainder theorem, 143
- Euler ϕ, 137
- suitable, 152

Montgomery, P., 34, 127
multiplication of ideals, 95
multiplicative group (extended), 352
multiplicative structure of $(\mathbb{Z}_K/\mathfrak{p}^k)^*$, 194, 196, 201
multiplier (Lagrange), 455
multiplier system, 317

N

Nakagawa, J., 418
Nakayama's lemma, 475
Neukirch, J., 133, 150
Newton
- iteration, 189, 500
- power sum recursion, 252

Newton's formulas, 447, 451
Newton's inequalities, 454
Niklasch, G., 467
Noether's theorem, 493
Noether, E., 493
Noetherian ring, 2
norm
- Archimedean, 3
- field, 3
- local, 151
- non-Archimedean, 3
- relative, 76, 153

norm group, 153
- of an Abelian extension, 215

norm-class group, 349
norm-default quotient, 349
norm-unit, 358
normal extension, 74
number field
- compositum, 56
- primitive, 446

O

odd character, 300
Odlyzko bounds, 431
Odlyzko, A., 430, 543
Olivier, M., v, 445
one-element representation, 205
order
- maximal, 17
- p-maximal, 102
- product, 108

order-ideal, 15
orthogonal idempotents, 55
Ostrowsky's theorem, 3

P

$\mathfrak{p}$-adic completion, 190
$\mathfrak{p}$-adic exponential, 190
$\mathfrak{p}$-adic integer, 190
$\mathfrak{p}$-adic logarithm, 190
$\mathfrak{p}$-maximal order, 102
$\mathfrak{p}$-radical, 102
pairing (perfect), 496
Pari, iii
Pari/GP, iii, 430, 525
perfect pairing, 496

$\wp$-function of Weierstrass, 325
Phragmén–Lindelöf theorem, 513
place of a number field, 3
Plouffe's inverter, 530
Plouffe, S., 530
Pohst, M., 1, 430, 445, 455
Poitou, G., 430, 543
polynomial
- absolute, 49
- defining, 49
- relative, 49
polynomial factorization in a finite field, 110
polynomial reduction, 290
Poonen, B., 76
$\mathfrak{f}_\infty$-positive, 307
power of an ideal, 96
power sum recursion, 252
presentation, 165
prime ideal
- decomposition, 111
- down, 116
- representation, 89
- uniformizer, 104
primitive
- character, 146, 300
- element, 54
- element theorem, 52, 54
- field, 446
- form, 394
- ideal, 319
- relative ideal, 86
principal ideal algorithm, 209
- in ray class groups, 210
product
- of ideals, 95
- of orders, 108
product formula, 3
projective module, 8, 46
pseudo–two-element representation, 87
pseudo-basis, 26
- integral, 78
- reduction modulo $\mathfrak{p}$, 37
pseudo-class group, 348
pseudo-class number, 348
pseudo-element, 26, 87, 348
pseudo-generating set, 26
pseudo-matrix, 28, 29
- image, 29
- integral, 43
- module, 29
pseudo-principal ideal, 348
pseudo-quadratic form, 122
- composition, 126
- inverse, 126
- reduction, 127
pseudo-unit, 353

Q

quadratic form (pseudo), 122
quasi-period, 326
quintuple-product identity, 345
quotient (norm-default), 349
quotient of two Abelian groups, 168

R

radicand, 118
ramification (tame), 483
ramification group, 439
ramification index, 82
ramified embedding, 73
ray class field, 139, 155
- real, 301
ray class group, 135, 136, 163, 172, 209, 347
ray group, 136
real cubic form (reduced), 411
real ray class field, 301
reciprocity law, 154
- Artin's, 153
- Shimura's, 315
reduced
- complex cubic form, 419
- cubic polynomial, 426
- Julia, 427
- real cubic form, 411

- swap, 128
- translation, 127
reduction
- ideal in a ray class, 212
- modulo $\mathfrak{p}$ of a pseudo-basis, 37
- modulo an ideal in HNF, 32
- modulo an ideal in LLL, 33
- of a pseudo-quadratic form, 127
- of an ideal, 366
- of elements, 211
- of ideals, 211
- partial LLL, 34, 127
- polynomial, 290
- relative ideal, 365
- relative polynomial, 110
regulator
- relative, for $i_{L/K}$, 358
- relative, for $\mathcal{N}_{L/K}$, 359
relative
- characteristic polynomial, 76
- class group, 348
- codifferent, 96
- defining polynomial, 49
- degree, 50
- different, 96
- discriminant, 79, 158
- discriminant ideal, 151
- element norm, 82
- embedding, 73
- equation, 49
- extension, iii, 49
- ideal norm, 80, 83
- ideal reduction, 366
- integral pseudo-basis, 78, 102
- norm, 76, 153
- norm of an element, 76
- norm of an ideal, 80
- norm-class group, 349
- norm-class number, 349
- norm-unit, 358
- polynomial, 49
- polynomial reduction, 110
- prime ideal decomposition, 111
- primitive ideal, 86
- pseudo-class group, 348
- pseudo-unit, 353
- regulator for $i_{L/K}$, 358
- regulator for $\mathcal{N}_{L/K}$, 359
- round 2 algorithm, 107
- trace, 76
- valuation, 124
representation
- Galois, 445
- HNF, 84
- of a group, 165
- of a prime ideal, 89
- of a subgroup, 166
- of an ideal, 83
- of elements of $(\mathbb{Z}_K/\mathfrak{m})^*$, 205
- of elements of $\mathbb{Z}_K/\mathfrak{p}$, 105
- one-element, 205
- pseudo–two-element, 87, 89, 91
- two-element, 95
residual degree, 82
resolvent
- cubic, 461
- Lagrange, 227, 248
reversion of an algebraic number, 65
ring (Noetherian), 2
Roblot, X., v, 298
root discriminant, 430, 545
root number, 300
round 2 algorithm, 107

S

Schertz, R., 314, 320
Schwarz, A., 445
section, 8
Selmer group, 231
separable algebra, 51
sequences, 530
Serre, J.-P., 430, 468, 543
Shafarevitch, I., 134, 349
Shimura, G., 315
Shintani, T., 418
short discrete logarithm, 238
Siegel, C.-L., 454

σ-function of Weierstrass, 328, 334
signature homomorphism, 206
signature of an Abelian extension, 157
Simon, D., iv, v, 130, 347, 352, 360, 377, 466
S-integral ideal, 372
Sloane, N., 445, 530
Smith normal form
- generalized, 355
- in Dedekind domains, 44
- of an Abelian group, 165
SNF, 42, 165
Smyth, C., 454
splitting field, 312
Stark
- conjecture, 297
- units, 297
Stark, H., 223, 297, 430, 543
Steinitz class, 6, 11
- additivity, 12
strong approximation theorem, 4
structure theorem
- for finitely generated modules, 13
- for projective modules, 11
- for torsion modules, 13
- for torsion-free modules, 8
subgroup
- congruence, 138
- enumeration, 179
- image, 172
- inverse image, 173
- kernel, 174
subresultant algorithm, 52, 60
suitable modulus, 152
suitable set S_0 for L/K, 378
sum of ideals, 95
superseeker, 530
swap-reduced quadratic form, 128
Sylvester matrix, 51, 63
syzygy, 394

T

Takagi existence theorem, 154
Takagi group, 153
tamely ramified extension, 483
Tate, J., 133, 150, 298
theta function, 329
Tollis, E., v
torsion
- element, 7
- module, 7
- submodule, 7
torsion-free, 7
tower (class field), 349
trace, 55
- relative, 76
transitivity
- of degrees, 50
- of relative discriminants, 114
- of the characteristic polynomial, 130
- of the different, 97
- of the ideal norm, 80
- of the trace and norm, 76
translation-reduced quadratic form 127
triple-product identity, 329
two-element representation, 21, 87, 95

U

uniformizer (of an ideal), 5
uniformizer of a prime ideal, 104
unit
- exceptional, 467
- virtual, 231
unramified embedding, 73
unramified extension, 133

V

valuation
- at a prime ideal, 92
- relative, 124
virtual unit, 231

W

weak approximation theorem, 2
Weierstrass σ-function, 328, 334
Weierstrass ζ-function, 326
Weierstrass $\wp$-function, 325
Wright, D., 450

Y

Yukie, A., 450

Z

Zassenhaus, H., 104
ζ-function of Weierstrass, 326
$\mathbb{Z}_K/\mathfrak{p}^k$ (additive structure), 193
$(\mathbb{Z}_K/\mathfrak{p}^k)^*$ (multiplicative structure), 194, 196, 201

Graduate Texts in Mathematics

(continued from page ii)

62 KARGAPOLOV/MERLZJAKOV. Fundamentals of the Theory of Groups.
63 BOLLOBAS. Graph Theory.
64 EDWARDS. Fourier Series. Vol. I 2nd ed.
65 WELLS. Differential Analysis on Complex Manifolds. 2nd ed.
66 WATERHOUSE. Introduction to Affine Group Schemes.
67 SERRE. Local Fields.
68 WEIDMANN. Linear Operators in Hilbert Spaces.
69 LANG. Cyclotomic Fields II.
70 MASSEY. Singular Homology Theory.
71 FARKAS/KRA. Riemann Surfaces. 2nd ed.
72 STILLWELL. Classical Topology and Combinatorial Group Theory. 2nd ed.
73 HUNGERFORD. Algebra.
74 DAVENPORT. Multiplicative Number Theory. 2nd ed.
75 HOCHSCHILD. Basic Theory of Algebraic Groups and Lie Algebras.
76 IITAKA. Algebraic Geometry.
77 HECKE. Lectures on the Theory of Algebraic Numbers.
78 BURRIS/SANKAPPANAVAR. A Course in Universal Algebra.
79 WALTERS. An Introduction to Ergodic Theory.
80 ROBINSON. A Course in the Theory of Groups. 2nd ed.
81 FORSTER. Lectures on Riemann Surfaces.
82 BOTT/TU. Differential Forms in Algebraic Topology.
83 WASHINGTON. Introduction to Cyclotomic Fields. 2nd ed.
84 IRELAND/ROSEN. A Classical Introduction to Modern Number Theory. 2nd ed.
85 EDWARDS. Fourier Series. Vol. II. 2nd ed.
86 VAN LINT. Introduction to Coding Theory. 2nd ed.
87 BROWN. Cohomology of Groups.
88 PIERCE. Associative Algebras.
89 LANG. Introduction to Algebraic and Abelian Functions. 2nd ed.
90 BRØNDSTED. An Introduction to Convex Polytopes.
91 BEARDON. On the Geometry of Discrete Groups.
92 DIESTEL. Sequences and Series in Banach Spaces.
93 DUBROVIN/FOMENKO/NOVIKOV. Modern Geometry—Methods and Applications. Part I. 2nd ed.
94 WARNER. Foundations of Differentiable Manifolds and Lie Groups.
95 SHIRYAEV. Probability. 2nd ed.
96 CONWAY. A Course in Functional Analysis. 2nd ed.
97 KOBLITZ. Introduction to Elliptic Curves and Modular Forms. 2nd ed.
98 BRÖCKER/TOM DIECK. Representations of Compact Lie Groups.
99 GROVE/BENSON. Finite Reflection Groups. 2nd ed.
100 BERG/CHRISTENSEN/RESSEL. Harmonic Analysis on Semigroups: Theory of Positive Definite and Related Functions.
101 EDWARDS. Galois Theory.
102 VARADARAJAN. Lie Groups, Lie Algebras and Their Representations.
103 LANG. Complex Analysis. 3rd ed.
104 DUBROVIN/FOMENKO/NOVIKOV. Modern Geometry—Methods and Applications. Part II.
105 LANG. $SL_2(\mathbf{R})$.
106 SILVERMAN. The Arithmetic of Elliptic Curves.
107 OLVER. Applications of Lie Groups to Differential Equations. 2nd ed.
108 RANGE. Holomorphic Functions and Integral Representations in Several Complex Variables.
109 LEHTO. Univalent Functions and Teichmüller Spaces.
110 LANG. Algebraic Number Theory.
111 HUSEMÖLLER. Elliptic Curves.
112 LANG. Elliptic Functions.
113 KARATZAS/SHREVE. Brownian Motion and Stochastic Calculus. 2nd ed.
114 KOBLITZ. A Course in Number Theory and Cryptography. 2nd ed.
115 BERGER/GOSTIAUX. Differential Geometry: Manifolds, Curves, and Surfaces.
116 KELLEY/SRINIVASAN. Measure and Integral. Vol. I.
117 SERRE. Algebraic Groups and Class Fields.
118 PEDERSEN. Analysis Now.
119 ROTMAN. An Introduction to Algebraic Topology.

120 ZIEMER. Weakly Differentiable Functions: Sobolev Spaces and Functions of Bounded Variation.
121 LANG. Cyclotomic Fields I and II. Combined 2nd ed.
122 REMMERT. Theory of Complex Functions. *Readings in Mathematics*
123 EBBINGHAUS/HERMES et al. Numbers. *Readings in Mathematics*
124 DUBROVIN/FOMENKO/NOVIKOV. Modern Geometry—Methods and Applications. Part III.
125 BERENSTEIN/GAY. Complex Variables: An Introduction.
126 BOREL. Linear Algebraic Groups. 2nd ed.
127 MASSEY. A Basic Course in Algebraic Topology.
128 RAUCH. Partial Differential Equations.
129 FULTON/HARRIS. Representation Theory: A First Course. *Readings in Mathematics*
130 DODSON/POSTON. Tensor Geometry.
131 LAM. A First Course in Noncommutative Rings.
132 BEARDON. Iteration of Rational Functions.
133 HARRIS. Algebraic Geometry: A First Course.
134 ROMAN. Coding and Information Theory.
135 ROMAN. Advanced Linear Algebra.
136 ADKINS/WEINTRAUB. Algebra: An Approach via Module Theory.
137 AXLER/BOURDON/RAMEY. Harmonic Function Theory.
138 COHEN. A Course in Computational Algebraic Number Theory.
139 BREDON. Topology and Geometry.
140 AUBIN. Optima and Equilibria. An Introduction to Nonlinear Analysis.
141 BECKER/WEISPFENNING/KREDEL. Gröbner Bases. A Computational Approach to Commutative Algebra.
142 LANG. Real and Functional Analysis. 3rd ed.
143 DOOB. Measure Theory.
144 DENNIS/FARB. Noncommutative Algebra.
145 VICK. Homology Theory. An Introduction to Algebraic Topology. 2nd ed.
146 BRIDGES. Computability: A Mathematical Sketchbook.
147 ROSENBERG. Algebraic *K*-Theory and Its Applications.
148 ROTMAN. An Introduction to the Theory of Groups. 4th ed.
149 RATCLIFFE. Foundations of Hyperbolic Manifolds.
150 EISENBUD. Commutative Algebra with a View Toward Algebraic Geometry.
151 SILVERMAN. Advanced Topics in the Arithmetic of Elliptic Curves.
152 ZIEGLER. Lectures on Polytopes.
153 FULTON. Algebraic Topology: A First Course.
154 BROWN/PEARCY. An Introduction to Analysis.
155 KASSEL. Quantum Groups.
156 KECHRIS. Classical Descriptive Set Theory.
157 MALLIAVIN. Integration and Probability.
158 ROMAN. Field Theory.
159 CONWAY. Functions of One Complex Variable II.
160 LANG. Differential and Riemannian Manifolds.
161 BORWEIN/ERDÉLYI. Polynomials and Polynomial Inequalities.
162 ALPERIN/BELL. Groups and Representations.
163 DIXON/MORTIMER. Permutation Groups.
164 NATHANSON. Additive Number Theory: The Classical Bases.
165 NATHANSON. Additive Number Theory: Inverse Problems and the Geometry of Sumsets.
166 SHARPE. Differential Geometry: Cartan's Generalization of Klein's Erlangen Program.
167 MORANDI. Field and Galois Theory.
168 EWALD. Combinatorial Convexity and Algebraic Geometry.
169 BHATIA. Matrix Analysis.
170 BREDON. Sheaf Theory. 2nd ed.
171 PETERSEN. Riemannian Geometry.
172 REMMERT. Classical Topics in Complex Function Theory.
173 DIESTEL. Graph Theory.
174 BRIDGES. Foundations of Real and Abstract Analysis.
175 LICKORISH. An Introduction to Knot Theory.
176 LEE. Riemannian Manifolds.
177 NEWMAN. Analytic Number Theory.
178 CLARKE/LEDYAEV/STERN/WOLENSKI. Nonsmooth Analysis and Control Theory.
179 DOUGLAS. Banach Algebra Techniques in Operator Theory. 2nd ed.

180 SRIVASTAVA. A Course on Borel Sets.
181 KRESS. Numerical Analysis.
182 WALTER. Ordinary Differential Equations.
183 MEGGINSON. An Introduction to Banach Space Theory.
184 BOLLOBAS. Modern Graph Theory.
185 COX/LITTLE/O'SHEA. Using Algebraic Geometry.
186 RAMAKRISHNAN/VALENZA. Fourier Analysis on Number Fields.
187 HARRIS/MORRISON. Moduli of Curves.
188 GOLDBLATT. Lectures on the Hyperreals: An Introduction to Nonstandard Analysis.
189 LAM. Lectures on Modules and Rings.
190 ESMONDE/MURTY. Problems in Algebraic Number Theory.
191 LANG. Fundamentals of Differential Geometry.
192 HIRSCH/LACOMBE. Elements of Functional Analysis.
193 COHEN. Advanced Topics in Computational Number Theory.
194 ENGEL/NAGEL. One-Parameter Semigroups for Linear Evolution Equations.

This book addresses a number of specific topics in computational number theory centered on class field theory and relative extensions of number fields. Most of the material is new from the algorithmic standpoint. The book is organized as follows. Chapters 1 and 2 contain the theory and algorithms concerning Dedekind domains and relative extensions of number fields, and in particular the generalization to the relative case of the round 2 and related algorithms. Chapters 3, 4, 5, and 6 contain the theory and complete algorithms concerning class field theory over number fields. The highlights are the algorithms for computing the structure of $(\mathbb{Z}_K/\mathfrak{m})^*$, of ray class groups, and relative equations for Abelian extensions based on complex multiplication or Stark's conjectures. Together with Chapter 10, which contains complete proofs of several results used in the rest of the book that cannot easily be found in the existing literature, Chapters 1 to 6 form a homogeneous subject matter, which can be used for a 6-month to 1-year graduate course in computational number theory. The other chapters deal with more miscellaneous subjects. Written by an authority with great practical and teaching experience in the field, this book together with the author's earlier book, *A Course in Computational Algebraic Number Theory* (GTM 138), will become the standard and indispensable reference on the subject.

ISBN 0-387-**98727**-4
www.springer-ny.com

ISBN 0-387-98727-4
EAN
9 780387 987279 >