



Universiteit
Leiden
The Netherlands

Constructing elliptic curves of prescribed order

Bröker, R.

Citation

Bröker, R. (2006, June 27). *Constructing elliptic curves of prescribed order*. Retrieved from <https://hdl.handle.net/1887/4425>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4425>

Note: To cite this publication please use the final published version (if applicable).

CONSTRUCTING ELLIPTIC CURVES OF PRESCRIBED ORDER

PROEFSCHRIFT

ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van de Rector Magnificus Dr. D. D. Breimer,
hoogleraar in de faculteit der Wiskunde en
Natuurwetenschappen en die der Geneeskunde,
volgens besluit van het College voor Promoties
te verdedigen op dinsdag 27 juni 2006
klokke 14:15 uur

door

REINIER MARTIJN BRÖKER

geboren te Geldermalsen
in 1979

Samenstelling van de promotiecommissie:

promotor: Prof. dr. P. Stevenhagen
referent: Prof. dr. R. Schoof (*Università di Roma Tor Vergata*)

overige leden: Prof. dr. S. J. Edixhoven
Dr. A. Enge (*École Polytechnique*)
Prof. dr. H. W. Lenstra, Jr.
Prof. dr. J. van Mill (*Vrije Universiteit*)
Prof. dr. S. M. Verduyn Lunel

Constructing elliptic curves of prescribed order

Bröker, Reinier, 1979 –
Constructing elliptic curves of prescribed order
AMS 2000 Subj. class. code: 14H52, 11G15
NUR: 921
ISBN-10: 90-9020447-4
ISBN-13: 978-90-9020447-5

THOMAS STIELTJES INSTITUTE
FOR MATHEMATICS



© R. Bröker, Leiden 2006.

The illustration on page 164 is due to Reid, Geleijnse and Van Tol and is used with their permission.

All rights reserved. No part of this publication may be reproduced in any form or by any electronic or mechanical means including information storage and retrieval systems without the prior written permission from the author.

Table of contents

1	Introduction	1
1	Background	1
2	Elliptic curves of prescribed order	3
3	Complex multiplication constructions	5
4	Class invariants	7
2	Elliptic curves of given order	11
1	Elliptic curves over finite fields	11
2	Does there exist a curve with exactly N points?	14
3	Naïve algorithm	18
4	Analysis	19
5	Timings and examples	22
3	Complex multiplication	25
1	Deuring lifting	25
2	Complex multiplication constructions	26
3	Complex analytic methods	31
4	Constructing supersingular elliptic curves	34

4	An efficient algorithm	37
1	Finding a small discriminant	37
2	An algorithm to solve problem 4.1	41
3	Heuristic run time analysis	44
4	Examples and practical considerations	52
5	A non-archimedean algorithm	57
1	Finding a small splitting prime	57
2	The canonical lift	59
3	Modular curves	63
4	Computing the canonical lift	65
5	Isogenous curves with isomorphic endomorphism rings	72
6	Computing the kernel polynomial	75
7	Algorithm for computing the canonical lift	83
8	Computing the Hilbert class polynomial	86
9	Example	88
6	Class invariants	91
1	Introduction	91
2	The modular function field	93
3	Class invariants	95
4	Shimura reciprocity over the ring class field	97
5	Shimura reciprocity	103
6	Class invariants in a non-archimedean setting	107
7	Finding a class invariant	113
8	Using modular polynomials	117
9	Further improvements	123

7	Examples	127
1	A cryptographic curve	127
2	Large group orders	137
3	Simple η -quotients	139
4	Double η -quotients	147
5	A large discriminant	152
	 References	 155
	 Samenvatting	 161
	 Curriculum Vitae	 167

1

Introduction


1.1 Background

This thesis deals with elliptic curves, and more specifically with some of their algorithmic aspects. In algorithmic practice, an elliptic curve E over a field K is often described by a Weierstraß equation, i.e., a specific model for the curve in the projective plane $\mathbf{P}_K^2$ over K . For $\text{char}(K) \neq 2, 3$ this model takes the simple form

$$Y^2Z = X^3 + aXZ^2 + bZ^3$$

with coefficients $a, b \in K$. The set $E(K)$ of K -rational points consists of the solutions $(x : y : z) \in \mathbf{P}^2(K)$ to this equation. It contains the point at infinity $O = (0 : 1 : 0)$. All other points lie in the affine plane $Z \neq 0$ and we usually give the affine equation $Y^2 = X^3 + aX + b$ for the curve, the point O at infinity being understood.

One of the key ingredients of most algorithms employing elliptic curves, is that for an elliptic curve E defined over a field K , the set $E(K)$ of K -rational points has a natural group operation for which O is the neutral element. If $K = \mathbf{R}$ is the field of real numbers, we can easily visualise the group law.



The picture shows the definition of the sum $P + Q$ for two points $P, Q \in E(\mathbf{R})$. Proving that this definition indeed turns $E(\mathbf{R})$ into an abelian group is not so easy

2 Introduction

in fact; the hard part is to show that the addition is associative. The usual proof proceeds via algebraic geometric lines, and does not use the Weierstraß equation.

In the 1980's, elliptic curves gained importance in algorithmic number theory. Given an elliptic curve E over a finite field $\mathbf{F}_q$, Hasse's theorem from 1933 states that the number $\#E(\mathbf{F}_q)$ of $\mathbf{F}_q$ -rational points of E is an element of the Hasse interval

$$\mathcal{H}_q = [q + 1 - 2\sqrt{q}, q + 1 + 2\sqrt{q}] \quad (1.1)$$

around $q + 1$. In 1985, Schoof [51] gave a deterministic polynomial time algorithm to compute $\#E(\mathbf{F}_q)$ from a standard representation of E by a Weierstraß model

$$Y^2 + a_1XY + a_3Y = X^3 + a_2X^2 + a_4X + a_6$$

over $\mathbf{F}_q$. This algorithm has subsequently been improved by Elkies and Atkin [50, 19, 44], and point counting is nowadays considered 'easy'.

In the same paper, Schoof gives a deterministic polynomial time algorithm to compute a square root of the reduction $\bar{x} \in \mathbf{F}_p$ of a *fixed* integer $x \in \mathbf{Z}$. This algorithm, which is an application of his point counting algorithm, is currently the only known *deterministic* polynomial time algorithm to compute modular square roots of a fixed integer $x \in \mathbf{Z}$. As this algorithm is quite impractical, its importance is mostly theoretical.

In 1987, Lenstra published a factoring algorithm based on elliptic curves [39]. Here one works with elliptic curves over the *ring* $\mathbf{Z}/N\mathbf{Z}$, where N is the integer we want to factor. This probabilistic algorithm is an extension of Pollard's $(p - 1)$ -method. Pollard's algorithm is efficient if $p - 1$ is smooth, i.e., not divisible by 'large' primes. The elliptic curve factoring algorithm is efficient for a curve E over $\mathbf{Z}/N\mathbf{Z}$ if the group order $\#E(\mathbf{Z}/N\mathbf{Z})$ is smooth. We have *many* elliptic curves E over $\mathbf{Z}/N\mathbf{Z}$ to choose from, and it is this flexibility that is crucial to the performance of the algorithm. In practice, it is the fastest known algorithm to find prime factors up to say 50 decimal digits of N .

The improved flexibility also plays a vital role in the elliptic curve primality test, proposed by Goldwasser and Killian in 1986. See [13, Section 14D] for a description of this algorithm. After the practical improvements made by Atkin in 1988, this algorithm is currently one of the fastest known algorithms to rigorously prove primality. A rigorous run time analysis is out of reach. The deterministic primality test by Agrawal, Kayal and Saxena [2] runs in time that is polynomially bounded in the input size $\log N$. This test is not fast in practice however.

1.2 Elliptic curves of prescribed order

The problem we consider in this thesis arises as a natural ‘inverse problem’ to the point counting problem considered by Schoof.

PROBLEM 1. *Given a finite field $\mathbf{F}_q$ and an integer $N \in \mathcal{H}_q$, find an elliptic curve $E/\mathbf{F}_q$ for which $E(\mathbf{F}_q)$ has order N .*

If $q = p$ is a prime number, every integer $N \in \mathcal{H}_p$ arises as group order of an elliptic curve $E/\mathbf{F}_p$. We prove this in theorem 2.5. For arbitrary prime powers $q = p^f$ this is not generally true: there are often not enough supersingular curves to cover the cases $N \equiv 1 \pmod p$.

There is no algorithm known that solves problem 1 (in the cases that a solution exists) in a time that is polynomially bounded in the input size $\log q \approx \log N$. The fastest algorithm known has an expected run time $\tilde{O}(N^{1/2})$, where the $\tilde{O}$ -notation indicates that we disregard logarithmic factors. This *naïve algorithm*, which simply tries random curves over $\mathbf{F}_q$ until we hit a curve with N points, is stated and analysed in chapter 2.

We can relax the conditions in problem 1 by considering the prime power q as part of the *output* instead of the *input*.

PROBLEM 2. *Given an integer $N \in \mathbf{Z}_{\geq 1}$, find a finite field $\mathbf{F}_q$ and an elliptic curve $E/\mathbf{F}_q$ for which $E(\mathbf{F}_q)$ has order N .*

This problem forms the core of the first half of this thesis. It is not only inspired by simply relaxing the conditions of problem 1, but also by cryptographic applications.

If the order of $E(\mathbf{F}_q)$ is a large prime, the discrete log problem in $E(\mathbf{F}_q)$ is considered to be hard. That is, if we are given $P, Q \in E(\mathbf{F}_q)$ it is hard to find an integer $k \in \mathbf{Z}$ with $kP = Q$. There are a few technical conditions to exclude ‘weak curves’ such as supersingular curves; we do not go into this here.

Curves for which the discrete log problem is hard can be used for a cryptographic system. A simple algorithm to find a ‘strong curve’ is to select a prime p of 60 digits and try random elliptic curves over $\mathbf{F}_p$ until we hit a curve of prime order N . Hasse’s theorem ensures that N is of the same size as p . This resembles the naïve algorithm to construct a curve with N points. Using this algorithm to construct a curve of prime order is heuristically polynomial time however. Indeed, by the prime number theorem we expect that one out of every $\log p$ integers of size

p is prime. If we treat the group orders of the curves we try as random integers of size $\log p$, we expect that we have to try about $\log p$ curves until we hit a curve of prime order. Since point counting is polynomial time, this yields a polynomial time algorithm.

For the hardness of the discrete log problem in $E(\mathbf{F})$, it is not so relevant over which finite field the curve E is defined. All we require is that $E(\mathbf{F})$ has prime order. It therefore suffices to prescribe the prime order N and ask for a curve with N points. Hence, for cryptographic purposes we are mostly interested in a solution to problem 2.

One of the main results of this thesis is that there does exist an efficient solution to problem 2 if N is provided to the algorithm in factored form. For practical applications, such as those in elliptic curve cryptography, it is unlikely that one will need or want to use elliptic curves for which the factorization of the group order is unknown, so requiring the factorization of N to be part of our input is not a severe restriction. Our solution to problem 2 for factored orders N is almost polynomial time, provided that one is willing to make a number of ‘standard heuristic assumptions’.

THEOREM 1.3. *There exists an algorithm that, on input of an integer $N \geq 1$ together with its factorization, returns a prime number p and an elliptic curve $E/\mathbf{F}_p$ with $\#E(\mathbf{F}_p) = N$ whenever such a pair (E, p) exists. Under standard heuristic assumptions, a pair (E, p) exists for all N , and the expected run time of the algorithm is polynomial in $2^{\omega(N)} \log N$. Here $\omega(N)$ denotes the number of distinct prime factors of N .*

The explicit description of this algorithm and the run time analysis are given in chapter 4. Although the run time in theorem 1.3 is not polynomial in the usual sense, it is polynomial in $\log N$ outside a zero density subset of $\mathbf{Z}_{\geq 1}$ consisting of very smooth input values N . Note that such N are not used in cryptographic applications, as the discrete logarithm problem in groups of smooth order is easy. If the input N is prime, an expected run time $O((\log N)^{4+\epsilon})$ can be achieved. An example of a cryptographic curve is given in chapter 7. The computation time for such a curve is less than one second.

It should not come as a surprise that our solutions to problem 2 are elliptic curves over prime fields. By Hasse’s theorem, we require that N is contained in the Hasse interval $\mathcal{H}_q$ of some prime power q . It is easy to see that the union of the Hasse intervals $\mathcal{H}_q$ over the prime powers q that are *not* primes is a zero density subset

of $\mathbf{Z}_{\geq 1}$. Solvability of problem 2 for all values of N is therefore in an informal sense ‘equivalent’ to the fact that the union of the Hasse intervals $\mathcal{H}_p$ over the *primes* p contains $\mathbf{Z}_{\geq 1}$.

Defining the Hasse interval $\mathcal{H}_N$ around arbitrary integers N as in formula (1.1), we have the equivalence

$$N \in \mathcal{H}_p \iff p \in \mathcal{H}_N, \quad (1.2)$$

and we see that we want every Hasse interval $\mathcal{H}_N$ around an integer N to contain a prime number p . This amounts to the statement that the size of the ‘gaps’ between consecutive primes around N does not exceed $4\sqrt{N}$. Although prime gaps of this size are not believed to exist, the best proven upper bound on their size is currently $O(N^\alpha)$, with $\alpha = 0.525 > \frac{1}{2}$. A more extensive treatment of prime gaps is given in chapter 2.

1.3 Complex multiplication constructions

Our construction method of a curve with prescribed order relies on complex multiplication (CM) techniques. CM-theory describes the abelian extensions of an imaginary quadratic number field K , i.e., extensions L/K with abelian Galois group $\text{Gal}(L/K)$. It was initiated by Kronecker in the second half of the 19-th century. In 1880, Kronecker gave a conjectural description of the abelian extensions of K . He stated that the abelian extensions of K are generated by values of suitable elliptic and modular functions. It was his *liebster Jugendtraum* (dearest dream of his youth) to prove this. Weber came close to proving this conjecture [65, §169], but overlooked a subtle sign condition. The first solution was given by Takagi in 1920 in his article on general class field theory [61].

The first main theorem of complex multiplication is concerned with the Hilbert class field H – the maximal totally unramified abelian extension – of an imaginary quadratic number field K . It states that H is generated over K by the j -invariant $j(E)$ of an elliptic curve E with endomorphism ring $\mathcal{O}_K$. Furthermore, we have an explicit description of the action of the Galois group $\text{Gal}(H/K)$ on $j(E)$, see chapter 3.

The name complex multiplication can be explained as follows. For an elliptic curve E defined over a number field, the endomorphism ring $\text{End}(E)$ is either isomorphic to $\mathbf{Z}$ or to an imaginary quadratic order $\mathcal{O}$. In the latter case, the curve is said to have *complex multiplication*. Over the field $\mathbf{C}$ of complex numbers we can also represent an elliptic curve as a compact Riemann surface $\mathbf{C}/\Lambda$, with $\Lambda \subset \mathbf{C}$ a

lattice of rank 2. The j -invariant of the curve $E = \mathbf{C}/\Lambda$ is then given by $j(\Lambda)$, where $j : \mathbf{H} \rightarrow \mathbf{C}$ denotes the modular j -function from the upper half plane $\mathbf{H}$ to the field of complex numbers. The first main theorem of CM-theory now states that we have

$$K(j(E)) = H,$$

with $E = \mathbf{C}/\mathcal{O}_K$ an elliptic curve with endomorphism ring $\mathcal{O}_K$. The explicit Galois action enables us to compute the minimal polynomial P_K of $j(E)$ over $\mathbf{Q}$. The polynomial P_K , which has *integer* coefficients, is called the class polynomial for $\mathcal{O}_K$.

CM-theory provides a link between the theory of elliptic curves and algebraic number theory. In chapter 3 we explain how we can use CM-theory to construct a curve of prescribed order N . To every choice of a prime $p \in \mathcal{H}_N$, we associate an imaginary quadratic field $K = K_{p,N}$. We can then construct a curve of order N over $\mathbf{F}_p$ as reduction of a curve E in characteristic zero with endomorphism ring $\mathcal{O}_K$. CM-theory tells us that we may take E to be defined over the Hilbert class field H of K . The prime p splits completely in H . Consider the class polynomial P_K for K . The reduction $\bar{P}_K \in \mathbf{F}_p[X]$ splits completely, and any of its roots is the j -invariant of a curve with N points over $\mathbf{F}_p$.

Every choice $p \in \mathcal{H}_N$ yields a quadratic field $K = K_{p,N}$ and once we have computed the class polynomial P_K for K , it is easy to construct a curve with N points over $\mathbf{F}_p$. Computing P_K takes time $O(|D|^{1+\varepsilon})$, where $D = \text{disc}(K_{p,N})$ is the field discriminant of K . For an arbitrary choice $p \in \mathcal{H}_N$, we expect that D is of the same size as N . This leads to an exponential time algorithm to construct a curve of order N that is even inferior to the naïve algorithm from chapter 2. In problem 1 we have no control over the prime p , but the situation is different in problem 2. In chapter 4 we explain how to pick a prime $p \in \mathcal{H}_N$ for which the field discriminant of $K_{p,N}$ is of almost polynomial size in $\log N$, rather than of exponential size. This is the key to the proof of theorem 1.3.

In chapter 5 we do not focus on the problem of constructing an elliptic curve of prescribed order any more, but concentrate on the problem of computing the class polynomial P_K for a given imaginary quadratic field K . This problem fits into the broader scope of class field theory. Only for $K = \mathbf{Q}$ and for imaginary quadratic fields K it is known how to explicitly compute the class fields of K . For general K , it is one of the unsolved Hilbert problems to find generators for the class fields.

1.4 Class invariants

Let $K = \mathbf{Q}(\sqrt{D})$ be imaginary quadratic of discriminant $D < 0$ and let E be an elliptic curve with endomorphism ring $\mathcal{O}_K$. CM-theory tells us that the minimal polynomial over $\mathbf{Q}$ of $j(E)$ has *integer* coefficients. Denoting this polynomial by P_D , we have

$$H \cong K[X]/(P_D),$$

with H the Hilbert class field of K . There is a classical algorithm to compute P_D . Let $j : \mathbf{H} \rightarrow \mathbf{C}$ be the modular j -function from the upper half plane $\mathbf{H}$ to the field of complex numbers. We explicitly know – see section 3.3 – in which points $\tau_I \in \mathbf{H}$ we should evaluate j to compute a root of $P_D \in \mathbf{C}[X]$. We compute all the roots $j(\tau_I)$ of $P_D \in \mathbf{C}[X]$ with high accuracy and expand the product

$$P_D = \prod_I (X - j(\tau_I)) \in \mathbf{C}[X].$$

If we have computed all roots with high enough accuracy – we have an explicit upper bound for the required precision – we round the coefficients of $P_D \in \mathbf{C}[X]$ to the nearest integer.

In 2002, Couveignes and Henocq published a method [12] to compute P_D by working over $\mathbf{Q}_p$ for a suitable prime p rather than over $\mathbf{C}$. Working over $\mathbf{Q}_p$ has the advantage that we need not worry about rounding errors. Their paper mostly gives the mathematical framework of the p -adic algorithm, and focuses not so much on an actual implementation. An actual implementation is far from straightforward, and all of chapter 5 is devoted to this. We present the algorithm in more detail and explain how one can implement it.

A serious drawback of computing P_D is that the coefficients of this polynomial are *huge*. Not only do they grow exponentially in size for $|D| \rightarrow \infty$, but – perhaps even worse – even for moderately small values of D the coefficients are tremendous. Consider the polynomial for $D = -23$:

$$P_{-23} = X^3 + 3491750X^2 - 5151296875X + 12771880859375 \in \mathbf{Z}[X].$$

History tells us that we should be able to do better. In his *Lehrbuch der Algebra* (1908) Weber explains that function values of ‘smaller’ functions than the j -function sometimes also generate the Hilbert class field [65]. Weber gives a modular function $\mathfrak{f} : \mathbf{H} \rightarrow \mathbf{C}$ of level 48 with the property that $\mathfrak{f}(\omega)$ generates the Hilbert

class field H of $\mathbf{Q}(\sqrt{-23})$ for an appropriate choice of generator ω for the $\mathbf{Z}$ -algebra $\mathcal{O}_K = \mathbf{Z}[\omega]$. We find that $\mathfrak{f}(\omega)$ is a root of

$$P_{-23}^{\mathfrak{f}} = X^3 - X - 1 \in \mathbf{Z}[X].$$

For a modular function $f : \mathbf{H} \rightarrow \mathbf{C}$ and a point $\tau \in \mathbf{H}$, the function value $f(\tau)$ is called a *class invariant* if we have

$$K(f(\tau)) = K(j(\tau)).$$

The function $\mathfrak{f}$ is an example of a function yielding class invariants. The logarithm of the coefficients of its Fourier expansion is 72 times smaller than the logarithm of the coefficients of the j -function. Moreover, the minimal polynomial $P_D^{\mathfrak{f}}$ of a class invariant $\mathfrak{f}(\tau)$ often has integer coefficients. We expect that the logarithmic height of the coefficients of $P_D^{\mathfrak{f}}$ is 72 times smaller than that of P_D . This is a *constant* factor, but it enables us to treat much larger discriminants.

Weber focuses on a few specific functions, such as $\mathfrak{f}$ and a cube root γ_2 of the j -function. He uses ad hoc methods to decide whether $f(\tau)$ is a class invariant, and if so, to compute the Galois conjugates of $f(\tau)$ under $\text{Gal}(H/K)$. For a general approach, we need to understand the Galois action of $\text{Gal}(K_{\text{ab}}/K)$ on values of modular functions. For the j -function this is rather simple, but for modular functions of higher level the situation becomes more complicated. Shimura reciprocity (1971) describes this Galois action, and this is the modern tool for working with class invariants.

Using Shimura reciprocity, it is now a rather mechanical process [26, 59, 49] to decide whether $f(\tau)$ is a class invariant, and if so compute its conjugates under $\text{Gal}(H/K)$. A precise description of Shimura reciprocity, including examples, is given in chapter 6.

The theory of class invariants is firmly rooted in a complex analytic setting. For the j -function we were able to work over the non-archimedean field $\mathbf{Q}_p$ rather than over $\mathbf{C}$. A natural question to consider is the following.

Question. *Can we also work with class invariants in a p -adic setting?*

In chapter 6 we develop a p -adic theory of class invariants, showing that the answer to the question above is *yes*. This combines both improvements to the classical algorithm of computing P_D by evaluating the j -function in suitable points $\tau \in \mathbf{H}$, i.e., we can use smaller functions and we do not have to worry about rounding errors. Partial results in this direction can already be found in [7, Section 5].

The technique we use consists of both Shimura reciprocity and a systematic use of modular curves. The main algorithmic tools are *modular polynomials*. For the j -function these polynomials are well-known, but they also exist for modular functions of higher level. See section 6.8 for a definition of modular polynomials that is inspired by geometry.

The last section of chapter 6, section 6.9, gives the algorithm to work with class invariants over $\mathbf{Q}_p$. This algorithm is illustrated by the examples in chapter 7.

Elliptic curves of given order

2.1 Elliptic curves over finite fields

A classical theorem of Hasse from 1933 states that for an elliptic curve $E/\mathbf{F}_q$, the order of the group $E(\mathbf{F}_q)$ is an integer in the Hasse interval

$$\mathcal{H}_q = [q + 1 - 2\sqrt{q}, q + 1 + 2\sqrt{q}] \quad (2.1)$$

around $q + 1$. The key in understanding this result lies in the endomorphism ring $\text{End}_{\mathbf{F}_q}(E)$ of E , and we give the main ingredients of the proof of Hasse's result.

For an elliptic curve E defined over $\mathbf{F}_q$, the Frobenius map $x \mapsto x^q$ on $\mathbf{F}_q$ induces an endomorphism of the curve $E/\mathbf{F}_q$:

$$F_q : E \rightarrow E \quad (x, y) \mapsto (x^q, y^q),$$

which is called the *Frobenius morphism*. Since the action of the Frobenius morphism on $E(\mathbf{F}_q)$ is raising the coordinates of a point to the q -th power, we have

$$\#E(\mathbf{F}_q) = \#\ker(1 - F_q).$$

The Frobenius morphism is purely inseparable of degree q . As the inseparable endomorphisms form an ideal of $\text{End}_{\mathbf{F}_q}(E)$, we see that $1 - F_q$ is separable. For non-constant separable morphisms, the number of points in the kernel equals the degree, i.e., we have

$$\#\ker(1 - F_q) = \deg(1 - F_q).$$

The ring $\text{End}_{\mathbf{F}_q}(E)$ has an involution that sends $\alpha \in \text{End}_{\mathbf{F}_q}(E)$ to its dual $\hat{\alpha}$. We compute $\deg(1 - F_q) = F_q \hat{F}_q + 1 - (F_q + \hat{F}_q) = q + 1 - (F_q + \hat{F}_q)$. The integer $t = (F_q + \hat{F}_q) \in \mathbf{Z}$ is called the *trace of Frobenius*. The inequality

$$|t| \leq 2\sqrt{q}$$

12 Elliptic curves of given order

now follows from a variant of the Cauchy-Schwarz inequality [58, Lemma 5.1.2].

The proof of Hasse's result shows that the Frobenius morphism $F_q : E \rightarrow E$, which is defined over $\mathbf{F}_q$, satisfies the relation

$$F_q^2 - tF_q + q = 0 \in \text{End}_{\mathbf{F}_q}(E).$$

For $F_q \notin \mathbf{Z}$, we have $\text{disc}(\mathbf{Z}[F_q]) = t^2 - 4q < 0$. Hence, for $t \neq \pm 2\sqrt{q}$, the order $\mathbf{Z}[F_q]$ is isomorphic to the imaginary quadratic order $\mathcal{O}$ of discriminant $t^2 - 4q < 0$. Let $\pi_q \in \mathcal{O}$ be the image of F_q under an isomorphism $\mathbf{Z}[F_q] \xrightarrow{\sim} \mathcal{O}$. Then π_q has norm $N(\pi_q) = \deg(F_q) = q$ and trace $\text{Tr}(\pi_q) = t$. For a curve with N points over $\mathbf{F}_q$, we have $(1 - \pi_q)(1 - \bar{\pi}_q) = N$ and $\pi_q \bar{\pi}_q = q$. This observation gives the symmetric relation

$$N \in \mathcal{H}_q \iff q \in \mathcal{H}_N,$$

where $\mathcal{H}_N$ is defined by the same formula as in (2.1).

The ring $\mathbf{Z}[F_q]$ is a subring of the endomorphism ring $\text{End}_{\bar{\mathbf{F}}_q}(E)$. If $\text{End}_{\bar{\mathbf{F}}_q}(E)$ is imaginary quadratic, then the curve E is called *ordinary*, otherwise E is said to be *supersingular*.

REMARK. For an ordinary elliptic curve $E/\mathbf{F}_q$, we have $\text{End}_{\bar{\mathbf{F}}_q}(E) = \text{End}_{\mathbf{F}_q}(E)$. Indeed, a necessary and sufficient condition for an endomorphism $\alpha \in \text{End}_{\bar{\mathbf{F}}_q}(E)$ to be defined over $\mathbf{F}_q$ is $\alpha \circ F_q = F_q \circ \alpha$. For an ordinary curve E , all endomorphisms $\alpha \in \text{End}_{\bar{\mathbf{F}}_q}(E)$ commute with F_q since the ring $\text{End}_{\bar{\mathbf{F}}_q}(E)$ is commutative. In this case, we will often write $\text{End}(E)$ for the endomorphism ring of E .

THEOREM 2.1. *Let q be a prime power and let $E/\mathbf{F}_q$ be an elliptic curve. If E is supersingular, then $\text{End}_{\bar{\mathbf{F}}_q}(E)$ is isomorphic to a maximal order in a quaternion algebra. Furthermore, E is supersingular if and only if $\text{char}(\mathbf{F}_q)$ divides the trace of the Frobenius morphism $F_q : E \rightarrow E$.*

PROOF. [58, Theorem 3.1] □

For the rest of this section we assume $\text{char}(\mathbf{F}_q) \neq 2, 3$. An elliptic curve over $\mathbf{F}_q$ is determined up to $\bar{\mathbf{F}}_q$ -isomorphism by its j -invariant $j(E) \in \mathbf{F}_q$. We can put any elliptic curve $E/\mathbf{F}_q$ into a Weierstraß form given by $Y^2 = X^3 + aX + b$, and $j(E)$ is defined as

$$j(E) = 1728 \frac{4a^3}{4a^3 + 27b^2} \in \mathbf{F}_q.$$

The j -invariant $j(E)$ determines the endomorphism ring $\text{End}_{\overline{\mathbf{F}}_q}(E)$, so we have supersingular and ordinary j -invariants.

Let $E : Y^2 = X^3 + aX + b$ and $E' : Y^2 = X^3 + a'X + b'$ be two elliptic curves over $\mathbf{F}_q$. The curves E and E' are isomorphic over an extension $L/\mathbf{F}_q$ if and only if there exists $c \in L^*$ with

$$a' = c^4 a \quad \text{and} \quad b' = c^6 b.$$

We see that if E and E' are isomorphic over an extension $L/\mathbf{F}_q$, they are isomorphic over an extension of degree at most 6 of $\mathbf{F}_q$, and isomorphic over a quadratic extension of $\mathbf{F}_q$ if ab is non-zero. The curves having $a = 0$ in their Weierstraß equation have j -invariant 0, whereas the curves with $b = 0$ have j -invariant 1728.

THEOREM 2.2. *Let $\text{char}(\mathbf{F}_q) > 3$ and let $j \in \mathbf{F}_q$. The number of elliptic curves (up to $\mathbf{F}_q$ -isomorphism) with j -invariant j is:*

- (a) 4 if $j = 1728$ and $q \equiv 1 \pmod{4}$;
- (b) 6 if $j = 0$ and $q \equiv 1 \pmod{3}$;
- (c) 2 otherwise.

PROOF. From the discussion above, we have to compute $\#(\mathbf{F}_q^*/\mathbf{F}_q^{*n})$ with $n = 4, 6, 2$ depending on the j -invariant. The result follows. $\square$

THEOREM 2.3. *Let E, E' be elliptic curves over $\mathbf{F}_q$. If E is ordinary, then E and E' are isomorphic over $\mathbf{F}_q$ if and only if we have $j(E) = j(E')$ and $\#E(\mathbf{F}_q) = \#E'(\mathbf{F}_q)$.*

PROOF. If E and E' are isomorphic over $\mathbf{F}_q$, they obviously have the same number of points over $\mathbf{F}_q$. Our proof of the other implication relies on a classical theorem of Tate [62, Theorem 1]. It states that if E and E' have the same number of points over $\mathbf{F}_q$, then they are $\mathbf{F}_q$ -isogenous. Let $\alpha : E \rightarrow E'$ be an isogeny that is defined over $\mathbf{F}_q$. Furthermore, let $\varphi : E' \rightarrow E$ be an isomorphism that is defined over $\overline{\mathbf{F}}_q$. We have

$$\varphi \circ \alpha \in \text{End}_{\overline{\mathbf{F}}_q}(E) = \text{End}_{\mathbf{F}_q}(E),$$

where the equality sign follows from the assumption that E is ordinary. For an element $\sigma \in \text{Gal}(\overline{\mathbf{F}}_q/\mathbf{F}_q)$, we have

$$\varphi^\sigma \circ \alpha = \varphi^\sigma \circ \alpha^\sigma = (\varphi \circ \alpha)^\sigma = \varphi \circ \alpha,$$

and hence $(\varphi^\sigma - \varphi) \circ \alpha = 0$. Since α is surjective, we have $\varphi^\sigma = \varphi$. $\square$

REMARK. Theorem 2.3 does not hold in general for supersingular curves. The supersingular curves $E : Y^2 = X^3 + 4$ and $E' : Y^2 = X^3 + 3$ over $\mathbf{F}_5$ have j -invariant $0 \in \mathbf{F}_5$. Both E and E' have trace of Frobenius 0, but they are not isomorphic over $\mathbf{F}_5$, as $4/3 = 3 \in \mathbf{F}_5^*$ is not a 6-th power.

The curves isomorphic over $\bar{\mathbf{F}}_q$, but not over $\mathbf{F}_q$, to a curve $E/\mathbf{F}_q$ are called the *twists* of E . The j -invariants 0 and 1728 are special. For $j = 0, 1728$ the endomorphism ring of an ordinary elliptic curve over $\mathbf{F}_q$ with j -invariant j equals $\mathbf{Z}[\zeta_3]$, $\mathbf{Z}[i]$ respectively. There are no other ordinary j -invariants with this property.

Let E be an ordinary elliptic curve over $\mathbf{F}_q$ with j -invariant $j(E) \neq 0, 1728$. The unique twist of E is called the *quadratic twist*. If E has $q+1-t$ points, then the quadratic twist of E has $q+1+t$ points. In order to prove this last statement, we let $E'/\mathbf{F}_q$ be an ordinary curve with endomorphism ring $\mathcal{O}$ of discriminant $t^2 - 4q < -4$. By theorem 2.3, it suffices to show that we have $t' = \pm t$ in the diagram below.

$$\begin{array}{ccccc} \mathbf{Z}[F_q] & \hookrightarrow & \mathcal{O} & \xrightarrow{\text{Tr}} & \mathbf{Z} \\ F_q & \longmapsto & \pi_q & \longmapsto & t' \end{array}$$

We know that $\pi_q \in \mathcal{O}$ has norm $q = p^f$. Since E' is ordinary, we have $p \nmid t' = (\pi_q + \bar{\pi}_q) \in \mathbf{Z}$. Hence, we may write $(\pi_q) = \mathfrak{p}^f$, where $\mathfrak{p}$ is an $\mathcal{O}$ -ideal lying over p . By assumption, $\mathcal{O}$ has unit group $\mathcal{O}^* = \{\pm 1\}$, and a generator of $\mathfrak{p}^f$ is therefore determined up to sign. This shows that we have $t' = \pm t$.

Schoof's algorithm [51] gives an efficient way of computing the order $\#E(\mathbf{F}_q)$ of a Weierstraß curve $E : Y^2 = X^3 + aX + b$ over $\mathbf{F}_q$. The main idea behind the algorithm is to compute the trace of Frobenius t modulo many small primes l . Since we have an upper bound $|t| \leq 2\sqrt{q}$ from Hasse's theorem, we can use the Chinese remainder theorem to reconstruct $t \in \mathbf{Z}$ from the values $t \bmod l$. The run time of Schoof's algorithm is polynomially bounded in the input size $\log q$.

2.2 Does there exist a curve with exactly N points?

This section gives necessary conditions for solvability of the leading problem in this thesis. This problem is the 'inverse' problem of the point counting problem considered by Schoof.

PROBLEM. Given an integer $N \in \mathbf{Z}_{\geq 1}$, find a finite field $\mathbf{F}_q$ and an elliptic curve $E/\mathbf{F}_q$ with N rational points over $\mathbf{F}_q$.

The order of the group $E(\mathbf{F}_q)$ is an integer in the Hasse interval

$$\mathcal{H}_q = [q + 1 - 2\sqrt{q}, q + 1 + 2\sqrt{q}] \quad (2.1)$$

around $q + 1$. From this we see that a necessary condition for the solvability of the problem is that every N is contained in some interval $\mathcal{H}_q$. In other words, we want the union $\bigcup_q \mathcal{H}_q$ over all prime powers q to contain $\mathbf{Z}_{\geq 1}$. The contribution to $\bigcup_q \mathcal{H}_q$ coming from true prime powers, i.e., prime powers which are not primes, is a zero density subset of $\mathbf{Z}_{\geq 1}$. Therefore, it is not unreasonable to restrict to *primes* $q = p$ in our problem.

Define $\mathcal{H}_N$ by the same formula as in (2.1) for arbitrary integers N . From the symmetric relation

$$N \in \mathcal{H}_p \iff p \in \mathcal{H}_N$$

from the previous section, we see that $\mathcal{H}_N$ contains a prime if the problem has a solution with $q = p$ prime. This implies that for solvability for all integers N and with q prime, we need that the distance between two consecutive primes near N is at most of size $4\sqrt{N}$.

If we denote the n -th prime by p_n , we want at least

$$p_{n+1} - p_n = O(\sqrt{p_n}) \quad (p_n \rightarrow \infty). \quad (2.2)$$

There is a big difference between proven results and practice regarding the truth of estimate (2.2). The prime number theorem asserts that, on average, the distance between p_{n+1} and p_n is of size $\log p_n$.

PRIME NUMBER THEOREM (2.4). *Denote by $\pi(x)$ the number of primes up to x . Then:*

$$\lim_{x \rightarrow \infty} \frac{\pi(x)}{x/\log x} = 1.$$

In practice one does find that the distance between p_{n+1} and p_n is of size $\log p_n$. Indeed, defining the *gap* between two consecutive primes a and b as $(b - a)/\log a$, the largest known gap occurs [48] between two primes a and b of 16 digits and has size 32.28.

Estimate (2.2) has led to much research in analytic number theory, but it has remained unproved to date. The classical result that there is a prime in the interval $(z, 2z)$ for every $z \in \mathbf{Z}_{\geq 1}$ was improved upon by Hoheisel [31] in 1930. Hoheisel was the first to prove the existence of a constant $\theta < 1$ with $p_n - p_{n-1} = O(p_n^\theta)$.

His initial value $\theta = \frac{32999}{33000}$ has since then been improved by many people. At this moment, the best result [4] known is $\theta = 0.525$.

Hoheisel's original proof and all subsequent improvements use properties of the zeroes of the Riemann zeta function $\zeta(s)$. This function is defined by $\sum_{n=1}^{\infty} n^{-s}$ for $s \in \mathbf{C}$ with $\operatorname{Re}(s) > 1$, and it can be extended analytically to $\mathbf{C} \setminus \{1\}$. It is therefore no surprise that we can do better than $\theta = 0.525$ by assuming the (generalized) Riemann hypothesis. In 1920 Cramér [14] proved, under GRH,

$$p_{n+1} - p_n = O(\sqrt{p_n} \log p_n).$$

See [32, Theorem 12.10] for a modern proof. Cramér's result is close to the expression in (2.2), but we still have an extra logarithmic factor.

We can do much better if we only insist that *nearly all* integers lie in some Hasse interval $\mathcal{H}_p$. Here nearly all is defined as in analytic number theory, i.e., nearly all integers x have property P precisely when

$$\lim_{x \rightarrow \infty} \frac{P(x)}{x} = 1$$

holds, where $P(x)$ denotes the number of integers up to x that have property P . The prime number theorem tells us for instance that nearly all integers are composite.

We define θ_0 by

$$\theta_0 = \inf_{\theta} \{ \text{for nearly all } n \text{ the interval } [n, n + n^{\theta}] \text{ contains a prime} \}.$$

The upper bound $\theta_0 \leq 19/77$ was shown in 1943 by Selberg [54]. We conclude that nearly all integers N arise as the order of an elliptic curve over a finite field.

If we are also willing to assume GRH, the situation is even better. In the same paper Selberg proved that, under GRH, nearly all intervals

$$[n, n + f(n)(\log n)^2]$$

contain a prime, provided $f(n) \rightarrow \infty$ for $n \rightarrow \infty$. The exponent 2 in the logarithm can be lowered to 1 if we moreover assume some vertical distribution of the zeroes on the critical line [30]. This last result implies that, under an extended version of GRH, for nearly all N we can find a prime $p(N)$ that is close to N , i.e., $|p(N) - N| \leq (\log N)^{1+\varepsilon}$ for every $\varepsilon > 0$. We conclude that it is safe to expect that every Hasse interval $\mathcal{H}_N$ around $N + 1$ contains a prime.

Fix a prime $p \in \mathcal{H}_N$, and define $t = p + 1 - N$. We are interested in the number of curves over $\mathbf{F}_p$ that have trace of Frobenius t . Hence, for a fixed integer t with $|t| \leq 2\sqrt{p}$, we want to count the set

$$\{E : E \text{ elliptic curve over } \mathbf{F}_p \text{ with } \text{Tr}(F_p) = t\} / \cong_{\mathbf{F}_p},$$

where we count every isomorphism class $[E]$ with weight $(\#\text{Aut}_{\mathbf{F}_p}(E))^{-1}$. Note that $\#\text{Aut}_{\mathbf{F}_p}(E)$ equals 6 or 4 if $\text{End}_{\mathbf{F}_p}(E)$ is isomorphic to $\mathbf{Z}[\zeta_3]$, $\mathbf{Z}[i]$ respectively, and $\#\text{Aut}_{\mathbf{F}_p}(E) = 2$ otherwise. We use the notation $\#'$ to indicate that we use this weighted cardinality.

Formulas for the number of curves with a prescribed trace of Frobenius go back to Deuring [16]. The answer involves the *Kronecker class number* of the imaginary quadratic order $\mathcal{O}$, which we proceed to define. Write $h'(\mathcal{O}) = h(\mathcal{O})/|\mathcal{O}^*| \in \mathbf{Q}$ for the ‘weighted’ class number.

Definition. The *Kronecker class number* $H'(\Delta)$ of the imaginary quadratic order $\mathcal{O}_\Delta$ of discriminant Δ is

$$H'(\Delta) = \sum_{\mathcal{O}_\Delta \subset \mathcal{O}' \subset \mathcal{O}_{\max}} h'(\mathcal{O}') \in \mathbf{Q},$$

where $h'(\mathcal{O}')$ denotes the weighted class number of $\mathcal{O}'$, and $\mathcal{O}_{\max}$ is the maximal order of $\mathbf{Q}(\sqrt{\Delta})$.

We have the following theorem relating the number of curves with trace of Frobenius t and the Kronecker class number.

THEOREM 2.5. *Let $\mathbf{F}_p$ be a finite prime field. Then the following equality holds:*

$$\#'\{E : E \text{ elliptic curve over } \mathbf{F}_p \text{ with } \text{Tr}(F_p) = t\} / \cong_{\mathbf{F}_p} = H'(t^2 - 4p) \in \mathbf{Q}.$$

PROOF. This is theorem 4.6 in [52]. We will give a proof, for $t \neq 0$, based on the Deuring lifting theorem in chapter 3. $\square$

In particular we see from theorem 2.5 that for any integer t with $|t| \leq 2\sqrt{p}$ there exists an elliptic curve over $\mathbf{F}_p$ with trace of Frobenius t . This does not hold in general [52] if we replace p by a prime power p^f . There are often not enough supersingular curves to cover the cases where t is divisible by p .

2.3 Naïve algorithm

We now formulate the first algorithm for solving our problem. Given that computing the trace of Frobenius of an elliptic curve $E/\mathbf{F}_p$ takes time polynomial in $\log p$, a natural idea is to choose a prime $p \in \mathcal{H}_N$ and construct random curves over $\mathbf{F}_p$ until we have found a correct one. This observation forms the basis of the *naïve algorithm*. We also implement an ‘early abort strategy’ in checking whether a curve is a curve with the correct number of points. From a theoretical point of view this is not very important, since it does not change the asymptotic run time of the algorithm. From a practical point of view it is very important however. We can handle much larger inputs, which is of importance since the naïve algorithm will also be used as a subalgorithm of the main algorithm in chapter 5.

Algorithm (*Naïve algorithm*). Input: an integer $N > 4$. Output: a prime $p \in [N + 1 - \sqrt{N}, N + 1 + \sqrt{N}]$ and an elliptic curve $E/\mathbf{F}_p$ with $|E(\mathbf{F}_p)| = N$ if such a pair (p, E) exists; failure otherwise.

1. Put $a \leftarrow \lceil N + 1 - \sqrt{N} \rceil$.
 - 1a. If $a > N + 1 + \sqrt{N}$, return failure and halt.
 - 1b. If a is prime, set $p \leftarrow a$, $t \leftarrow p + 1 - N$ and go to step 2.
 - 1c. Put $a \leftarrow a + 1$ and go to step 1a.
2. Pick a random element $b \in \mathbf{F}_p^* \setminus \{-\frac{27}{4}\}$.
 - 2a. Define $E_b : Y^2 = X^3 + bX - b$ and $P = (1, 1) \in E_b(\mathbf{F}_p)$.
 - 2b. If $(p + 1 - t)P = O_{E_b}$, compute the trace of Frobenius u for E_b . If $u = t$, return E_b .
 - 2c. If $t \neq 0$ and $(p + 1 + t)P = O_{E_b}$, compute the trace of Frobenius u for E_b . If $u = -t$, return the quadratic twist of E_b .
 - 2d. Return to step 2.

Before we analyse the run time of the algorithm, we give some remarks on the individual steps. From theorem 2.5 we see that if we find a prime p in step 1, there exists an elliptic curve $E/\mathbf{F}_p$ with $|E(\mathbf{F}_p)| = N$. We look for primes in a smaller set than the entire Hasse interval $\mathcal{H}_N$. The reason is that if we would take a prime p close to $N + 1 \pm 2\sqrt{N}$, the associated discriminant $\Delta = t^2 - 4p$ would be very small in absolute value. There are $H'(\Delta)$ curves (up to isomorphism) with trace of Frobenius t , cf. theorem 2.5, and if $|\Delta|$ is very small, then $H'(\Delta)$ is also very small. Hence the probability of ‘hitting’ a correct curve in step 2 would be very small. (More precise statements are provided in the analysis in section 2.4.)

In step 2 we may assume that there exists a curve $E/\mathbf{F}_p$ with N points and with $j(E) \neq 0, 1728$. For b ranging over $\mathbf{F}_p^* \setminus \{-\frac{27}{4}\}$, the j -invariant of the curve E_b attains every value of $\mathbf{F}_p^* \setminus \{1728\}$. For $j \neq 0, 1728$, there are two non-isomorphic curves E, E' with j -invariant j , cf. theorem 2.2. If E has $p + 1 - t$ points, then E' has $p + 1 + t$ points. Both possibilities are tested in steps 2b and 2c.

2.4 Analysis

We proceed with the run time analysis of the algorithm. The run time will be exponential in $\log N$. We use the $\tilde{O}$ -notation to indicate that factors that are of logarithmic size in the main term have been disregarded. More precisely, for two functions $f, g : \mathbf{Z}_{>0} \rightarrow \mathbf{R}_{>0}$, we say that f is $\tilde{O}(g)$ if there exist $N, c \in \mathbf{Z}_{>0}$ such that for all $n \geq N$ we have

$$f(n) \leq g(n)(\log(3 + g(n)))^c.$$

The only case where the algorithm will return ‘failure’ is when the interval $[N + 1 - \sqrt{N}, N + 1 + \sqrt{N}]$ contains no primes. It will then have done $1 + \lfloor 2\sqrt{N} \rfloor$ primality tests, and since primality testing is polynomial time [2], the total run time will be $\tilde{O}(N^{1/2})$.

We now assume that $[N + 1 - \sqrt{N}, N + 1 + \sqrt{N}]$ contains a prime p . Finding one will take time $\tilde{O}(N^{1/2})$; in practice one expects that the distance between $\lceil N + 1 - \sqrt{N} \rceil$ and the next prime is only a power of $\log N$, leading to a heuristic run time that is polynomial in $\log N$. This difference turns out not to be important for the total run time of the algorithm.

In step 2 we have to compute the twists of an elliptic curve. As noted in section 2.1, this boils down to finding a representative for $\mathbf{F}_p^*/\mathbf{F}_p^{*2}$. Doing this probabilistically, we expect that we have to try 2 random elements of $\mathbf{F}_p^*$ before we have a non-square. This can clearly be done in time polynomial in $\log p$. Once we have the twists, we have to compute their group orders. Using Schoof’s algorithm [51], this takes time $\tilde{O}((\log p)^5)$.

To analyse step 2, we need good bounds on the number of curves over $\mathbf{F}_p$ with trace of Frobenius t . From theorem 2.5 we see that this amounts to finding good bounds for the Kronecker class number $H'(t^2 - 4p)$. This is done in [39]; the result is the following lemma.

LEMMA 2.6. *There exist effectively computable constants $c_1, c_2 \in \mathbf{R}_{>0}$ such that for every $z \in \mathbf{Z}_{>1}$ there exists $\Delta^* = \Delta^*(z) < -4$ with*

$$\frac{c_1 \sqrt{-\Delta}}{\log z} \leq H'(\Delta) \leq c_2 \cdot \sqrt{-\Delta} \cdot \log |\Delta| \cdot (\log \log |\Delta|)^2$$

for all negative discriminants $-z \leq \Delta < 0$, except that the left inequality may be invalid if Δ^* is equal to the fundamental discriminant Δ_0 associated to Δ . If GRH holds true, there is no need to exclude an exceptional value Δ^* for Δ_0 .

PROOF. See [39, Proposition 1.8] and the discussion preceding it. $\square$

COROLLARY 2.7. *There exist effectively computable constants $c_1, c_2 \in \mathbf{R}_{>0}$ such that the following is true. Let p be a prime, and let t be an integer with $|t| \leq \sqrt{p}$.*

(i) *We have an upper bound*

$$H'(t^2 - 4p) \leq c_1 \cdot \sqrt{p} \cdot \log p \cdot (\log \log p)^2.$$

(ii) *If GRH holds true, we have a lower bound*

$$H'(t^2 - 4p) \geq c_2 \cdot \sqrt{p} / \log p.$$

(iii) *Let $\Delta' < 0$ be a discriminant with $|\Delta'| \leq 10p$. If the fundamental discriminants associated to $\Delta = t^2 - 4p$ and Δ' are distinct, then at least one of the estimates*

$$H'(\Delta) \geq c_2 \cdot \sqrt{p} / \log p \quad \text{or} \quad H'(\Delta') \geq c_2 \cdot \sqrt{p} / \log p$$

is valid without the assumption of GRH.

PROOF. We apply lemma 2.6 with $z = 4p$. Part (i) follows immediately. For part (ii) we simply note that the assumption $|t| \leq \sqrt{p}$ implies $|t^2 - 4p| \geq 3p$. For part (iii) we apply lemma 2.6 with $z = 10p$. Note that we have $-z \leq \Delta, \Delta' < 0$. By assumption, at least one of the fundamental discriminants associated to Δ and Δ' is not equal to the exceptional value Δ^* . $\square$

Returning to the analysis of step 2 of the algorithm, we see that, under the assumption of GRH, we expect to find a correct curve after $\tilde{O}(p^{1/2})$ tries.

THEOREM 2.8. *If GRH holds true, the naïve algorithm has an expected run time of $\tilde{O}(N^{1/2})$.*

The assumption of GRH may sound a bit heavy. The naïve algorithm is supposed to be *practical* for relatively small N however, so assuming GRH is not much of a problem. From a more theoretical point of view it is of course inconvenient, but the assumption of GRH can be replaced by another assumption.

From part (iii) of corollary 2.7 we see that it suffices to find two primes $p, q \in [N+1-\sqrt{N}, N+1+\sqrt{N}]$ with the property that their associated fundamental discriminants Δ_0, Δ'_0 are distinct. In step 2 of the naïve algorithm we may then work with both $\mathbf{F}_p$ and $\mathbf{F}_q$. We first apply steps 2a–2c with an element $b \in \mathbf{F}_p^*$, then we apply 2a–2c with an element $b \in \mathbf{F}_q^*$, then from $\mathbf{F}_p^*$ again, etc., until we find a curve with N points. The expected run time of this algorithm is $\tilde{O}(N^{1/2})$.

We now analyse how many primes the interval $[N+1-\sqrt{N}, N+1+\sqrt{N}]$ must contain to guarantee the existence of two primes with the property that their associated fundamental discriminants are distinct. Fix a fundamental discriminant $\Delta < -4$. We want to have a good upper bound for the number of solutions (p, f) to

$$(p+1-N)^2 - 4p = \Delta f^2 \quad (2.3)$$

with $p \in [N+1-\sqrt{N}, N+1+\sqrt{N}]$ prime. Just as the relation $p \in \mathcal{H}_N \Leftrightarrow N \in \mathcal{H}_p$ is symmetric in p and N , we have $(p+1-N)^2 - 4p = (N+1-p)^2 - 4N$. Writing $u = N+1-p$, we have to count the number of solutions (u, f) to

$$N = \frac{u+f\sqrt{\Delta}}{2} \cdot \frac{u-f\sqrt{\Delta}}{2} \in \mathcal{O}_\Delta, \quad (2.4)$$

with $N+1-u \in [N+1-\sqrt{N}, N+1+\sqrt{N}]$ prime. Since we do not know anything about the class group of $\mathcal{O}_\Delta$, we cannot say much on the number of elements of norm N . Instead of looking at equation (2.4), we count the number $\rho(N)$ of solutions to

$$N = I\bar{I},$$

with $I \subset \mathcal{O}_\Delta$ an ideal. For primes N we have $\rho(N) = 2$ if N splits, $\rho(N) = 1$ if N ramifies and $\rho(N) = 0$ if N remains inert in $\mathcal{O}_\Delta$. Since we want to derive an upper bound, we now assume that all prime divisors p of N split in $\mathcal{O}_\Delta$. For $N = p^k$ we have $\rho(N) = k+1$. The function $\rho(N)$ is multiplicative and consequently we have $\rho(N) = d(N)$, with $d(N)$ the number of divisors of N .

Since we assumed $\Delta < -4$, a possible generator of an ideal $I \subset \mathcal{O}_\Delta$ is determined up to sign. We see that we have at most $2d(N)$ solutions (p, f) to equation (2.3). Hence, if the interval $[N+1-\sqrt{N}, N+1+\sqrt{N}]$ contains more than $2d(N)$ primes, we can apply the modified naïve algorithm described above.

Unfortunately, the number of divisors $d(N)$ of an integer N grows faster than any power of $\log N$ by [28, Theorem 314]. For every $\varepsilon > 0$, we do have [28, Theorem 315]

$$d(N) = O(N^\varepsilon).$$

THEOREM 2.9. *If the interval $[N+1-\sqrt{N}, N+1+\sqrt{N}]$ contains more than $2d(N)$ primes, the modified naïve algorithm described above has an expected run time of $\tilde{O}(N^{1/2})$.*

The assumption that an interval of length $2\sqrt{N}$ contains at least $2d(N)$ primes is not known to be implied by GRH. As noted in section 2.2, the assumption of GRH implies

$$p_{n+1} - p_n = O(\sqrt{p_n} \log p_n),$$

where p_n is the n -th prime. GRH is not known to imply the existence of a single prime in our interval, let alone $2d(N)$ primes.

The advantage of the modified algorithm is the following. Suppose that we find two primes $p, q \in [N+1-\sqrt{N}, N+1+\sqrt{N}]$ with different associated fundamental discriminants. We now have an *unconditional* expected run time at our disposal. This is quite a contrast with the first algorithm.

2.5 Timings and examples

The condition $N > 4$ in the naïve algorithm ensures that we have $p \geq 5$ for the resulting curve $E/\mathbf{F}_p$ with N points. For completeness sake, we give curves with $N = 1, \dots, 4$ points. The curve with 1 point is defined over $\mathbf{F}_3$, the 3 other curves are defined over $\mathbf{F}_5$.

N	curve
1	$Y^2 = X(X-1)(X-2) + 2$
2	$Y^2 = X^3 + 2X$
3	$Y^2 = X^3 + 4X + 2$
4	$Y^2 = X^3 + X$

As an example of the algorithm, we construct a curve with exactly $N = 10^3$ points. By the prime number theorem, we expect to find roughly $4\sqrt{N}/\log N \approx 18$ primes in the Hasse interval $\mathcal{H}_N$. The interval $\mathcal{H}_N = [938, 1064]$ contains 20 primes.

In step 1, the prime $p = (N+1 - \lfloor \sqrt{N} \rfloor) + 1 = 971$ is selected. Define $t = p+1 - N = -28$. In step 2 we select random values $b \in \mathbf{F}_p^*$ and test whether $E_b : Y^2 = X^3 + bX - b$ has trace of Frobenius $\pm t$. For $b = 237$, the point $P = (1, 1) \in E_b(\mathbf{F}_p)$ is annihilated by $p+1+t$. The trace of Frobenius of E_b is 28 and consequently, the quadratic twist

$$Y^2 = X^3 + 4 \cdot 237X - 8 \cdot 237$$

of E_b has exactly N points.

The naïve algorithm is intended to be practical for relatively small values of N . To test its practical performance, we constructed elliptic curves with exactly $10^7, 10^8, \dots, 10^{13}$ points. To eliminate most of the probabilistic effects, we did this 50 times. The table below gives the average run time in seconds on our standard, 32-bit 2.8 GHz, PC.

N	run time
10^7	< 1
10^8	< 1
10^9	5
10^{10}	104
10^{11}	169
10^{12}	539
10^{13}	1754

The difference in time needed to construct a curve with 10^{11} and 10^{12} points respectively is reasonable in accordance with the expected run time ($169 \cdot \sqrt{10} \approx 534$). Likewise for curves with 10^{12} and 10^{13} points. Something strange seems to be happening for curves for 10^9 and 10^{10} points however. This is probably a classical case where mathematics forgets the laws of computer science: 10^9 still fits in 32 bits, whereas 10^{10} has just crossed this barrier. Computers are far more efficient with numbers of 32 bits than they are with larger numbers.

It is of course a bit dangerous to draw conclusions from this table. It suggests however that with better hardware and improved code (written in assembly for instance), it should be possible to construct a curve with say 10^{20} points in a few hours.

Complex multiplication

3.1 Deuring lifting

This chapter deals with a classical deterministic algorithm for constructing an elliptic curve with exactly N points. We fix a prime $p \in \mathcal{H}_N$ for the remainder of this section. Let $E/\mathbf{F}_p$ be a curve with N points. In chapter 2 we have seen that N satisfies

$$N = p + 1 - t,$$

where t denotes the trace of the Frobenius morphism $F_p : E \rightarrow E$. The quadratic ring $\mathbf{Z}[F_p]$ has discriminant $\Delta = t^2 - 4p < 0$, and the endomorphism ring $\text{End}_{\mathbf{F}_p}(E)$ contains a subring isomorphic to the imaginary quadratic order $\mathcal{O}_\Delta$. Conversely, let $E'/\mathbf{F}_p$ be a curve with $\mathbf{Z}[F'_p] \cong \mathcal{O}_\Delta$, where F'_p is the Frobenius morphism of E' . As an element of norm p in $\mathcal{O}_\Delta$ is determined up to complex conjugation and multiplication by units in $\mathcal{O}_\Delta$, we see that one of the twists of E' has trace t and therefore N points.

This argument shows that finding an elliptic curve E with $\text{End}_{\mathbf{F}_p}(E) \supseteq \mathcal{O}_\Delta$ is equivalent to finding a twist of a curve having $N = p + 1 - t$ points, where we write $\Delta = t^2 - 4p$. As noted in chapter 2, it is very easy to compute the twists of a curve in a probabilistic way. It therefore suffices to find a curve E with $\text{End}_{\mathbf{F}_p}(E) \supseteq \mathcal{O}_\Delta$.

We will not construct such a curve directly in characteristic p , but obtain it as the reduction of a curve in characteristic 0. The following theorem tells us that we can lift an elliptic curve in characteristic p together with an endomorphism.

THEOREM 3.1. (*Deuring lifting*) *Let $E/\mathbf{F}_p$ be an elliptic curve and let $\alpha \in \text{End}_{\mathbf{F}_p}(E)$. Then there exist an elliptic curve A defined over a number field K , an endomorphism $\beta \in \text{End}_K(A)$ and a prime $\mathfrak{P}|p$ of K such that the following is true. The curve A has good reduction at $\mathfrak{P}$. For the reduction $\bar{A} = A \bmod \mathfrak{P}$, there exists an isomorphism $\varphi : \bar{A} \xrightarrow{\sim} E$, and for the induced map $\varphi_* : \text{End}(\bar{A}) \xrightarrow{\sim} \text{End}(E)$ we have $\varphi_*(\bar{\beta}) = \alpha$.*

PROOF. [37, Theorem 13.14]

□

COROLLARY 3.2. *If $E/\mathbf{F}_p$ is ordinary, we can choose A in the Deuring lifting theorem with $\text{End}_K(A) \cong \text{End}_{\mathbf{F}_p}(E)$.*

PROOF. Choose $\alpha \in \text{End}_{\mathbf{F}_p}(E)$ with $\text{End}_{\mathbf{F}_p}(E) = \mathbf{Z}[\alpha]$. We apply the Deuring lifting theorem to the pair (E, α) , yielding an elliptic curve A defined over a number field K . Let G be the reduction modulo $\mathfrak{P}$ of the endomorphism ring $\text{End}_K(A)$. Since endomorphisms reduce injectively, we have an inclusion

$$G \hookrightarrow \text{End}_{\mathbf{F}_p}(\bar{A}) \xrightarrow[\varphi_*]{\sim} \text{End}_{\mathbf{F}_p}(E).$$

The map $G \rightarrow \text{End}_{\mathbf{F}_p}(\bar{A})$ is surjective by our choice of α . □

It is well known [63], that elliptic curves in characteristic 0 have endomorphism rings of rank at most 2 over $\mathbf{Z}$. If the rank equals 2, the curve is said to be a CM-curve, where CM is an abbreviation for *complex multiplication*. Let $E/\mathbf{F}_p$ be a supersingular elliptic curve. Since the endomorphism ring of E is free of rank 4 over $\mathbf{Z}$, we cannot lift the entire endomorphism ring to characteristic zero. Let $\alpha \notin \mathbf{Z}$ be an endomorphism of E . Then α is quadratic over $\mathbf{Z}$, so also in this case we get a CM-‘lift’ of E by applying the Deuring lifting theorem to the pair (E, α) .

3.2 Complex multiplication constructions

The theory of complex multiplication provides us with a means of constructing a curve in characteristic zero with prescribed endomorphism ring. Before we can state the first main theorem of complex multiplication, we need some definitions.

Let K be a field for which there exists an elliptic curve E/K with $\text{End}_K(E) \cong \mathcal{O} = \mathcal{O}_\Delta$. We write $\mathcal{O} = \mathbf{Z}[\alpha]$ for some $\alpha \in \mathcal{O}$. The minimal polynomial $f_{\mathbf{Z}}^\alpha$ of α splits in $K[X]$. We fix a root of $f_{\mathbf{Z}}^\alpha \in K[X]$, and view K as an $\mathcal{O}$ -algebra. There are two isomorphisms $\mathcal{O} \xrightarrow{\sim} \text{End}_K(E)$, and it is important to pin down one of these isomorphisms. We will always consider the *normalized* isomorphism, i.e., the unique isomorphism φ with $\varphi(\alpha)^*\omega = \alpha\omega$ for all $\alpha \in \mathcal{O}$ and all invariant differentials $\omega \in E$. Such a pair (E, φ) is called a *normalized elliptic curve*. Two normalized elliptic curves (E, φ) and (E', φ') are said to be isomorphic if there exists an isomorphism $\tau : E \rightarrow E'$ of elliptic curves with $\tau^{-1}\varphi'(\alpha)\tau = \varphi(\alpha)$ for all $\alpha \in \mathcal{O}$. As there will hardly be any risk of confusion, we usually write E instead of (E, φ) and just speak of an elliptic curve instead of a normalized one.

Let $I \subseteq \text{End}_K(E)$ be an ideal with $N(I)$ coprime to $\text{char}(K)$ and define

$$E[I] = \{P \in E(\bar{K}) \mid \forall \alpha \in I : \alpha(P) = 0\},$$

the group of I -torsion points of E . There exist an elliptic curve E^I and a separable isogeny $\phi : E \rightarrow E^I$ with $\ker(\phi) = E[I]$ by [58, Proposition 3.4.12]. The curve E^I is unique up to K -isomorphism. We get a quotient map $E \rightarrow E^I$ for every ideal $I \subset \mathcal{O}$ coprime to $\text{char}(K)$. The definition of E^I does depend on the choice of an isomorphism $\mathcal{O} \xrightarrow{\sim} \text{End}_K(E)$.

Next we focus on the case that $K = \mathbf{C}$ is the field of complex numbers. A complex elliptic curve with endomorphism ring $\mathcal{O} \subset \mathbf{C}$ is isomorphic to a curve $E_{\mathfrak{a}} = \mathbf{C}/\mathfrak{a}$ for an invertible $\mathcal{O}$ -ideal $\mathfrak{a}$. For an invertible $\mathcal{O}$ -ideal I , the isogeny

$$\begin{array}{ccc} \mathbf{C}/\mathfrak{a} & \rightarrow & \mathbf{C}/(I^{-1}\mathfrak{a}) \\ z & \mapsto & z \end{array}$$

has kernel $E_{\mathfrak{a}}[I]$. We have $E_{\mathfrak{a}}^I \cong E_{I^{-1}\mathfrak{a}}$, and the curve $E_{\mathfrak{a}}^I$ has endomorphism ring $\mathcal{O}$. Let $\text{Ell}_{\Delta}(\mathbf{C})$ be the set of j -invariants of complex elliptic curves with endomorphism ring $\mathcal{O} = \mathcal{O}_{\Delta}$. We have a well-defined map $\rho_I : \text{Ell}_{\Delta}(\mathbf{C}) \rightarrow \text{Ell}_{\Delta}(\mathbf{C})$ sending $j(E)$ to $j(E^I)$. The inverse of ρ_I is given by $\rho_{\bar{I}}$, with $\bar{I}$ the complex conjugate of I . Consequently, the map ρ_I is injective. The map ρ_I gives an action of the group $\mathcal{I}(\mathcal{O})$ of invertible fractional $\mathcal{O}$ -ideals on the set $\text{Ell}_{\Delta}(\mathbf{C})$.

Let $\mathfrak{a}, \mathfrak{b} \subset \mathcal{O}$ be two invertible $\mathcal{O}$ -ideals. We view $\mathfrak{a}, \mathfrak{b}$ as lattices in $\mathbf{C}$. The complex elliptic curves $E_{\mathfrak{a}} = \mathbf{C}/\mathfrak{a}$ and $E_{\mathfrak{b}} = \mathbf{C}/\mathfrak{b}$ are isomorphic if and only if the lattices $\mathfrak{a}$ and $\mathfrak{b}$ are homothetic. In other words: we have $j(\mathbf{C}/\mathfrak{a}) = j(\mathbf{C}/\mathfrak{b})$ if and only if the equality $[\mathfrak{a}] = [\mathfrak{b}]$ holds in the Picard group $\text{Pic}(\mathcal{O})$. The action of $\mathcal{I}(\mathcal{O})$ given by the map $\rho_I : \text{Ell}_{\Delta}(\mathbf{C}) \rightarrow \text{Ell}_{\Delta}(\mathbf{C})$ factors through the quotient map $\mathcal{I}(\mathcal{O}) \twoheadrightarrow \text{Pic}(\mathcal{O})$. We get an action of $\text{Pic}(\mathcal{O})$ on $\text{Ell}_{\Delta}(\mathbf{C})$. This action is simply transitive. The transitivity follows from the equality $\rho_{\mathfrak{b}^{-1}\mathfrak{a}}(j(\mathbf{C}/\mathfrak{a})) = j(\mathbf{C}/\mathfrak{b})$. It is clear that the action is free. We have made $\text{Ell}_{\Delta}(\mathbf{C})$ into a principal homogeneous $\text{Pic}(\mathcal{O})$ -space, or $\text{Pic}(\mathcal{O})$ -torsor. In particular, we see that $\text{Ell}_{\Delta}(\mathbf{C})$ is a *finite* set of cardinality $h(\Delta)$.

Let now K be a number field, and let L/K be a finite abelian extension with discriminant $\Delta_{L/K}$. Let $\mathfrak{p}$ be an $\mathcal{O}_K$ -ideal that is coprime to $\Delta_{L/K}$ and let $\mathfrak{P}|\mathfrak{p}$ be a prime of L . We have an extension of finite fields $(\mathcal{O}_L/\mathfrak{P})/(\mathcal{O}_K/\mathfrak{p})$. This extension is cyclic, and the Galois group is generated by the Frobenius automorphism $x \mapsto x^{N(\mathfrak{p})}$. Since $\mathfrak{P}$ is unramified, there is a unique element $\sigma_{\mathfrak{p}} \in \text{Gal}(L/K)$ mapping to this Frobenius, i.e., $\sigma_{\mathfrak{p}}$ is determined by the condition

$$\sigma_{\mathfrak{p}}(x) \equiv x^{N(\mathfrak{p})} \pmod{\mathfrak{p}\mathcal{O}_L}.$$

This $\sigma_{\mathfrak{p}}$ is called the *Artin symbol* for $\mathfrak{p}$. The map $\mathfrak{p} \mapsto \sigma_{\mathfrak{p}}$ extends multiplicatively

to a homomorphism

$$[\cdot, L/K] : \mathcal{I}(\Delta_{L/K}) \rightarrow \text{Gal}(L/K)$$

from the group $\mathcal{I}(\Delta_{L/K})$ of fractional $\mathcal{O}_K$ -ideals coprime to $\Delta_{L/K}$ to $\text{Gal}(L/K)$.

Let now K be an imaginary quadratic field and $\mathcal{O} = \mathcal{O}_f = \mathbf{Z} + f\mathcal{O}$ the unique order of index $f \geq 1$ in the maximal order $\mathcal{O}_K$. Class field theory tells us that there is a unique abelian extension $H_{\mathcal{O}}/K$ inside a fixed algebraic closure $\bar{K}$, which is unramified outside (f) , such that the Artin map induces an isomorphism

$$\text{Pic}(\mathcal{O}) \xrightarrow{\sim} \text{Gal}(H_{\mathcal{O}}/K).$$

The field $H_{\mathcal{O}}$ is called the *ring class field* for $\mathcal{O}$. The ring class field for $\mathcal{O} = \mathcal{O}_K$ is called the *Hilbert class field* of K . It is the maximal unramified abelian extension of K .

The isomorphism $\text{Pic}(\mathcal{O}) \xrightarrow{\sim} \text{Gal}(H_{\mathcal{O}}/K)$ induced by the Artin map yields the following lemma.

LEMMA 3.3. *Let K be an imaginary quadratic number field and let $\mathcal{O} \subset K$ be the order of index f in $\mathcal{O}_K$. Let $\mathfrak{p}$ be a prime of $\mathcal{O}$ that is coprime to f . Then:*

$$\mathfrak{p} \text{ is principal in } \mathcal{O} \iff \mathfrak{p} \text{ splits completely in } H_{\mathcal{O}}.$$

PROOF. Immediate from the discussion above. □

After these preparations, we can state the first main theorem of complex multiplication.

THEOREM 3.4. *Let $\mathcal{O}$ be an order in an imaginary quadratic field K and write $E = \mathbf{C}/\mathcal{O}$. Then $K(j(E)) = H_{\mathcal{O}}$ and the Galois action of an ideal class $[\mathfrak{a}] \in \text{Pic}(\mathcal{O})$ on $j(E)$ is given by*

$$j(E)^{[\mathfrak{a}, H_{\mathcal{O}}/K]} = j(E^{\mathfrak{a}}).$$

PROOF. [37, Section 10.3]. □

This theorem is the first tool for computing the j -invariant of a curve E with endomorphism ring $\mathcal{O} = \mathcal{O}_{\Delta}$. Let $E/H_{\mathcal{O}}$ be an elliptic curve with endomorphism ring $\mathcal{O}$. Consider the polynomial

$$P_{\Delta} = \prod_{j(E) \in \text{Ell}_{\Delta}(\mathbf{C})} (X - j(E)) \in \mathbf{Q}[X],$$

which is the minimal polynomial of $j(E)$ over $\mathbf{Q}$. The polynomial P_Δ depends only on Δ , and not on the choice of E . The polynomial P_Δ is called the *Hilbert class polynomial* for the order $\mathcal{O}$. The following theorem tells us that P_Δ has *integer* coefficients.

THEOREM 3.5. *Let $E/\mathbf{C}$ be an elliptic curve with $\text{End}(E) \cong \mathcal{O}_\Delta$. Then $j(E)$ is an algebraic integer, i.e.,*

$$P_\Delta \in \mathbf{Z}[X].$$

PROOF. There are at least three different proofs of this theorem. The complex analytic proof proceeds via the same ‘modular polynomials’ that we will use in chapter 5. The ‘good reduction’ proof of Serre and Tate uses local class field theory and the ‘bad reduction’ proof of Serre is based on the observation that if $j(E)$ would not be integral at a prime $\mathfrak{p}$, the curve E would not have complex multiplication. The first two proofs can be found in [57, Section 2.6] and the third proof can be found in [57, Section 5.6]. $\square$

For a prime $p \in \mathcal{H}_N$, write $N = p + 1 - t$ and $\Delta = t^2 - 4p$. For $t \neq 0$, the Hilbert class polynomial P_Δ splits into linear factors in $\mathbf{F}_p[X]$. Indeed, we can write $p = \frac{t+\sqrt{\Delta}}{2} \cdot \frac{t-\sqrt{\Delta}}{2} \in \mathcal{O}$. This implies that the ideal (p) splits into two principal ideals in $\mathcal{O}$ and lemma 3.3 gives us that (p) splits completely in the ring class field $H_\mathcal{O}$.

The roots of $P_\Delta \in \mathbf{F}_p[X]$ are the j -invariants of the elliptic curves over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$. Furthermore, by the Deuring lifting theorem, every curve $E/\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$ arises as the reduction of a curve $A/H_\mathcal{O}$ with endomorphism ring $\mathcal{O}$. Hence, an elliptic curve $E/\mathbf{F}_p$ has endomorphism ring $\mathcal{O}$ if and only if $j(E) \in \mathbf{F}_p$ is a zero of $P_\Delta \in \mathbf{F}_p[X]$.

The theory developed so far can be used to give a proof of theorem 2.5 for ordinary curves, i.e., that we have

$$\#\{E : E \text{ elliptic curve over } \mathbf{F}_p \text{ with } \text{Tr}(F_p) = t \neq 0\} / \cong_{\mathbf{F}_p} = H'(t^2 - 4p).$$

PROOF OF THEOREM 2.5. Assume $t \neq 0$, and write $\Delta = t^2 - 4p$. The prime p splits completely in $H_{\mathcal{O}_\Delta}$, and consequently also in $H_{\mathcal{O}'}$ for any overorder $\mathcal{O}' \supseteq \mathcal{O}_\Delta$. The Hilbert class polynomials $P_{\Delta'}$ for $H_{\mathcal{O}'}$ therefore split completely in $\mathbf{F}_p[X]$. The roots of $P_{\Delta'} \in \mathbf{F}_p[X]$ are the j -invariants of curves over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}'$. We get

$$\#\{E : E \text{ elliptic curve over } \mathbf{F}_p \text{ with } \text{Tr}(F_p) = t\} / \cong_{\mathbf{F}_p} \leq H'(t^2 - 4p).$$

For the other inequality, let $E/\mathbf{F}_p$ be a curve with trace of Frobenius t . By the Deuring lifting theorem it is the reduction of a curve $A/H_{\mathcal{O}'}$ with $\text{End}_{H_{\mathcal{O}'}}(A) \cong \text{End}_{\mathbf{F}_p}(E)$ for some overorder $\mathcal{O}'$. This concludes the proof. $\square$

Section 3 of this chapter gives an algorithm for computing the Hilbert class polynomial P_Δ based on complex analytic methods. A non-archimedean approach is given in chapter 5. Assuming that we can compute P_Δ , we have the following algorithm for constructing an elliptic curve of prescribed order N .

Algorithm. (*CM algorithm*) Input: an integer $N > 6$ and a prime $p \in \mathcal{H}_N$. Output: an elliptic curve $E/\mathbf{F}_p$ with $|E(\mathbf{F}_p)| = N$.

1. Compute the Hilbert class polynomial $P_\Delta \in \mathbf{Z}[X]$ for $\Delta = (p+1-N)^2 - 4p$.
2. Compute a root $j \in \mathbf{F}_p$ of $\bar{P}_\Delta \in \mathbf{F}_p[X]$.
3. Put $a \leftarrow 27j/(4(1728-j))$ and $E : Y^2 = X^3 + aX - a$ for $j \neq 0, 1728$. For $j = 0$, put $E : Y^2 = X^3 + 1$ and for $j = 1728$, put $E : Y^2 = X^3 + X$.
4. Return a twist of E with N points.

THEOREM 3.6. *The CM algorithm will return an elliptic curve over $\mathbf{F}_p$ with exactly N points.*

PROOF. Immediate from the discussion above. $\square$

The main contribution in the run time comes from step 1, i.e., computing the Hilbert class polynomial P_Δ . The run time for both the complex analytic and the non-archimedean approach is $O(|\Delta|^{1+o(1)})$, as we will see in section 3.3 and in chapter 5. Since we have $\Delta = O(N)$, this leads to the following run time.

Run time. *The CM-algorithm has run time $O(N^{1+\varepsilon})$ for every $\varepsilon > 0$.*

This run time is far worse than the run time for the probabilistic version of the naïve algorithm from chapter 2. We can improve the algorithm by noting that it suffices to compute the Hilbert class polynomial P_D for $D = \text{disc}(\mathbf{Q}(\sqrt{\Delta}))$. This usually has little effect on the run time, since the squarefree part of an integer x is typically of the same size as x itself. In our problem we have the freedom to choose the finite field $\mathbf{F}_p$ however. In chapter 4 we explain how to pick a prime p such that $D = \text{disc}(\mathbf{Q}(\sqrt{t^2 - 4p}))$ is of almost polynomial size in $\log N$, rather than of size $O(N)$.

3.3 Complex analytic methods

The classical way of computing P_Δ for a discriminant $\Delta < 0$ proceeds via complex analytic techniques. Let K be the imaginary quadratic field $\mathbf{Q}(\sqrt{\Delta})$ and let $H_\mathcal{O}$ be the ring class field corresponding to the order $\mathcal{O} = \mathcal{O}_\Delta$. We can compute P_Δ as

$$P_\Delta = \prod_{j(E) \in \text{Ell}_\Delta(\mathbf{C})} (X - j(E)) \in \mathbf{Z}[X],$$

and in this section we explain how we can explicitly compute the finite set $\text{Ell}_\Delta(\mathbf{C})$.

Every complex elliptic curve is as a Riemann surface isomorphic to a torus $\mathbf{C}/\Lambda$ for a lattice $\Lambda \subset \mathbf{C}$. More precisely, we can embed $\mathbf{C}/\Lambda$ in $\mathbf{P}^2(\mathbf{C})$ as a Weierstraß curve

$$Y^2 = 4X^3 - g_2(\Lambda)X - g_3(\Lambda),$$

with $g_2(\Lambda) = 60G_4(\Lambda)$, $g_3(\Lambda) = 140G_6(\Lambda)$, and $G_i(\Lambda)$ is the i -th Eisenstein series attached to Λ . A short computation yields that the j -invariant of the curve obtained equals

$$j(\Lambda) = 1728 \frac{g_2(\Lambda)^3}{g_2(\Lambda)^3 - 27g_3(\Lambda)^2} \in \mathbf{C}.$$

After possibly applying a homothety, we may assume $1 \in \Lambda$ and write $\Lambda = \mathbf{Z} + \tau\mathbf{Z}$ for some τ in the upper half plane $\mathbf{H}$. We define $j : \mathbf{H} \rightarrow \mathbf{C}$ by $j(\tau) = j(\mathbf{Z} + \tau\mathbf{Z})$.

The group $\text{SL}_2(\mathbf{Z})$ acts on $\mathbf{H}$ via

$$z \mapsto \frac{az + b}{cz + d} \quad \text{for} \quad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}_2(\mathbf{Z}).$$

The equality of lattices $\mathbf{Z} + \tau\mathbf{Z} = (a\tau + b)\mathbf{Z} + (c\tau + d)\mathbf{Z}$ yields that j is $\text{SL}_2(\mathbf{Z})$ -invariant. In particular, it has a Fourier expansion. It is a classical result that the Fourier expansion of j has integral coefficients. It starts with $q^{-1} + 744 + 196884q$, where $q = \exp(2\pi i\tau)$.

Viewing $\mathcal{O}$ as a lattice in $\mathbf{C}$, the elliptic curve $\mathbf{C}/\mathcal{O}$ has endomorphism ring $\mathcal{O}$. Furthermore, every ideal $I \subset \mathcal{O}$ is a lattice in $\mathbf{C}$ and the curve $\mathbf{C}/I$ has endomorphism ring $\mathcal{O}$ if I is an invertible $\mathcal{O}$ -ideal. This shows that we can compute the Hilbert class polynomial P_Δ as

$$P_\Delta = \prod_{[\mathfrak{a}] \in \text{Pic}(\mathcal{O})} (X - j(\mathfrak{a})) \in \mathbf{Z}[X],$$

where j is now the complex analytic j -function.

We use the standard representation of ideals by binary quadratic forms. This representation is carried out in detail in [10, Section 5.2]; we recall the basic statements here. Let F_Δ^+ be the set of integral positive definite primitive binary quadratic forms of discriminant $\Delta < 0$. We write $[a, b, c]$ for the form $ax^2 + bxy + cy^2 \in F_\Delta^+$. A matrix $A = \begin{pmatrix} p & q \\ r & s \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z})$ acts on F_Δ^+ via $f(x, y)A = f(px + qy, rx + sy)$. As $-1 \in \mathrm{SL}_2(\mathbf{Z})$ acts trivially, we get an action of $\mathrm{PSL}_2(\mathbf{Z})$ on F_Δ^+ . We denote by $\mathcal{F}_\Delta^+$ the set of equivalence classes for F_Δ^+ under this $\mathrm{PSL}_2(\mathbf{Z})$ -action. The map

$$\begin{aligned} \varphi: F_\Delta^+ &\rightarrow I_\Delta \\ [a, b, c] &\mapsto a\mathbf{Z} + \frac{-b+\sqrt{\Delta}}{2}\mathbf{Z} \end{aligned}$$

from F_Δ^+ to the set of fractional ideals I_Δ induces a bijection

$$\varphi: \mathcal{F}_\Delta^+ \rightarrow \mathrm{Pic}(\mathcal{O}_\Delta).$$

In order to use this isomorphism effectively, we agree on a standard representative for an equivalence class in $\mathcal{F}_\Delta^+$. A positive definite quadratic form $[a, b, c]$ is *reduced* if $|b| \leq a \leq c$ and moreover $b \geq 0$ if one of the two inequalities is an equality. This condition is equivalent to saying that the imaginary quadratic number $\tau = \frac{-b+\sqrt{\Delta}}{2a}$ associated to $[a, b, c]$ lies in the standard fundamental domain

$$\left\{ \tau \in \mathbf{H} \mid (\mathrm{Re}(\tau) \in [-\frac{1}{2}, \frac{1}{2}), |\tau| > 1) \quad \text{or} \quad (|\tau| = 1 \text{ and } \mathrm{Re}(\tau) \in [-\frac{1}{2}, 0]) \right\}$$

for $\mathbf{H}$ under the action of $\mathrm{PSL}_2(\mathbf{Z})$. Every class of positive definite quadratic forms contains exactly one reduced form. We see that we can compute P_Δ as

$$P_\Delta = \prod_{[a, b, c] \in \mathcal{F}_\Delta^+} \left(X - j\left(\frac{-b+\sqrt{\Delta}}{2a}\right) \right) \in \mathbf{Z}[X].$$

As we know that P_Δ has integer coefficients, we only have to approximate the j -values in the product with high enough accuracy. We give an estimate for the required precision. For $z = (-b + \sqrt{\Delta})/2a$ we have $|q| = |\exp(2\pi iz)| = \exp(-\pi\sqrt{|\Delta|}/a)$. With a close analysis of the size of the Fourier coefficients for the j -function, one can show [21, 38, 51] that we have $|j(z) - 1/q| \leq 2100$ if z lies in the fundamental domain. Using this estimate, we get the upper bound

$$k = \frac{\pi\sqrt{|\Delta|}}{\log 10} \sum_{[a, b, c] \in \mathcal{F}_\Delta^+} \frac{1}{a}$$

for the number of decimal digits of the constant term $P_\Delta(0) \in \mathbf{Z}$. The number of decimal digits of the largest coefficient of P_Δ is bounded by

$$\log\left(\binom{h}{\lfloor h/2 \rfloor} \cdot \exp(k)\right) \leq 2h + k,$$

where $h = h(\Delta)$ is the degree of P_Δ . As in [1] or [53], we estimate

$$\sum_{[a,b,c] \in \mathcal{F}_\Delta^+} \frac{1}{a} = O((\log |\Delta|)^2).$$

We conclude that we have the estimate $O(\sqrt{|\Delta|}(\log |\Delta|)^2)$ for the required precision in the computation of P_Δ .

In practice, we can compute P_Δ for $|\Delta|$ of size at most 10^{12} in a reasonable amount of time. Often, the constant term of P_Δ is the largest in size. Furthermore, if the constant term is not the largest coefficient, the size of the largest coefficient differs only by a small amount from the constant term. For discriminants down to -10^{12} it is safe to perform the computation with $k + 10$ digits precision.

There are several ways to compute $j(\tau)$. One can for instance use the recursive formulas given for the Fourier coefficients given in [42] or work with the Dedekind η -function as in [3]. In [20] it is noted that it is asymptotically faster to use multi-evaluation to compute all the j -values we want at once. We refer to that paper for the details and give the more naïve algorithm here. This algorithm is much faster *in practice*, i.e., for discriminants down to -10^{12} .

Algorithm. (*Complex analytic class polynomial*) Input: a negative discriminant Δ . Output: the Hilbert class polynomial $P_\Delta \in \mathbf{Z}[X]$.

1. Make a list L of reduced quadratic forms of discriminant Δ .
2. Put $P \leftarrow 1$ and $k \leftarrow \lfloor \frac{\pi\sqrt{|\Delta|}}{\log 10} \sum_{[a,b,c] \in L} \frac{1}{a} \rfloor + \log\left(\binom{h}{\lfloor h/2 \rfloor}\right)$, with $h = h(\Delta)$.
3. For every $[a,b,c] \in L$ do the following:
 Set $P \leftarrow P \cdot (X - j(\frac{-b+\sqrt{\Delta}}{2a}))$, where the j -value is computed with k digits accuracy.
4. Round the coefficients of P to the nearest integer and return P .

One can make a small modification by noting that the complex roots of P_Δ come in conjugate pairs. We can therefore save some evaluations of the j -function. This is done for instance in [10, Section 7.6]. A rigorous run time analysis of this algorithm is not so easy. This analysis has only been undertaken under the simplifying assumption that rounding errors do not play a role in expanding the polynomial P in step 3. The proof of the run time below is given in [20].

Run time. Assume that the precision used in step 2 of the complex analytic algorithm is high enough to neutralise possible rounding errors. Then the algorithm has run time $O(|\Delta|^{3/2+\varepsilon})$ for every $\varepsilon > 0$. With the multi-evaluation modification from [20] the run time becomes $O(|\Delta|^{1+\varepsilon})$ for every $\varepsilon > 0$.

REMARK. The run time for the multi-evaluation approach is in a certain sense best possible. The polynomial P_Δ has degree $|\text{Pic}(\mathcal{O}_\Delta)|$, which grows like $\sqrt{|\Delta|}$ for $|\Delta|$ tending to infinity. Furthermore, the coefficients of P_Δ are of size $\sqrt{|\Delta|}$. We see that just writing down the polynomial P_Δ already takes time at least $O(|\Delta|)$.

3.4 Constructing supersingular elliptic curves

Constructing a supersingular elliptic curve over $\mathbf{F}_p$, which will have $p + 1$ points, is much easier than constructing an ordinary curve of prescribed order. As supersingular curves often are exceptions in the theory developed in the next chapters, this section gives an algorithm to construct a supersingular elliptic curve. The following theorem is fundamental.

THEOREM 3.7. *Let E be a CM curve defined over a number field L with endomorphism ring $\text{End}_L(E) \cong \mathcal{O}$, where $\mathcal{O}$ is an order in an imaginary quadratic field K . Let $\mathfrak{P}|p$ be a prime of L where E has good reduction. Then the reduction $E \bmod \mathfrak{P}$ is supersingular if and only if p does not split in K .*

PROOF. [37, Theorem 13.12] □

Let D be a fundamental discriminant such that p is inert in the imaginary quadratic field $K = \mathbf{Q}(\sqrt{D})$. Since (p) is a principal prime ideal of $\mathcal{O}_K$, lemma 3.3 tells us that (p) splits completely in the Hilbert class field of K . We see that the Hilbert class polynomial P_D splits completely over $\mathbf{F}_{p^2}$. Its roots are j -invariants of supersingular elliptic curves. In fact [63], any supersingular j -invariant lives in $\mathbf{F}_{p^2}$.

To ensure that P_D also has a root in $\mathbf{F}_p$, we demand that the class number $h_K = \deg(P_D)$ be odd. We use genus theory [13, Section 6] to determine the parity of the class number h_K . Let $p_1, \dots, p_n$ be the odd prime factors of D and define $L = K(\sqrt{p_1^*}, \dots, \sqrt{p_n^*})$ with $p_i^* = (-1)^{(p_i-1)/2} p_i$. The field L is called the *genus field* of K . It is the largest unramified extension of K that is abelian over $\mathbf{Q}$.

The Galois group of the extension L/K is isomorphic to the 2-Sylow subgroup of $\text{Pic}(\mathcal{O}_D)$. This means that the class number h_K is odd if and only if we have an equality $L = K$.

We conclude:

$$\begin{aligned} h_K \text{ is odd} &\iff K = \mathbf{Q}(i) \text{ or } K = \mathbf{Q}(\sqrt{-2}) \text{ or} \\ &K = \mathbf{Q}(\sqrt{-q}) \text{ with } q \text{ prime and congruent to } 3 \bmod 4. \end{aligned}$$

This observation leads to the following algorithm.

Algorithm. Input: a prime $p > 3$. Output: a supersingular curve over $\mathbf{F}_p$.

1. If $p \equiv 3 \bmod 4$, return $Y^2 = X^3 - X$.
2. Let q be the smallest prime congruent to 3 mod 4 with $\left(\frac{-q}{p}\right) = -1$.
3. Compute $P_{-q} \in \mathbf{Z}[X]$.
4. Compute a root $j \in \mathbf{F}_p$ of $\bar{P}_{-q} \in \mathbf{F}_p[X]$.
5. If $q = 3$, return $Y^2 = X^3 - 1$. Else, put $a \leftarrow 27j/(4(1728 - j)) \in \mathbf{F}_p$ and return $Y^2 = X^3 + aX - a$.

The correctness of this algorithm is clear from the discussion preceding it. The main point in the run time analysis is step 2. We know that p is congruent to 1 mod 4, so we have $\left(\frac{-q}{p}\right) = \left(\frac{p}{q}\right)$. We therefore want q to be inert in $\mathbf{Q}(\sqrt{p})$ and the condition that q should be congruent to 3 mod 4 translates into the condition that q be inert in $\mathbf{Q}(i)$. The field $L = \mathbf{Q}(\sqrt{p}, i)$ is of degree 4 over $\mathbf{Q}$ and has Galois group $V_4 = \langle \sigma \rangle \times \langle \tau \rangle$, where σ and τ are the non-trivial elements of $\text{Gal}(\mathbf{Q}(\sqrt{p})/\mathbf{Q})$, $\text{Gal}(\mathbf{Q}(i)/\mathbf{Q})$ respectively. The prime q is inert in both $\mathbf{Q}(\sqrt{p})$ and in $\mathbf{Q}(i)$ if and only if the Frobenius of q equals $\sigma\tau \in V_4$.

Just as in chapter 2, there is a big difference between practice and proven results regarding the smallest prime q with prescribed Frobenius $v \in V_4$. The Chebotarev density theorem tells us that the set of primes with prescribed Frobenius v has density 1/4. The error estimates in the proof [35] are very weak, i.e., we can only derive that the smallest prime that has Frobenius $\sigma\tau \in V_4$ is $O(p^\alpha)$ for some $\alpha > 0$. If we assume GRH however, life improves dramatically. Under GRH, there exists an effectively computable constant c such that there exists a prime $q \in \mathbf{Z}$ that is inert in both $\mathbf{Q}(\sqrt{p})$ and in $\mathbf{Q}(i)$ with

$$q \leq c(\log d_L)^2,$$

where $d_L = 2^4 p^2$ is the discriminant of $L/\mathbf{Q}$.

The degree of the class polynomial P_{-q} equals the class number of $\mathbf{Q}(\sqrt{-q})$ and grows like $q^{1/2+o(1)}$. Finding a root $j \in \mathbf{F}_p$ of $\bar{P}_{-q} \in \mathbf{F}_p[X]$ in step 4 takes time $\tilde{O}(\deg(P_{-q})(\log p)^2) = \tilde{O}((\log p)^3)$, cf. [24, Section 14.5]. We summarize the analysis in the following theorem.

THEOREM 3.8. *There exists an algorithm which has as input a prime number p and as output a supersingular elliptic curve over $\mathbf{F}_p$. If GRH holds true, the run time of the algorithm is $\tilde{O}((\log p)^3)$.*

Examples. For $p = 10^{20} + 39$, the elliptic curve given by

$$Y^2 = X^3 - X$$

is supersingular as p is congruent to 3 mod 4. For $p = 10^{20} + 129$, which is congruent to 1 mod 12, the prime $q = 7$ is inert in both $\mathbf{Q}(\sqrt{p})$ and $\mathbf{Q}(i)$. The Hilbert class polynomial for q equals $X + 3375 \in \mathbf{Z}[X]$, so an elliptic curve with j -invariant $-3375 \in \mathbf{F}_p$ is supersingular. The smallest prime $p > 10^{100}$ with $p \equiv 1 \pmod{12}$ is $p = 10^{100} + 1293$. In this case, the prime $q = 11$ is inert in both $\mathbf{Q}(\sqrt{p})$ and $\mathbf{Q}(i)$. An elliptic curve with j -invariant $-32768 \in \mathbf{F}_p$ is supersingular.

An efficient algorithm

4.1 Finding a small discriminant

The complex multiplication algorithm of chapter 3 to construct an ordinary elliptic curve $E/\mathbf{F}_p$ with $N = p + 1 - t$ points has a run time which is dominated by the time needed to construct the Hilbert class polynomial P_Δ , with $\Delta = t^2 - 4p < 0$. We can save some work by computing the class polynomial P_D for the fundamental discriminant $D = \text{disc}(\mathbf{Q}(\sqrt{\Delta}))$ rather than that for Δ itself. As $p = \pi\bar{\pi} \in \mathcal{O}_\Delta$ splits in the same way in the maximal order $\mathcal{O}_D \supset \mathcal{O}_\Delta$ as it does in $\mathcal{O}_\Delta$, ordinary elliptic curves $E/\mathbf{F}_p$ with endomorphism ring $\text{End}(E) = \mathcal{O}_D$ are just as good for our purposes, and we may everywhere replace Δ by D in the algorithm. If Δ has a large square factor, this can be a considerable improvement since the polynomial P_D is then much smaller than P_Δ .

In our problem we usually have *many* primes $p \in \mathcal{H}_N$ to choose from, and every prime p leads to a field discriminant $D(p) = \text{disc}(\mathbf{Q}(\sqrt{\Delta}))$ with

$$\Delta = \Delta(p, N) = t^2 - 4p = (p + 1 - N)^2 - 4p. \quad (4.1)$$

This is exactly the difference with the problem of constructing an elliptic curve with N points over a *prescribed* prime field $\mathbf{F}_p$ that was mentioned in the introduction. In the latter case we have no control over the discriminant $\Delta(p, N)$, which will typically be of the same order of magnitude as N and without large square factors. The resulting run time $\tilde{O}(N)$ is then inferior to the $\tilde{O}(N^{1/2})$ of the naïve probabilistic method from chapter 2.

As explained in chapter 2, we cannot prove the existence of a prime $p \in \mathcal{H}_N$. This means that we will have to rely on *heuristics* for the minimal value of $D(p)$ for p ranging over $\mathcal{H}_N$. The first thing that comes to mind is to choose $p \in \mathcal{H}_N$ as close as possible to one of the end points of $\mathcal{H}_N$. The trace $t = p + 1 - N$ is then very close to $2\sqrt{p}$. By the prime number theorem, we expect that we can choose p for which

$|t| - 2\sqrt{p}$ is of size $\log N$. This makes $\Delta = (t - 2\sqrt{p})(t + 2\sqrt{p})$ of size $O(N^{1/2+o(1)})$ and reduces the run time to $O(N^{1/2+o(1)})$, just as for the naïve probabilistic method.

More generally, one can examine which primes p at distance at most N^α from the end points of $\mathcal{H}_N$ give rise to values of Δ with large square factors. Heuristically, there are about $N^\alpha / \log N$ such primes, giving rise to discriminants of size $N^{\alpha+1/2}$. Among the discriminants of this size, those of the form $\Delta = f^2 D$ with $|D| < N^\beta$ constitute a fraction of order of magnitude

$$P(\alpha, \beta) = N^{-(\alpha+1/2)} \sum_{\substack{|D| < N^\beta \\ \text{squarefree}}} \sqrt{\frac{N^{\alpha+1/2}}{|D|}} \approx N^{\frac{1}{2}(\beta-\alpha)-\frac{1}{4}}.$$

The number of discriminants $\Delta = f^2 D$ with $|D| < N^\beta$ we expect to find from p 's no further than N^α from the end points of $\mathcal{H}_N$ is therefore

$$P(\alpha, \beta) \cdot \frac{N^\alpha}{\log N} = \frac{1}{\log N} \cdot N^{\frac{1}{2}(\alpha+\beta)-\frac{1}{4}},$$

which tends to infinity with N exactly when we have $\alpha + \beta > 1/2$. Rough as this heuristic analysis may be, it ‘explains’ why in the example $N = 10^{30}$ given in [7, Section 6] to illustrate the non-archimedean approach to computing class polynomials, examining the primes p at distance $< 10^6$ from the end points of $\mathcal{H}_N$ leads to a fundamental discriminant $D \approx -10^8$. As examining the primes in an interval of length N^α to achieve $|D| < N^\beta$ gives rise to a run time $\tilde{O}(N^{\max\{\alpha, \beta\}})$, we can achieve a heuristic run time $O(N^{\frac{1}{4}+\varepsilon})$ by taking $\alpha = \beta = \frac{1}{4} + \varepsilon$. Although this is still exponential, this method of selecting p already enables us to deal with values of N the naïve method cannot handle.

The extreme case $(\alpha, \beta) = (\varepsilon, 1/2)$ corresponds to taking p as close as possible to the end points of $\mathcal{H}_N$, a case we already discussed. The other extreme $(\alpha, \beta) = (1/2, \varepsilon)$ indicates that it should be possible to find D of *subexponential* size in terms of our input length $\log N$. This suggests that a fruitful approach to constructing a curve of prescribed order N by the complex multiplication method consists in efficiently minimizing the fundamental discriminant D involved.

It turns out that we can actually determine the ‘minimal’ imaginary quadratic fundamental discriminant D that can be used to construct an elliptic curve of order N in a relatively straightforward way. It uses the ‘symmetry’ between the order N of the point group $E(\mathbf{F}_p)$ and the order p of $\mathbf{F}_p$ itself, which are norms of the quadratic integers $1 - \pi = 1 - F_p$ and $\pi = F_p$, respectively. This symmetry is already familiar to us from section 2.1. In the case of the discriminant $\Delta = (\pi - \bar{\pi})^2 = ((1 - \pi) - (1 - \bar{\pi}))^2$

in (4.1), it takes the form

$$\Delta(p, N) = (p + 1 - N)^2 - 4p = (N + 1 - p)^2 - 4N.$$

We now fix N and try to write $\Delta = \Delta(p)$ as

$$\Delta(p) = (N + 1 - p)^2 - 4N = f^2 D \quad (4.2)$$

for ‘small’ $D < 0$. This comes down to solving the positive definite equation

$$x^2 - Df^2 = 4N \quad (4.3)$$

in integers x and f in such a way that the number $p = N + 1 - x$ is prime. This leads us to the following problem.

PROBLEM 4.1. *Given an integer $N \geq 1$, find the smallest squarefree integer $d \geq 1$ together with an algebraic integer $\alpha \in K = \mathbf{Q}(\sqrt{-d})$ such that*

- (i) $N_{K/\mathbf{Q}}(\alpha) = N$;
- (ii) $p = N_{K/\mathbf{Q}}(1 - \alpha) = N + 1 - \text{Tr}_{K/\mathbf{Q}}(\alpha)$ is prime.

The prime p occurring in condition (ii) has the property that there exists an elliptic curve $E/\mathbf{F}_p$ having N points and endomorphism ring $\text{End}(E)$ isomorphic to the ring of integers $\mathcal{O}_K$ of $K = \mathbf{Q}(\sqrt{-d})$. Once we find the solution (α, d) to problem 4.1, we can use it to construct a curve with N points: take $p = N_{K/\mathbf{Q}}(1 - \alpha)$ and construct an elliptic curve over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}_K$ for which $1 - \alpha \in \mathcal{O}_K$ is the Frobenius, using the class polynomial for the order $\mathcal{O}_K$. This elliptic curve will have $N = N_{K/\mathbf{Q}}(\alpha)$ points, as desired.

We cannot prove that a solution (α_0, d_0) to problem 4.1 exists, let alone that it can be found in time polynomial in $\log N$. In the next section we will give a heuristic analysis showing that it is reasonable to expect that a solution d_0 to problem 4.1 is of size $O((\log N)^2)$. Moreover, finding all algebraic integers $\alpha \in K = \mathbf{Q}(\sqrt{-d})$ of norm N for all squarefree d up to d_0 can be done efficiently using the 1908 algorithm of Cornacchia in the case that we have the prime factorization of N at our disposal. Together this will lead to the following theorem.

THEOREM 4.2. *There exists an algorithm that, on input of an integer $N \geq 1$ together with its factorization, returns a prime number p and an elliptic curve $E/\mathbf{F}_p$ with $\#E(\mathbf{F}_p) = N$ whenever such a pair (E, p) exists. Under standard heuristic assumptions, a pair (E, p) exists for all N , and the expected run time of the algorithm is polynomial in $2^{\omega(N)} \log N$. Here $\omega(N)$ denotes the number of distinct prime factors of N .*

Although the run time in theorem 4.2 is not polynomial in the usual sense, it is polynomial in $\log N$ outside a zero density subset of $\mathbf{Z}_{\geq 1}$ consisting of very smooth input values N . Note that such N are not used in cryptographic applications, as the discrete logarithm problem in groups of smooth order tends to be easy.

COROLLARY 4.3. *If the input values N in theorem 4.2 are restricted to be prime numbers or, more generally, to be in the density 1 subset of $\mathbf{Z}_{\geq 1}$ consisting of those N having $\omega(N) < 2 \log \log N$, then the expected run time is polynomial in $\log N$.*

The factorization of N is used by the algorithm in theorem 4.2 to reduce square root extractions of small integers modulo N to square root extractions modulo the prime factors of N . It is here that the approximate number $2^{\omega(N)}$ of such roots enters the run time of the algorithm.

The precise exponents in the run time depend on one's willingness to accept fast multiplication techniques and probabilistic subroutines in the algorithm. For instance, the square root extractions of small integers modulo the prime factors of N can be done efficiently by probabilistic means or, much less efficiently, but still in time polynomial in $2^{\omega(N)} \log N$, by a deterministic algorithm [51]. Similarly, one may require for the prime number p returned by the algorithm that its primality is proved by a deterministic AKS-type polynomial time algorithm, or employ a faster probabilistic algorithm to do so. If we insist on guaranteed correct output, i.e., a *proven* prime p as the characteristic of our curve E , but allow fast multiplication and probabilistic subroutines of the kind mentioned above, the heuristic run time of our algorithm is $O(2^{\omega(N)} (\log N)^{4+\varepsilon})$ for every $\varepsilon > 0$ (Corollary 4.7.). In the cryptographically relevant case where N is prime, this becomes $O((\log N)^{4+\varepsilon})$ (Corollary 4.5). Sections 4.2–4.4 are taken almost verbatim from [6].

4.2 An algorithm to solve problem 4.1

As indicated in section 2.1, it is currently not possible to prove rigorously that *any* pair (α, d) meeting the conditions of problem 4.1 exists at all, let alone that there is a pair with small d that can be found efficiently. We will however argue in the next section why it is reasonable to expect that the smallest integer d solving problem 4.1 exists for all $N \geq 1$, and why this d is even rather small in terms of N , of size at most $\tilde{O}((\log N)^2 + 2^{\omega(N)})$. Given this expectation, it makes sense to solve problem 4.1 in a straightforward way using an algorithm that, on input of a factored number N , tries for increasing squarefree numbers $d \in \mathbf{Z}_{\geq 1}$ to

- find the integral ideals in $K = \mathbf{Q}(\sqrt{-d})$ of norm N ;
- determine the generators of those ideals that are principal;
- test for each generator α found whether $N_{K/\mathbf{Q}}(1 - \alpha)$ is prime.

As soon as a prime value $p = N_{K/\mathbf{Q}}(1 - \alpha)$ is encountered for some d , this is the minimal d we are after, and (α, d) is a solution to problem 4.1.

Before we describe an actual algorithm, we look at the three individual tasks to be performed, and the run time of the various subroutines involved. These run times depend on the time $O(L^{1+\mu})$ needed to multiply two L -bit integers. We have $\mu = 1$ for ordinary multiplication, and $\mu = \varepsilon > 0$ for any fast multiplication method. We will give our run times using $\mu = \varepsilon > 0$ for a fast multiplication method.

Task 1: Finding the integral ideals in $\mathbf{Q}(\sqrt{-d})$ of norm N .

Write the ring of integers of $\mathbf{Q}(\sqrt{-d})$ as $\mathbf{Z}[\omega]$, with $\omega = \omega_d$ a zero of

$$f = f_{\mathbf{Q}}^{\omega} = \begin{cases} X^2 - X + \frac{1+d}{4} & \text{if } -d \equiv 1 \pmod{4}; \\ X^2 + d & \text{otherwise.} \end{cases} \quad (4.4)$$

Then every ideal of norm N in $\mathbf{Z}[\omega]$ can be uniquely written as kI , with k a positive integer for which k^2 divides N , and I a *primitive* ideal of $\mathbf{Z}[\omega]$ of norm $N_0 = N/k^2$. This last condition means that $\mathbf{Z}[\omega]/I$ is, as a group, cyclic of order N_0 , and it implies that we have $I = (N_0, \omega - r)$ for some integer $r \in \mathbf{Z}$ satisfying $f(r) \equiv 0 \pmod{N_0}$. Finding all ideals of norm N therefore amounts to finding, for each square divisor $k^2 | N$, the roots of f modulo $N_0 = N/k^2$. It is here that we require the factorization of N , not only because this implicitly encodes a list of square divisors $k^2 | N$, but also because it enables us to find the roots of f modulo N_0 . Indeed, finding these roots is done by finding the roots of f modulo the prime powers $p^{\text{ord}_p(N_0)}$ dividing N_0 , and combining these in all possible ways, using the Chinese remainder theorem, to obtain the roots modulo N_0 . Note that f has *no* roots modulo N_0 if N_0 is divisible

by a prime p that is inert in $\mathbf{Z}[\omega]$, or by the square p^2 of a prime p that ramifies in $\mathbf{Z}[\omega]$.

As finding a root of f modulo an integer essentially amounts to extracting a square root of $-d$ modulo that integer, we need to extract square roots of $-d$ modulo the prime powers dividing N_0 . This easily reduces to extracting square roots of $-d$ modulo each of the primes dividing N_0 . This can be done efficiently by employing a variant of the probabilistic Cantor-Zassenhaus algorithm [24, Section 14.5], and leads to an expected run time $O((\log p)^{2+\varepsilon})$ to extract square roots modulo a prime p . For any selection of square roots $(\sqrt{-d} \bmod p^{\text{ord}_p(N_0)})$, the Chinese remainder theorem lifts these to a square root modulo N_0 in time $O(\omega(N)(\log N)^2)$.

Task 2: Finding generators for principal ideals of norm N .

For each ideal $kI = k \cdot (N_0, \omega - r) \subset \mathbf{Z}[\omega]$ of norm N found, we use the 1908 algorithm of Cornacchia described in [50, pp. 229–232] or [8] to find a generator of I , if it exists. This algorithm performs a number of steps of the Euclidean algorithm to the basis elements N_0 and $\omega - r$ of the $\mathbf{Z}$ -lattice $I = (N_0, \omega - r) \subset \mathbf{Z}[\omega]$ in order to decide whether I is a principal ideal. If it is, a generator $\alpha = k\alpha_0$ of kI of norm N is found. For $d \neq 1, 3$, the unique other generator of I is $-\alpha$. For the special values $d = 1$ and $d = 3$ there are 4 and 6 generators for each principal ideal I , respectively, obtained by multiplying α by 4th and 6th roots of unity. The run time of Cornacchia’s algorithm on input $k \cdot (N_0, \omega - r)$ is of order $O((\log N)^{2+\varepsilon})$.

Task 3: Testing which algebraic integers α of norm N lead to prime elements $1 - \alpha$.

For each of the elements α of norm N found in the previous step 2, we need to test whether the norm $N + 1 - \text{Tr}(\alpha)$ of $1 - \alpha$ is a prime number. As most α ’s will have norms that are not prime, a cheap compositeness test such as the Miller-Rabin test, which takes time $\tilde{O}((\log N)^2)$, can be used to discard most α ’s. Once we find an α for which $N + 1 - \text{Tr}(\alpha)$ is a probable prime, we do a true primality test to *prove* primality of $p = N + 1 - \text{Tr}(\alpha)$. This can be done deterministically in time polynomial in $\log N$ by the 2002 result of Agrawal, Kayal and Saxena [2]. Recent speed-ups of the test [41] take time $O((\log N)^{6+\varepsilon})$, whereas probabilistic versions [5] have expected run time $O((\log N)^{4+\varepsilon})$.

Using the various subroutines specified in the tasks above, we formulate an algorithm to solve problem 4.1. A slightly more practical algorithm that we use to actually find elliptic curves with a given number of points does not exactly follow the outline below; it is discussed in section 4.4. The version in this section is phrased to facilitate the heuristic run time estimate in section 4.3.

Algorithm.

Input: a factored integer $N = \prod_{i=1}^t p_i^{e_i}$. Output: a solution (d, α) to problem 4.1.

1. Put $d \leftarrow 1$.
2. If d is not squarefree, put $d \leftarrow d + 1$ and go to step 2. Otherwise, define $\omega = \omega_d$ and $f = f_Q^\omega$ as in (4.4).
3. Determine the splitting behavior in $\mathbf{Z}[\omega]$ of all prime divisors of N .
 - 3a. Put $k_1 \leftarrow 1$. For every prime divisor p_i of N that is inert in $\mathbf{Z}[\omega]$, put

$$k_1 \leftarrow k_1 p_i^{\lfloor e_i/2 \rfloor}$$

in case e_i is even. In case e_i is odd, put $d \leftarrow d + 1$ and go to step 2.

- 3b. For every prime divisor p_i of N that ramifies in $\mathbf{Z}[\omega]$, put

$$k_1 \leftarrow k_1 p_i^{\lfloor e_i/2 \rfloor}.$$

4. Put $N_1 \leftarrow N/k_1^2$. For every root $(r \bmod N_1)$ of f and for every square divisor $k_2^2 \mid N_1$ do the following.
 - 4a. Put $k \leftarrow k_1 k_2$ and $N_0 \leftarrow N/k^2 = N_1/k_2^2$. Use Cornacchia to find a generator of $(N_0, \omega - r) \subset \mathbf{Z}[\omega]$, in case it exists.
 - 4b. If a generator is found, test for all (2, 4 or 6) generators α_0 whether the norm $N + 1 - \text{Tr}(k\alpha_0)$ of $k\alpha_0 \in \mathbf{Z}[\omega]$ is prime. If it is, return d and $\alpha = k\alpha_0$ and halt.
5. Put $d \leftarrow d + 1$ and go to step 2.

The determination of the splitting behavior of the primes $p_i \mid N$ in $\mathbf{Z}[\omega]$ in step 3 amounts to computing the Kronecker symbol $\left(\frac{D}{p_i}\right)$ for $D = \text{disc}(\mathbf{Q}(\sqrt{-d}))$. For $p > 2$ this is simply the Legendre symbol, which is easily evaluated by combining quadratic reciprocity with the Euclidean algorithm. The factor k_1 computed in this step is the minimal ‘imprimitivity factor’ dividing all ideals of norm N in $\mathbf{Z}[\omega]$. It reflects the fact that primitive ideals are not divisible by inert primes, or by squares of ramified primes.

The computation of the roots of f modulo N_1 in step 4 is done by computing the roots of f modulo the various prime powers dividing N_1 , and combining these in all possible ways using the Chinese remainder theorem. For the ramified primes p_i dividing N_1 , which occur with exponent 1, there is a unique (double) root of f modulo p_i . For splitting primes p_i , the polynomial f has exactly 2 different roots modulo p_i , and these lift uniquely to $\mathbf{Z}_{p_i}$. Finding the roots of f modulo these p_i is non-trivial as it involves the extraction of a square root $\sqrt{-d}$ modulo p_i . Refining these roots to roots modulo $p_i^{e_i}$ is much faster, and an easy application of Hensel’s

lemma. The number of distinct roots modulo N_1 is $2^s \leq 2^{\omega(N)}$, with s the number of $p_i|N$ that split in $\mathbf{Z}[\omega]$.

Step 4 computes the possible generators of the primitive parts of ideals of norm N in $\mathbf{Z}[\omega]$. It is not completely optimized as it does not take into account that different roots of f modulo N_1 may coincide modulo N_0 , and give rise to the same ideal $(N_0, \omega - r)$ in step 4a. It also unnecessarily treats the complex conjugate $(N_0, \omega - r')$ of every ideal $(N_0, \omega - r)$, whose generators (if any) are of course the complex conjugates of the generators of $(N_0, \omega - r)$.

4.3 Heuristic run time analysis

In this section, we present a heuristic run time analysis of the algorithm in the previous section, and numerical data supporting this analysis.

Assumption 1. For the elements $\alpha = k\alpha_0 \in \mathbf{Z}[\omega]$ of norm N that we find in step 4a of our algorithm, the norm of $1 - \alpha$ will be an element of the Hasse interval $\mathcal{H}_N$ that, apart from being congruent to 1 mod k , does not appear to have any predictable primality properties. Based on the prime number theorem, a reasonable assumption is therefore that for varying d , r and N_0 , the norms found in step 4b will be prime with ‘probability’ at least $1/\log N$. In other words, the number of times we expect to execute step 4b of our algorithm before we find a prime value is of order of magnitude $\log N$.

Assumption 2. The input for step 4b is provided by step 4a, which finds the generators of those ideals of norm N in $\mathbf{Z}[\omega]$ that are principal. The likelihood for a ‘random’ ideal in $\mathbf{Z}[\omega]$ to be principal is $1/h_d$, with h_d the class number of the ring of integers $\mathbf{Z}[\omega] \subset \mathbf{Q}(\sqrt{-d})$. As we have no indication that the primitive ideals of norm N_0 arising in step 4a behave differently from random ideals in $\mathbf{Z}[\omega]$, it seems reasonable that they will be principal with ‘probability’ around $1/h_d$.

The class number h_d behaves somewhat irregularly as a function of d , but its growth rate $d^{\frac{1}{2}+o(1)}$ was already found by Siegel. In order to bound the number of times we execute the steps 4a and 4b, we need to bound the integers d we encounter in step 2, i.e., to find an upper bound B_N for the minimal integer d that occurs in a solution to problem 4.1. Clearly, such an upper bound will be of heuristical nature, based on the two ‘randomness assumptions’ above. As our algorithm consists of a loop over $d = 1, 2, 3, \dots$, and d has to be factored in step 2 to find if it is squarefree, the value of B_N is of great importance in estimating the run time, and the success of our method depends on B_N being ‘small’ as a function of N .

Elliptic curves of prime order. In the case our input number N is prime, our algorithm is similar to the first step of the elliptic curve primality proving algorithm ECPP. On input N , this algorithm looks for an imaginary quadratic field K of small discriminant containing an element α of norm N with the property that $N_{K/\mathbf{Q}}(1 - \alpha) = N + 1 - \text{Tr}_{K/\mathbf{Q}}(\alpha)$ is *twice* a probable prime number N' . If $\alpha \in K$ is found, N becomes the order of the finite field $\mathbf{F}$ and $2N'$ the number of points of an elliptic curve over $\mathbf{F}$. As $\#\mathbf{F}$ and $\#E(\mathbf{F})$ occur symmetrically in all considerations, this problem is almost identical to our problem 4.1. In fact, since finding a prime around a large number N is heuristically just as difficult as finding twice a prime around N , the heuristic run time for our algorithm on prime input N is *identical* to the heuristic run time for the first step of ECPP on input N . In accordance with the results in [45, Section 3], we obtain the following.

THEOREM 4.4. *Let N be a prime number. Under the heuristic assumptions 1 and 2, the integer d solving problem 4.1 is of size $\tilde{O}((\log N)^2)$, and our algorithm can be expected to find it in time $O((\log N)^{4+\varepsilon})$.*

COROLLARY 4.5. *Under the heuristic assumptions 1 and 2, an elliptic curve with prime order N can be constructed in time $O((\log N)^{4+\varepsilon})$.*

PROOF OF 4.5. We first use our algorithm to find d , α and $p = N - 1 + \text{Tr}(\alpha)$ solving problem 4.1 for N ; the time $O((\log N)^{4+\varepsilon})$ needed for this dominates the steps that follow. We then construct the class polynomial P_D for $D = \text{disc}(\mathbf{Q}(\sqrt{-d}))$ in time $\tilde{O}(d) = \tilde{O}((\log N)^2)$. As P_D has degree $h_d \approx \sqrt{d}$, finding a root j of P_d in $\mathbf{F}_p$ takes time $\tilde{O}(\deg(P_d)(\log p)^2) = \tilde{O}((\log N)^3)$, cf. [24, Section 14.5]. An elliptic curve E with j -invariant j and its quadratic twist E' will have $N = p + 1 - \text{Tr}(\alpha)$ or $p + 1 + \text{Tr}(\alpha)$ points. Matching the group order with the curve can be done efficiently by determining which of the two quantities annihilates random points on the curve. We know that only one of them does for either E or E' for all $p > 229$ by [50, Theorem 3.2]. $\square$

PROOF OF 4.4. For prime input N , our algorithm is rather simple. For increasing values of d , it singles out those d for which N is not inert in $\mathbf{Z}[\omega_d]$ in step 3; in step 4, it computes the primes over N in $\mathbf{Z}[\omega_d]$ and determines whether these are principal with a generator α for which $1 - \alpha$ is a prime element.

The ring $\mathbf{Z}[\omega_d]$ contains elements α of norm N if and only if N splits into principal primes of norm N . For primes N coprime to $2d$, this means that N has to split completely in the Hilbert class field H_d of $\mathbf{Q}(\sqrt{-d})$. Our assumption 2, which states that primitive ideals of norm N should be principal in $\mathbf{Z}[\omega]$ with ‘probability’

$1/h_d$, now reminds us of the Chebotarev density *theorem*, which tells us that one out of every $2h_d = [H_d : \mathbf{Q}]$ primes splits completely in H_d . For $d > 3$, it leads us to expect with ‘probability’ $1/(2h_d)$ that there are (up to conjugation) exactly two integral elements α and $-\alpha$ of norm N . With complementary probability $1 - (2h_d)^{-1}$, there are no elements of norm N . Thus, a value d can be expected to yield an ‘on average’ number of $1/h_d$ elements of norm N .

The average statement that the number of algebraic integers $\alpha \in \mathbf{Q}(\sqrt{-d})$ of norm N is asymptotically a fraction $1/h_d$ of the pairs (d, N) tried is implied by Chebotarev’s theorem in case we fix d and let the prime N vary. We are however in the case where N is fixed and d varies. This is certainly different, but for varying d up to a bound B that is *small* with respect to N , it is assumption 2 that we will find approximately $\sum_{d < B} 1/h_d$ elements of fixed norm N . This is reasonable, provided that the fields H_d are ‘close’ to being linearly independent over $\mathbf{Q}$.

It is not exactly true that the Hilbert class fields H_d for the squarefree integers $d < B$ we encounter form a linearly disjoint family of number fields: the genus fields $G_d \subset H_d$ have many non-trivial intersections. However, in this family of fields, which has about $(6/\pi^2)B$ elements, there is a subfamily of fields H_d coming from the prime numbers $d \equiv 3 \pmod{4}$ that is linearly disjoint over $\mathbf{Q}$. This follows from the fact that for these primes d , the field H_d is ramified only at d , so every field H_d is linearly disjoint from the compositum of the other fields H_d in the subfamily. As the given subfamily has asymptotically $B/(2 \log B)$ elements, we can treat the family of fields H_d with $d < B$ as being linearly independent at the cost of allowing for lower order (logarithmic) factors in our estimates. We can estimate the asymptotic size of the sum $\sum_{d < B} 1/h_d$ for squarefree $d < B$ to be a positive constant times $\sum_{0 < d < B} \frac{1}{\sqrt{d}} \approx \int_0^B \frac{dt}{\sqrt{t}} = 2\sqrt{B}$.

We find that for B tending to infinity, assumption 2 implies that the number of elements of prime norm N coming from $d < B$ is bounded from below by some universal constant times $\sqrt{B}/\log B$. By assumption 1, we expect to need about $\log N$ elements of norm N in step 5b. Thus, for prime values N tending to infinity, the size B_N of the minimal d solving problem 4.1 can be expected to be of size $\tilde{O}((\log N)^2)$. Note that B_N is small with respect to N , as required in our heuristical argument.

For the run time of the algorithm, we obtain $O((\log N)^{4+\epsilon})$ exactly as in [45]. The main term in the run time comes from computing $\tilde{O}((\log N)^2)$ values of $\sqrt{-d}$ modulo N , which each take time $O((\log N)^{2+\epsilon})$, and from proving (as in [5]) that the output is correct, i.e., that we have found α of norm N for which $N + 1 - \text{Tr}(\alpha)$ is indeed prime. $\square$

Numerical support. The table below shows the number of solutions $x, y \in \mathbf{Z}_{\geq 1}$ to the equation $x^2 + dy^2 = 4N$ for d ranging over the squarefree integers $d \in [1, B]$ for various B . For N we took the 5 primes following 10^{100} and 10^{200} . Note that the spacing of primes around 10^{100} and 10^{200} is in accordance with assumption 1.

$\downarrow N$ $B \rightarrow$	1000	4000	16000	64000
$p_1 = 10^{100} + 267$	30	57	125	232
$p_2 = 10^{100} + 949$	41	87	161	304
$p_3 = 10^{100} + 1243$	22	51	93	173
$p_4 = 10^{100} + 1293$	39	72	145	316
$p_5 = 10^{100} + 1983$	29	57	123	245
$q_1 = 10^{200} + 357$	46	91	190	354
$q_2 = 10^{200} + 627$	24	51	98	210
$q_3 = 10^{200} + 799$	24	47	90	184
$q_4 = 10^{200} + 1849$	47	81	170	376
$q_5 = 10^{200} + 2569$	73	140	275	532

We see that the growth rate is indeed roughly proportional to $c_N \sqrt{B}$, for some constant c_N : the numbers double if we quadruple B .

The data show that the size of N , when large with respect to B , is irrelevant: only the class of the primes over N in the class group of $\mathbf{Z}[\omega]$ is important, not the size of N .

Figure 1 below shows the number of solutions for p_2 and p_3 . Inspecting the data, we see that the growth rate is indeed close to $\sqrt{B}$. The fluctuation in the graphs is caused by the somewhat irregular behaviour of h_d . On a logarithmic scale, the graphs do look like straight lines with slope 1/2, see figure 2.




Figure 1




Figure 2

There are clear differences in the constants c_N for various N . These can be explained by looking at the contributions coming from composite d , which we could afford to neglect in our analysis, but which play an important role in practical situations. For

solvability of (4.3), it is clear that N has to be a square modulo all primes dividing d . For even d , we also have the condition $(\frac{N}{2}) = 1$. If we have $(\frac{N}{p}) = 1$ for many small primes p , there will most likely be more composite d yielding solutions to (4.3). The most striking difference in the table occurs for p_3 and q_5 . Looking at the Kronecker symbols $(\frac{d}{p})$ for the first eight primes $p \leq 20$, we only have $(\frac{d}{p}) = -1$ for $p = 3, 11$. For p_3 this occurs for $p = 2, 3, 5, 13, 17$. This explains why q_5 ‘outperforms’ p_3 . The differences in the constants c_N disappear if we only consider primes $d \equiv 3 \pmod{4}$ in our table. For p_3 we get 53 solutions up to $B = 64000$ in this case and for q_5 we get 50 solutions.

Whereas the number of generators of norm N found in step 5a for $d < B$ increases regularly, and is roughly proportional to $\sqrt{B}$, assumption 1 tells us that the number of times we have to test for primality in step 5b before we hit a prime number is $\log N$ *on average*. As a consequence, we expect that the minimal $d = d(N)$ solving problem 4.1 is of size $O((\log N)^{2+\epsilon})$, but not that $d(N)$ increases very regularly with N for prime values N . For instance, the primes p_1 and p_5 above have rather similar curves exhibiting the number of solutions found in step 5a, but the corresponding minimal discriminants 643 and 303267 are quite far apart: they are the smallest and largest values found for the p_i . However, the average value of d for the first 100 primes larger than 10^{100} and the first 100 primes larger than 10^{200} are $82170 \approx (\log(10^{100}))^{2.08}$ and $396030 \approx (\log(10^{200}))^{2.10}$, respectively. Their quotient 4.8 is not too far from the factor 4 we expect.

Elliptic curves of arbitrary order. The assumptions 1 and 2 at the beginning of the section also provide a heuristic run time analysis for arbitrary input N .

Assume first that N is squarefree, say $N = \prod_{i=1}^{\omega(N)} p_i$ with p_i prime. In step 3a, all d are discarded for which one of the primes p_i is inert in $\mathbf{Z}[\omega_d]$, so we will only be working in step 4 with those d for which none of the $\omega(N)$ Kronecker symbols $(\frac{d}{p_i})$ equals -1 . This can be a set of integers of density as small as $2^{-\omega(N)}$ inside the set of all squarefree integers, and in case N is in the zero-density subset of integers satisfying the equivalent inequalities

$$2^{\omega(N)} > (\log N)^2 \iff \omega(N) > \frac{2}{\log 2} \log \log N \doteq 2.88539 \log \log N$$

it is clear that we can no longer expect the integer d solving problem 4.1 to be of size at most $(\log N)^{2+\epsilon}$.

Despite the scarcity of suitable d for large values of $\omega(N)$, it is still the case that we expect the number of elements of norm N coming from $d < B$ to grow at least as fast as some universal constant times $\sqrt{B}/\log B$ if B tends to infinity. Indeed,

looking as before at the prime numbers $d \equiv 3 \pmod{4}$ (not dividing N) up to B , we see that there are ideals of norm N only for a fraction $2^{-\omega(N)}$ of them. However, for each d meeting the $\omega(N)$ quadratic conditions, the number of ideals I of norm N equals $2^{\omega(N)}$: we can take $I = \prod_{i=1}^{\omega(N)} \mathfrak{p}_i$, with $\mathfrak{p}_i$ one of the two primes dividing p_i in $\mathbf{Z}[\omega_d]$. This means that the growth with B of the number of ideals of norm N coming from $d < B$ is *independent* of the value of $\omega(N)$: with increasing $\omega(N)$ they occur for fewer d , but the decrease in contributing d is exactly compensated by the number of ideals provided by such d . Our expected number of elements of norm N coming from $d < B$ is therefore unchanged with respect to the case of primes N discussed before.

The problem with the asymptotic growth $\sqrt{B}/\log B$ of elements of norm N coming from a thin subset of $d < B$ is that B may have to be large to observe this growth rate: clearly the expected number $2^{-\omega(N)}B$ of contributing $d < B$ should not be too small. As we want to take $B \approx (\log N)^2$, we can only use our previous estimate for the expected size of the integer d solving problem 4.1 in the case $2^{\omega(N)} \ll (\log N)^2$. In the ‘opposite’ case $2^{\omega(N)} \gg (\log N)^2$, finding a *single* quadratic ring $\mathbf{Z}[\omega_d]$ in which all primes $p_i|N$ split completely is what the algorithm needs to achieve: there will be $2^{\omega(N)}$ ideals of norm N in this ring, of which assumption 2 tells us we can expect $2^{\omega(N)}/h_d \approx 2^{\omega(N)}/\sqrt{d}$ to be principal. As the smallest d satisfying the $\omega(N)$ quadratic conditions imposed by the p_i is expected to be of order of magnitude $2^{\omega(N)}$, we will find $2^{\omega(N)/2} \gg \log N$ elements α of norm N in $\mathbf{Z}[\omega_d]$. By assumption 1 this will lead to a prime element $1 - \alpha$.

THEOREM 4.6. *Under the heuristic assumptions 1 and 2, the integer d solving problem 4.1 is of size $\tilde{O}((\log N)^2 + 2^{\omega(N)})$, and our algorithm can be expected to find it in time $O(2^{\omega(N)}(\log N)^{4+\varepsilon})$.*

COROLLARY 4.7. *Under the heuristic assumptions 1 and 2, an elliptic curve of prescribed order N can be constructed in time $O(2^{\omega(N)}(\log N)^{4+\varepsilon})$.*

PROOF OF 4.7. Analogous to the proof of 4.5. □

PROOF OF 4.6. We saw that for squarefree N , the size of the integer d solving problem 4.1 is of size $\tilde{O}((\log N)^2)$ in case $2^{\omega(N)}$ is of smaller magnitude. If it is bigger, the term $2^{\omega(N)}$ becomes dominant and determines the expected size $\tilde{O}(2^{\omega(N)})$ of d .

If N is not squarefree, the algorithm has an increased number of possibilities to find ideals and elements of norm N for each value of d . Primes occurring to even exponents are no longer an obstruction if they are inert in $\mathbf{Z}[\omega_d]$: they get absorbed

in k_1 in step 3 and no longer occur in N_1 in step 4. Splitting primes occurring to higher exponents lead to square divisors $k_2^2|N_1$ in step 4, and to various ideals $(N_0, \omega - r)$ that can be tested for principality in step 4a. The extra ways to find elements of norm N is an advantage as it will lead to a smaller bound B_N for the minimal d solving problem 4.1. In particular, B_N will be of size $\tilde{O}((\log N)^2 + 2^{\omega(N)})$ for all N .

In order to estimate the run time of the algorithm, we observe that by assumption 1, step 4b will be executed about $\log N$ times until a probable prime norm is found, and a true primality proof taking expected time $O((\log N)^{4+\varepsilon})$ is needed. This is the dominant term in the time spent on step 4b. The number of times Cornacchia's algorithm is executed in step 4a to yield the $\log N$ generators going into step 4b is by assumption 2 no more than $O(\sqrt{B_N} \log B_N \log N)$, as the class numbers h_d for $d < B_N$ are no bigger than $\sqrt{B_N} \log B_N$. As Cornacchia's algorithm takes time $O((\log N)^{2+\varepsilon})$, we expect to spend time $O(\sqrt{B_N} \log B_N (\log N)^{3+\varepsilon})$ in step 4a.

In order to find the roots $(r \bmod N_1)$ of f in step 4, we first extract the square roots $\sqrt{-d}$ modulo each of the primes p_i that split in $\mathbf{Z}[\omega_d]$. This takes time at most $O(\omega(N)(\log N)^{2+\varepsilon})$. For each choice of square roots, there is a root $(r \bmod N_1)$ of f that can be found using the Chinese remainder theorem, in time $\omega(N)(\log N)^2$. Each time we apply the Chinese remainder theorem, we use the root $(r \bmod N_1)$ obtained in Cornacchia's algorithm in step 4a. The number of times we apply the Chinese remainder theorem is therefore bounded by the number of times $O(\sqrt{B_N} \log B_N \log N)$ we apply Cornacchia's algorithm. We find that the total time spent on finding roots $(r \bmod N_1)$ is no more than $O(\sqrt{B_N} \log B_N \omega(N)(\log N)^3)$. Taking all parts of step 4 together, the total time spent in this step becomes $O(\sqrt{B_N} \log B_N \omega(N)(\log N)^{3+\varepsilon})$. This is $O((\log N)^{4+\varepsilon})$ in case $2^{\omega(N)} \ll (\log N)^2$, and $O(2^{\omega(N)/2}(\omega(N))^2(\log N)^{4+\varepsilon})$ in general.

Outside step 4, no substantial computing is done, only some administration for the relatively small integer d , which takes values up to B_N . In cases where B_N is of order of magnitude $2^{\omega(N)} \gg (\log N)^2$, doing this administration is not negligible because of the large number of values taken on by d . Taking this into account, we find that the heuristic run time is bounded in all cases by $O(2^{\omega(N)}(\log N)^{4+\varepsilon})$. $\square$

Numerical support. Figure 3 below shows how the number of solutions $x, y \in \mathbf{Z}_{\geq 1}$ to the equation $x^2 + dy^2 = 4N$ for d ranging over all squarefree integers $d \in [1, B]$ varies with B for different number $\omega(N)$ of prime factors of N . The graphs are given for $N = N_1, N_2, N_3, N_{10}$, where N_k is the product of the first k primes larger than 10^{10} . We see that the graphs for N_1, N_2 and N_3 behave quite simi-

larly. This is what we expected if the number of solutions is independent of $\omega(N)$. The graph for N_{10} appears to be quite different from the others, and this is because $2^{\omega(N_{10})} = 2^{10} = 1024$ is here of the same order of magnitude as the values of B in the graph. There are here fewer d for which we have a solution to $x^2 + dy^2 = 4N_{10}$, but if we do have a solution, we immediately get many. For instance, the first ‘jump’ in the graph occurs for the prime value $d = 1949$ and we get 28 solutions for this d . This is in nice accordance with the heuristics, which tell us to expect the first solutions to occur for $d \approx 2^{10} = 1024$, and to be about $2^5 = 32$ in number.




Figure 3




Figure 4

The irregularity of the graph for N_{10} disappears if we look at values of B that are large in comparison to $2^{\omega(N_{10})}$. Figure 4 shows the graph for N_{10} for B up to 10^7 . It is now similar in nature to that of N_1 , and exhibits the familiar $\sqrt{B}$ -profile.

The graph in figure 5 below illustrates the dependence on the number of square divisors of N . It shows the number of solutions for N_1 , $3^2 \cdot N_1$, $3^2 \cdot 5^2 \cdot N_1$ and $3^2 \cdot 5^2 \cdot 7^2 \cdot N_1$. If N has square divisors, we potentially test the principality of more ideals in step 4 of our algorithm, so we expect to obtain more solutions. Replacing N_1 for example by $3^2 \cdot N_1$, we expect to get on average a double amount of solutions for $d \equiv 1 \pmod{3}$. The gain is a constant factor > 1 that increases with the number of square divisors of N .




Figure 5

4.4 Examples and practical considerations

The description of the algorithm in section 4.2 is intended to facilitate the run time estimate in section 4.3. It does not address practical issues that are important in computing large examples. In this section, we explain how we find curves of large order N , where N is either prime or equal to a power of 10.

Elliptic curves of large prime order. From the description of the algorithm we gave in the previous section, and more in particular its relation to ECPP, it is clear that one should be able to construct a curve having a large prime number N of points in all cases where ECPP, as described in [45], can prove primality of a number of the same size. To do so, it makes sense to apply an idea attributed to J. Shallit in [45] to speed up the computation. This idea starts from the observation that for large prime numbers N , the algorithm spends a lot of time evaluating $(\sqrt{-d} \bmod N)$ for all squarefree d up to $B_N \approx (\log N)^2$ having $(\frac{-d}{N}) = 1$. We noticed already in the previous section that if the equation

$$x^2 + dy^2 = 4N$$

admits integral solutions, then N is a square modulo all primes dividing the discriminant $D = \text{disc}(\mathbf{Q}(\sqrt{-d}))$. It reflects the fact that if N splits completely in the Hilbert class field H_d of $K = \mathbf{Q}(\sqrt{-d})$, then it certainly splits completely in the genus field $G_d \subset H_d$ of K . As G_d is obtained by adjoining to K the square roots of $p^* = (-1)^{(p-1)/2}p$ for all odd prime divisors $p|d$, we have $(\frac{p^*}{N}) = (\frac{N}{p}) = 1$ in this case.

Once we know that those d providing solutions are essentially products of primes having the right quadratic character with respect to N , the idea suggests itself to look at those d only that are constructed as products of such primes. Creating d from a ‘basis’ of primes p with $(\frac{p^*}{N}) = 1$ allows us to compute $\sqrt{p^*} \bmod N$ for such p , and store the values in a list. For $p = 2$, one uses the square roots of -1 , 2 and -2 that can be extracted modulo N . For each d constructed from our basis of primes, $\sqrt{-d} \bmod N$ can be obtained by multiplying the square roots of primes modulo N we stored. Considering only products of two primes from our basis allows us to reduce the number of square root extractions modulo N from $O((\log N)^2)$ to $O(\log N)$, at the expense of extra multiplications modulo N and an increased storage requirement. In practice, we consider d with at most 3 prime divisors. One thing we lose in this approach is the guarantee that we really find the *smallest* solution d to problem 4.1.

EXAMPLE. Take $N = \text{nextprime}(10^{2004}) = 10^{2004} + 4863$, the exponent 2004 being the year we found our method. For this N , we have $\log(N) \doteq 4614.3$ and $(\log(N))^2 \doteq 2.13 \cdot 10^7$. There are 324 primes p less than 5000 with $(\frac{p^*}{N}) = 1$, and we compute and store $\sqrt{2} \bmod N$ and all square roots $\sqrt{p^*} \bmod N$. We now have $\binom{325}{3} = 5668650$ squarefree values of d at our disposal having up to 3 prime divisors from our base, and we know N to split completely in all genus fields G_d .

The 104415-th value of d we tried was $d = 59 \cdot 523 \cdot 2579 = 79580203$. For this value of d , we found a solution

$$x = 1885782 \dots 693127$$

to $x^2 + dy^2 = 4N$ for which

$$p = N + 1 - x = 999999 \dots 99999811421 \dots 8311737$$

is a 2004-digit prime. In each case, the dots represent 990 digits that we omitted.

The class polynomial P_{-d} has degree 1536 and coefficients up to 41984 digits. Modulo p , the polynomial P_{-d} splits completely. Taking j to be the smallest positive integer satisfying $P_{-d}(j) \equiv 0 \bmod p$ we put $a = \frac{27j}{4(1728-j)} \in \mathbf{F}_p$. Then the curve given by

$$E_a : Y^2 = X^3 + aX - a$$

has CM by $\mathcal{O}_{-d}$. As the point $(1, 1) \in E_a(\mathbf{F}_p)$ does not have order N , the quadratic twist $E'_a : Y^2 = X^3 + 9aX - 27a$ of E_a has N points. This can be verified by picking a random point $P \in E'_a(\mathbf{F}_p)$ and checking that we have $N \cdot P = 0$.

The value of d we find here is in fact the smallest d solving problem 4.1 for our N . Our algorithm did 565 primality tests before we found the solution above. Finding d and p took about 10 minutes on our standard, 32-bit 2.8 GHz, PC, and another 3 hours were needed to find and factor P_{-d} . Once we find j , the final result is almost immediate. If we trust the input value N as being a true prime number, there is no need to prove that p is prime. As in ECPP, this follows from the fact that E'_a has a non-trivial point that is annihilated by N .

Elliptic curves of 10-power order. We indicated in our analysis in section 4.3 that for input values of N having a large number of square divisors, the integer d solving problem 4.1 will be much smaller than the upper bound for squarefree N occurring in theorem 4.6. This can be illustrated by looking at the values $N = 10^k$ for $k \geq 1$, which have $\log N \approx 2.3k$. As none of the prime divisors 2 and 5 of N is inert in the field $\mathbf{Q}(i)$ and the prime 5 is split, there are already many solutions

to the norm equation $x^2 + y^2 = N$ for the very first value $d = 1$. In fact, as we have $h_d = 1$ there is no need for a Cornacchia algorithm, and the elements of norm $N = 2^k 5^k$ in $\mathbf{Z}[i]$ are the $4k + 4$ elements $\alpha_{s,t} = i^s(1+i)^k(2+i)^t(2-i)^{k-t}$ with $s \in \{0, 1, 2, 3\}$ and $t \in \{0, 1, \dots, k\}$. Up to conjugacy, we have about $2k = .87 \log N$ elements, so we expect that for a positive fraction of all k -values, $d = 1$ gives rise to a prime p and a twist E of the curve $Y^2 = X^3 + X$ having exactly 10^k points over $\mathbf{F}_p$. As the graph below indicates, this fraction appears to be close to 0.92.




Figure 6

EXAMPLE. Take $k = 2004$. We find that for $(s, t) = (2, 499), (0, 527), (0, 671)$, the element $\alpha_{s,t} = i^s(1+i)^{2004}(2+i)^t(2-i)^{2004-t}$ of norm 10^{2004} has the property that $p = N_{\mathbf{Q}(i)/\mathbf{Q}}(1 - \alpha_{s,t})$ is prime. The curve $Y^2 = X^3 + X$ having $j = 0$ and CM by $\mathbf{Z}[i]$ has 4 twists over $\mathbf{F}_p$ for each of these p , but in all cases $Y^2 = X^3 + X$ is the curve having 10^{2004} points. This follows from a result in [60] going back to Gauß. It says that if we choose the prime element $\pi = a + bi$ dividing a prime $p \equiv 1 \pmod{4}$ in $\mathbf{Z}[i]$ to satisfy $\pi \equiv 1 \pmod{(1+i)^3}$, then the curve $Y^2 = X^3 + X$ has exactly $p + 1 - \left(\frac{-1}{\pi}\right)_4(\pi + \bar{\pi}) = p + 1 - 2i^{1-a}a$ points over $\mathbf{F}_p$. In our case, $\pi = 1 - \alpha_{s,t}$ and a are congruent to 1 modulo $(1+i)^{2004} = -2^{1002}$, so we already know that $Y^2 = X^3 + X$ is the right curve before actually computing p .

For the purpose of constructing curves having $N = 10^k$ points, there are small values of d that conjecturally work for almost all values of k , not just for a positive fraction of them. These d have the property that 2 and 5 *both* split completely in $\mathbf{Q}(\sqrt{-d})$, i.e., they satisfy $d \equiv 31, 39 \pmod{40}$. For such d , the number of ideals of norm N grows quadratically in k , and hence in $\log N$. If we fix d , and hence h_d , the number of elements of norm N in $\mathbf{Q}(\sqrt{-d})$ will also grow quadratically in $\log N$, and our assumption 2 implies that such d will work for all but finitely many k .

EXAMPLE. Let ρ be a zero of $X^3 + X + 1$. Then ρ is the value of the Weber function $f(z) = \zeta_{48}^{-1} \cdot \frac{\eta(\frac{z+1}{2})}{\eta(z)}$ at $-23 - 1/\omega_{31}$, and a generator of the Hilbert class field of $\mathbf{Q}(\sqrt{-31})$. An elliptic curve $E_j/\mathbf{Q}(\rho)$ having j -invariant $j = (\rho^{24} - 16)^3/\rho^{24}$ has endomorphism ring $\mathbf{Z}[\omega_{31}]$. We may take

$$E_j : Y^2 = X^3 + 3j(1728 - j)X + 2j(1728 - j)^2$$

which has good reduction outside $2, 3, 11, 17, 23, 31$. For all values $1 \leq k \leq 1000$ except $k = 1, 2$, there exist primes of the form

$$p = x^2 + 31y^2 = 10^k - 1 + 2x. \quad (4.5)$$

To find them, we write $(\omega_{31} + 1) = \mathfrak{p}_2\mathfrak{p}_5$ and note that a $\mathbf{Z}[\omega_{31}]$ -ideal

$$\mathfrak{p}_2^s \cdot \bar{\mathfrak{p}}_2^{k-s} \cdot \mathfrak{p}_5^t \cdot \bar{\mathfrak{p}}_5^{k-t}$$

of norm 10^k is principal if and only if we have $s \equiv t \pmod{3}$. We use Cornacchia's algorithm to find the generators α for the principal ideals and test whether $N(1 - \alpha)$ is prime. For primes satisfying (4.5), either the reduction $\bar{E}_j/\mathbf{F}_p$ of E_j over a prime over p in $\mathbf{Q}(\rho)$ or its quadratic twist has exactly 10^k rational points over $\mathbf{F}_p$. It is likely that $k = 1, 2$ are the *only* values of k for which no prime p of the form (4.5) exists, but this is probably very hard to prove.

A non-archimedean algorithm

5.1 Finding a small splitting prime

The main task of the CM algorithm from chapters 3 and 4 is to compute the Hilbert class polynomial P_Δ for a suitable discriminant $\Delta < 0$. The classical approach proceeds by evaluating the modular function $j : \mathbf{H} \rightarrow \mathbf{C}$ in points $\tau \in \mathbf{H}$ corresponding to the ideal classes of $\mathcal{O}_\Delta$, as explained in section 3.3. This has the disadvantage that rounding errors may occur in expanding the product

$$P_\Delta = \prod_{[a,b,c] \in \mathcal{F}_\Delta^+} \left(X - j\left(\frac{-b + \sqrt{\Delta}}{2a}\right) \right) \in \mathbf{Z}[X].$$

We can avoid this problem by working in a p -adic setting, where the prime p that we will use for this purpose in the present chapter bears no relation to the prime p occurring in chapters 2–4. The approach presented in this chapter is based on the work of Couveignes and Henocq [12]. For the problem of constructing a curve with prescribed order N , we are mainly interested in *fundamental* discriminants Δ . The computation of the Hilbert class polynomial P_Δ is however interesting in its own right. Hence, we will present an algorithm that computes P_Δ for *any* discriminant $\Delta < -4$, fundamental or not. For completeness sake, we note that we have $P_{-3} = X$ and $P_{-4} = X - 1728$.

We fix a discriminant $\Delta < -4$, and let $\mathcal{O} = \mathcal{O}_\Delta$ be the order of discriminant Δ . For any prime p , we can embed the ring class field $H_\mathcal{O}$ corresponding to the order $\mathcal{O}$ in an algebraic closure $\bar{\mathbf{Q}}_p$ of the non-archimedean field $\mathbf{Q}_p$. All the computations within the algorithm will be done in the subfield $H_\mathcal{O}$ of $\bar{\mathbf{Q}}_p$. The first we do is finding a ‘small’ prime $p \neq 2, 3$ with the property that $H_\mathcal{O}$ can be embedded in $\mathbf{Q}_p$, i.e., a small prime that splits completely in $H_\mathcal{O}$.

By lemma 3.3, a prime p splits completely in $H_\mathcal{O}$ if and only if (p) splits into two principal ideals in $\mathcal{O}$. A prime p splits into two principal ideal ideals in $\mathcal{O}$ if and

only if we can solve the equation

$$4p = t^2 - u^2\Delta \quad (5.1)$$

in integers t, u . We see that we have a lower bound $p > |\Delta|/4$.

In order to find a prime p that splits completely in $H_{\mathcal{O}}$ we can first take $u = 1$ in equation 5.1. We let t range over $1, 2, \dots, B(\Delta)$ and test whether $\frac{t^2 - u^2\Delta}{4}$ is prime. Here $B(\Delta)$ is some upper bound, depending on Δ . If we do not find a solution to equation 5.1 with $u = 1$, we try $u = 2, 3, \dots$, etc. However, for $\Delta \equiv 0 \pmod{4}$, we take t even to ensure that $\frac{t^2 - u^2\Delta}{4}$ is an integer. For $\Delta \equiv 1 \pmod{4}$, the integers t and u should have the same parity.

We give a simple heuristic showing that we may expect to find a solution to equation 5.1 with p of size $O(|\Delta| \log |\Delta|)$. The integers $\frac{t^2 - u^2\Delta}{4}$ are more or less random integers of size $|\Delta|$. By the prime number theorem, we expect that one out of every $\log |\Delta|$ integers of size $|\Delta|$ is prime. Hence, if we take $B(\Delta) = \log \Delta$, we expect to find a prime p for a small value of u .

This is more than what can be rigorously proved. If we assume GRH however, the following lemma tells us that we can indeed find a prime p of size $\tilde{O}(|\Delta|)$ that splits completely in $H_{\mathcal{O}}$.

LEMMA 5.1. *If GRH holds true, there exists an effectively computable constant $c \in \mathbf{R}_{>0}$ such that for every $\Delta < 0$ there exists a prime $p \in \mathbf{Z}$ that splits completely in $H_{\mathcal{O}}$ and that satisfies*

$$p \leq c \cdot |\Delta| (\log |\Delta|)^4.$$

PROOF. We want to apply the effective Chebotarev density theorem [35], which requires the assumption of GRH. It states that there is an effectively computable constant $c \in \mathbf{R}_{>0}$ such that for every $\Delta < 0$ there exists a prime p that splits completely in $H_{\mathcal{O}}$ and that satisfies

$$p \leq c \cdot (\log |\text{disc}(H_{\mathcal{O}}/\mathbf{Q})|)^2,$$

where $\text{disc}(H_{\mathcal{O}}/\mathbf{Q})$ is the discriminant of $H_{\mathcal{O}}/\mathbf{Q}$. The lemma follows if we can bound this discriminant appropriately.

For fundamental discriminants we have $\text{disc}(H_{\mathcal{O}}/\mathbf{Q}) = |\Delta|^{h(\Delta)}$, but for non-fundamental discriminants the situation is more complicated. We will compute the discriminant $\text{disc}(H_{\mathcal{O}}/\mathbf{Q})$ via the relation

$$\text{disc}(H_{\mathcal{O}}/\mathbf{Q}) = N_{K/\mathbf{Q}}(\text{disc}(H_{\mathcal{O}}/K)) \cdot \text{disc}(K/\mathbf{Q})^{[H_{\mathcal{O}}:K]},$$

where we write $K = \mathbf{Q}(\sqrt{\Delta})$. Write $\Delta = f^2 D$ with $D = \text{disc}(\mathbf{Q}(\sqrt{\Delta}))$. Then $H_{\mathcal{O}}/K$ is an abelian extension of conductor dividing f and degree $h(\Delta)$. From the conductor discriminant formula [46, Theorem 11.9] we see that the K -ideal $\text{disc}(H_{\mathcal{O}}/K)$ is a divisor of $f^{h(\Delta)}$. We have $N_{K/\mathbf{Q}}(f^{h(\Delta)}) = f^{2h(\Delta)}$ and we estimate

$$\text{disc}(H_{\mathcal{O}}/\mathbf{Q}) \leq f^{2h(\Delta)} \cdot |D|^{h(\Delta)} = |\Delta|^{h(\Delta)}.$$

Using the upper bound $h(\Delta) \leq \sqrt{|\Delta|} \log |\Delta|$ from [40, Section 2], we conclude

$$\text{disc}(H_{\mathcal{O}}/\mathbf{Q}) \leq |\Delta|^{\sqrt{|\Delta|} \log |\Delta|}. \quad \square$$

This lemma tells us that, if GRH is true, there is a solution (t, u, p) to equation 5.1 with $t = O(\sqrt{|\Delta|} \cdot (\log |\Delta|)^2)$ and $u = O((\log |\Delta|)^2)$. Note that p has size $O(|\Delta|(\log |\Delta|)^4)$.

Let $\text{Ell}_{\Delta}(\mathbf{Q}_p)$ be the set of j -invariants of elliptic curves over $\mathbf{Q}_p$ with endomorphism ring $\mathcal{O} = \mathcal{O}_{\Delta}$. From the discussion in section 3.2, we see that this is a *finite* set of cardinality $h(\Delta)$. We can compute P_{Δ} as

$$P_{\Delta} = \prod_{j(E) \in \text{Ell}_{\Delta}(\mathbf{Q}_p)} (X - j(E)) \in \mathbf{Z}[X],$$

and in the remainder of this chapter we explain how to compute the finite set $\text{Ell}_{\Delta}(\mathbf{Q}_p)$.

5.2 The canonical lift

Let $p \geq 5$ be a prime that splits completely in the ring class field corresponding to the order $\mathcal{O} = \mathcal{O}_{\Delta}$. Let $\text{Ell}_{\Delta}(\mathbf{F}_p)$ be the set of j -invariants of elliptic curves over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$. The set $\text{Ell}_{\Delta}(\mathbf{F}_p)$ is a finite set of cardinality $h(\Delta)$. Since an element $j \in \text{Ell}_{\Delta}(\mathbf{Q}_p)$ is integral, we can consider its image under the reduction map $\mathbf{Z}_p \rightarrow \mathbf{F}_p$, and we obtain a natural bijection

$$\mu : \text{Ell}_{\Delta}(\mathbf{Q}_p) \rightarrow \text{Ell}_{\Delta}(\mathbf{F}_p).$$

For an ordinary elliptic curve $\bar{E}/\mathbf{F}_p$, the value $\mu^{-1}(j(\bar{E})) \in \mathbf{Q}_p$ is uniquely determined by $\bar{E}$. The following lemma gives us a stronger result.

LEMMA 5.2. *Let $p \geq 5$ be prime, and let $\bar{E}$ be an ordinary elliptic curve over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$. Then there exists an elliptic curve $\tilde{E}/\mathbf{Q}_p$ with endomorphism ring $\mathcal{O}$ that reduces to $\bar{E}/\mathbf{F}_p$. The curve $\tilde{E}$ is unique up to isomorphism over $\mathbf{Q}_p$.*

PROOF. The prime p splits completely in the ring class field of the order $\mathcal{O}$. The existence of $\tilde{E}/\mathbf{Q}_p$ follows immediately from the fact that $\mu : \text{Ell}_\Delta(\mathbf{Q}_p) \rightarrow \text{Ell}_\Delta(\mathbf{F}_p)$ is bijective. The j -invariant $j(\tilde{E}) \in \mathbf{Q}_p$ determines the curve $\tilde{E}$ up to twisting with elements of $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2}$. For $p > 2$, we have $\mathbf{Q}_p^*/\mathbf{Q}_p^{*2} \cong \mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/2\mathbf{Z}$ and we take $\{1, v, p, pv\}$ as set of representatives. Here, $v \in \mathbf{Z}_p^*$ is any element that reduces to a non-square modulo p . If we twist $\tilde{E}$ by p or by pv , the twist has bad reduction modulo p . If we twist $\tilde{E}$ by v , the twist $\tilde{E}'/\mathbf{Q}_p$ reduces to the quadratic twist of $\bar{E}/\mathbf{F}_p$. $\square$

The curve $\tilde{E}/\mathbf{Q}_p$ in lemma 5.2 is called the *canonical lift* of $\bar{E}/\mathbf{F}_p$. One of the main steps in the non-archimedean algorithm will be to compute, on input of an ordinary curve $\bar{E}/\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$, the j -invariant $\mu^{-1}(j(\bar{E})) \in \mathbf{Q}_p$ of its canonical lift. The explicit description of the algorithm in theorem 5.3 below is given in section 5.7.

THEOREM 5.3. *There exists an algorithm which has as input*

- ◊ a prime $p \geq 5$
- ◊ an ordinary j -invariant $j \in \mathbf{F}_p$
- ◊ a positive integer k

and as output the canonical lift $\mu^{-1}(j) \in \mathbf{Q}_p$ of $j \in \mathbf{F}_p$ in k digits accuracy. If GRH holds true, the expected run time of this algorithm is for every $\varepsilon > 0$ bounded by

$$c_\varepsilon \left(\exp((\log p)^{1/2+\varepsilon}) \times \log k \right)^4 \times k.$$

for some effectively computable constant $c_\varepsilon > 0$.

In the remainder of this section we explain how we can find, on input of a discriminant $\Delta < -4$, an ordinary curve $\bar{E}/\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$.

► Fundamental discriminants

Let $\Delta < -4$ be a fundamental discriminant. We apply the method explained in section 5.1 to find a solution (t, u, p) to the equation

$$t^2 - u^2\Delta = 4p$$

with p prime. In practice, we can find a solution with $u = 1$ for $\Delta \not\equiv 1 \pmod{8}$. For $\Delta \equiv 1 \pmod{8}$, this is never possible. The reason is that t needs to have the same

parity as Δ and is consequently odd. Hence, $t^2 - \Delta$ is divisible by 8. This is only possible for $\Delta = -7$, where we may take $p = 2$. We assumed $p > 3$ however. For $\Delta \equiv 1 \pmod 8$ we can in practice always take $u = 2$.

The run time of any algorithm computing P_Δ will be at least $O(|\Delta|)$. Since we have $p = O(|\Delta|^{1+o(1)})$, we can afford to apply the naïve algorithm from chapter 2 to find an ordinary curve $\bar{E}/\mathbf{F}_p$ whose Frobenius morphism F_p has trace $t \neq 0$. We have

$$u = [\mathcal{O} : \mathbf{Z}[F_p]],$$

and we see that for $u = 1$, the curve $\bar{E}$ must have endomorphism ring $\mathcal{O}$. For $u = 2$, the ring $\mathbf{Z}[F_p]$ is properly contained in $\mathcal{O}$ with index 2. We either have $\text{End}(\bar{E}) = \mathcal{O}$ or $\text{End}(\bar{E}) = \mathbf{Z}[F_p] \subsetneq \mathcal{O}$. We explain how we can compute the endomorphism ring $\text{End}(\bar{E})$ in this case, by looking at the $\mathbf{F}_p$ -rational 2-torsion of $\bar{E}$.

Note that we have

$$\bar{E}[2] \subset \bar{E}[F_p - 1] = \bar{E}(\mathbf{F}_p) \iff 2 \mid (F_p - 1) \in \text{End}(\bar{E}) \iff [\text{End}(\bar{E}) : \mathbf{Z}[F_p]] = 2.$$

Hence, we have $\text{End}(\bar{E}) = \mathcal{O}$ if and only if the 2-torsion $\bar{E}[2]$ is $\mathbf{F}_p$ -rational. If this is not the case, then $F_p \in 1 + 2\mathcal{O}$ has even trace t , so $\#\bar{E}(\mathbf{F}_p) = p + 1 - t$ is even and $\bar{E}$ has a single $\mathbf{F}_p$ -rational 2-torsion point P . The endomorphism ring of the isogenous curve $\bar{E}/\langle P \rangle$ contains the ‘same’ subring $\mathbf{Z}[F_p]$ as $\text{End}(\bar{E})$, but as there are no rational 2-isogenies to curves with endomorphism ring $\mathbf{Z}[F_p]$ by lemma 5.11(ii), the curve $\bar{E}/\langle P \rangle$ must have endomorphism ring $\mathcal{O}$.

Algorithmically, we compute the 2-torsion point $P = (p_1, p_2) \in \bar{E}(\mathbf{F}_p)$ for the curve $\bar{E} : Y^2 = X^3 + aX - a$ that we get from the naïve algorithm. We make a change of variables $(X, Y) \rightarrow (X - p_1, Y - p_2)$ to move the point P to $(0, 0)$. Let $Y^2 = X(X^2 + cX + d)$ be the new Weierstraß equation. A Weierstraß equation for $\bar{E}/\langle P \rangle$ is now given by $Y^2 = X(X^2 - 2cX + (c^2 - 4d))$, cf. [9, Chapter 14].

► Arbitrary discriminants and analysis

Let $\Delta < -4$ be an arbitrary discriminant. If GRH holds true, there exists a solution (t, u, p) to equation 5.1 with $t = O(|\Delta|^{1/2+o(1)})$, $u = O(|\Delta|^{o(1)})$ and $p = O(|\Delta|^{1+o(1)})$. Take any such solution. Applying the naïve algorithm from chapter 2 yields an ordinary curve $\bar{E}/\mathbf{F}_p$ with trace of Frobenius $t \neq 0$. As $u^2\Delta$ need not be fundamental, this does not ensure $\text{End}(\bar{E}) = \mathcal{O}$. We can apply Kohel’s algorithm [34] to compute the exact endomorphism ring of $\bar{E}$. The run time of Kohel’s algorithm is $O(p^{1/3+o(1)})$. We see that we can find a curve with trace of Frobenius t and

test whether it has endomorphism ring $\mathcal{O}$ in time $O(p^{1/2+o(1)})$. This observation leads to the following algorithm. If GRH holds true, the prime p we find in step 1 is of size $O(|\Delta|^{1+o(1)})$.

Algorithm. (*Endomorphism ring algorithm*)

Input: a negative discriminant $\Delta < -4$. Output: a prime p and an ordinary elliptic curve $\bar{E}/\mathbf{F}_p$ with $\text{End}(\bar{E}) = \mathcal{O}_\Delta$.

1. Find a solution (t, u, p) to equation 5.1, using the method described in section 5.1, with p prime and with $p \nmid t$.
2. Apply the naïve algorithm from chapter 2 to find an elliptic curve $\bar{E}/\mathbf{F}_p$ with trace of Frobenius t .
3. Test $\text{End}(\bar{E}) = \mathcal{O}_\Delta$ using Kohel's algorithm. If so, return p and $\bar{E}$. Otherwise, go to step 2.

To analyse the run time, it remains to give an estimate for the number of times we have to do step 2. The probability that an elliptic curve with trace of Frobenius t has endomorphism ring $\mathcal{O}$ equals

$$n = \frac{\#\{E/\mathbf{F}_p \mid \text{End}(E) = \mathcal{O}_\Delta\}/\cong}{\#\{E/\mathbf{F}_p \mid \text{End}(E) \supseteq \mathcal{O}_{u^2\Delta}\}/\cong}.$$

Recall that the $\#'$ -symbol means that we count every isomorphism class $[E]$ with weight $1/\#(\text{Aut}_{\mathbf{F}_p} E) = 1/2$.

From chapter 2 we know $n = \frac{h'(\Delta)}{H'(u^2\Delta)}$, where h' and H' denote the weighted class number and the weighted Kronecker class number respectively. From the formulas given in [39, Section 1.6] for the Kronecker class number we derive

$$n \geq \left(\frac{\varphi(f)}{f}\right)^2 \cdot \frac{1}{u},$$

where φ denotes the Euler- φ function and f is the index $[\mathcal{O}_{\max} : \mathcal{O}_{u^2\Delta}]$. Theorem 328 in [28] gives that $\liminf_{f \rightarrow \infty} \frac{\varphi(f) \log \log f}{f}$ is finite. Combining this with the estimate $u = O((\log |\Delta|)^2)$, we conclude

$$n \geq c_\varepsilon \cdot \frac{1}{(\log |\Delta|)^{2+\varepsilon}}$$

for some effectively computable constant c_ε , depending on ε . This lower bound for the probability n yields the following run time for the endomorphism ring algorithm.

Run time. *If GRH holds true, the endomorphism ring algorithm has expected run time $O(|\Delta|^{1/2+\varepsilon})$ for every $\varepsilon > 0$.*

5.3 Modular curves

The description of the algorithm to compute the canonical lift makes use of modular curves. We briefly recall the basic notions here. For more information, see e.g. [56] or [15].

Define the subgroup

$$\Gamma(N) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z}) \mid a \equiv d \equiv 1 \pmod{N}, b \equiv c \equiv 0 \pmod{N} \right\}$$

of $\mathrm{SL}_2(\mathbf{Z})$ for integers $N \in \mathbf{Z}_{>0}$. A subgroup $\Gamma \subseteq \mathrm{SL}_2(\mathbf{Z})$ is called a *congruence subgroup* if it contains $\Gamma(N)$ for some N . The group $\mathrm{SL}_2(\mathbf{Z})$ acts on the upper half plane $\mathbf{H}$ by linear fractional transformations as in section 3.3.

Let now $\Gamma \subset \mathrm{SL}_2(\mathbf{Z})$ be a congruence subgroup. The quotient $Y_\Gamma = \Gamma \backslash \mathbf{H}$ has the structure of a Riemann surface. See e.g. [56, Section 1.5] for the definition of the complex structure. Write $\bar{\mathbf{H}} = \mathbf{H} \cup \mathbf{P}^1(\mathbf{Q})$ and $X_\Gamma = \Gamma \backslash \bar{\mathbf{H}}$. We give X_Γ the structure of a compact Riemann surface by completing Y_Γ with the finitely many Γ -orbits of $\mathbf{P}^1(\mathbf{Q})$ called the *cusps* of X_Γ .

Since X_Γ is a compact Riemann surface, the field $\mathbf{C}(X_\Gamma)$ of meromorphic functions on X_Γ is a finitely generated extension of $\mathbf{C}$ of transcendence degree 1. There is an anti-equivalence between the category of finitely generated field extensions of $\mathbf{C}$ of transcendence degree 1 and the category of smooth complete curves over $\mathbf{C}$ with surjective morphisms. Via this anti-equivalence, we view X_Γ as a smooth complete curve over $\mathbf{C}$.

For $\Gamma = \mathrm{SL}_2(\mathbf{Z})$, we write $Y_\Gamma = Y(1)$. Then $Y(1)$ is an affine curve with completion $X(1) = X_\Gamma$. Points on $Y(1)$ have a natural interpretation as isomorphism classes of complex elliptic curves, which we make explicit by viewing $Y(1)$ as the quotient $\mathrm{SL}_2(\mathbf{Z}) \backslash \mathbf{H}$. We map the $\mathrm{SL}_2(\mathbf{Z})$ -orbit of $\tau \in \mathbf{H}$ to the isomorphism class of the complex elliptic curve $E_\tau = \mathbf{C}/\langle 1, \tau \rangle$. The j -function $\mathbf{H} \rightarrow \mathbf{C}$ gives a map $j : Y(1) \rightarrow \mathbf{C}$, and the extension of this map to $X(1)$ gives an isomorphism $X(1) \xrightarrow{\sim} \mathbf{P}_{\mathbf{C}}^1$.

We now focus on the case $\Gamma = \Gamma_0(N)$, with

$$\Gamma_0(N) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z}) \mid c \equiv 0 \pmod{N} \right\} \supset \Gamma(N),$$

and write $Y_\Gamma = Y_0(N)$. Then $Y_0(N)$ is an affine curve with completion $X_0(N) = X_\Gamma$. The function $j_N : \mathbf{H} \rightarrow \mathbf{C}$ given by $j_N(\tau) = j(N\tau)$ induces a map $j_N : X_0(N) \rightarrow \mathbf{P}_{\mathbf{C}}^1$. The function field of $X_0(N)$ is $\mathbf{C}(j, j_N)$.

Points on $Y_0(N)$ have a natural moduli interpretation as isomorphism classes of pairs (E, G) , where E is a complex elliptic curve and $G \subset E[N]$ is a cyclic subgroup

of order N . The notion of isomorphism is the following: two pairs $(E, G), (E', G')$ are isomorphic if and only if there exists an isomorphism of elliptic curves $\varphi : E \rightarrow E'$ with $\varphi(G) = G'$. We make this moduli interpretation explicit by viewing $Y_0(N)$ as the quotient $\Gamma_0(N) \backslash \mathbf{H}$. The $\Gamma_0(N)$ -orbit of $\tau \in \mathbf{H}$ is then mapped to the isomorphism class of the pair

$$(E_\tau, \langle 1/N \rangle),$$

consisting of the complex elliptic curve $E_\tau = \mathbf{C}/\langle 1, \tau \rangle$ and the cyclic subgroup generated by $1/N \in E_\tau[N]$. This moduli interpretation plays an important role in the sequel and we will often denote a point of $Y_0(N)$ by a pair (E, G) . We have $j(E_\tau) = j(\tau)$ and $j(E_\tau/\langle \frac{1}{N} \rangle) = j_N(\tau) = j(N\tau)$.

The identity map $\mathbf{H} \rightarrow \mathbf{H}$ induces a map $f : Y_0(N) \rightarrow Y(1)$:

$$\begin{array}{ccc} \mathbf{H} & \longrightarrow & Y_0(N) \\ \text{id} \downarrow & & \downarrow f \\ \mathbf{H} & \longrightarrow & Y(1) \end{array}$$

that can be extended to $f : X_0(N) \rightarrow X(1)$. In terms of the moduli interpretation, we have $f(E, G) = E$, i.e., f is just the ‘forgetful map’. Via this map f , the curve $X_0(N)$ is a cover of the curve $X(1)$ of degree $\psi(N) = N \prod_{p|N} (1 + \frac{1}{p}) = [\text{SL}_2(\mathbf{Z}) : \Gamma_0(N)]$. The j -invariant gives an isomorphism $X(1) \xrightarrow{\sim} \mathbf{P}_{\mathbf{C}}^1$, and we may view $X_0(N)$ in a natural way as a cover of $\mathbf{P}_{\mathbf{C}}^1$. This cover $X_0(N)/\mathbf{P}_{\mathbf{C}}^1$ is ramified only above $j = 0, 1728, \infty$.

The functions j and j_N are related by a polynomial relation $\Phi_N(j, j_N) = 0$, with $\Phi_N \in \mathbf{Z}[X, Y]$. There is a smooth complete curve $X_0(N)_{\mathbf{Q}}$ over $\mathbf{Q}$ with function field $\mathbf{Q}(j, j_N)$. For an extension $K/\mathbf{Q}$, the set $Y_0(N)_{\mathbf{Q}}(K)$ of K -valued points of $Y_0(N)_{\mathbf{Q}}$ consists of the $\bar{K}$ -isomorphism classes of pairs (E, G) , where E is an elliptic curve over K and $G \subset E[\bar{K}]$ is a cyclic subgroup of order N of $E[N]$ that is defined over K . We will often write $X_0(N)$ instead of $X_0(N)_{\mathbf{Q}}$ to denote the modular curve $X_0(N)$ over $\mathbf{Q}$.

We also want to consider isogenies between elliptic curves over finite fields. It would be most natural to properly define what the reduction of $X_0(N)$ is modulo a prime. Unfortunately, this algebraic geometric notion is not so easy. One can prove that the ‘reduction’ $X_0(N)_{\mathbf{F}_p}$ is again smooth for primes p not dividing N . See [17, Theorem 8.2.1] for a precise statement. We are mostly interested in a consequence of this result. Namely, let K be an algebraically closed field of characteristic $p \nmid N$ and let E, E' be two elliptic curves over K . The modular polynomial Φ_N has integer coefficients and we let $\bar{\Phi}_N \in \mathbf{F}_p[X, Y]$ be its reduction modulo p . Then there exists

an isogeny $E \rightarrow E'$ of degree N with cyclic kernel if and only if we have

$$\bar{\Phi}_N(j(E), j(E')) = 0.$$

In chapter 6 we will also consider the curve $Y(N) = Y_{\Gamma(N)}$. Points on $Y(N)$ have a natural interpretation as isomorphism classes of pairs (E, P, Q) , where E is a complex elliptic curve and $P, Q \in E[N]$ form a basis for the N -torsion $E[N]$ that satisfies $e_N(P, Q) = \exp(2\pi i/N)$. Here, we normalise the Weil pairing e_N on the complex curve E_τ such that we have $e_N(1/N, \tau/N) = \exp(2\pi i/N)$. The notion of isomorphism is similar to the case of $Y_0(N)$: two triples $(E, P, Q), (E', P', Q')$ are isomorphic if there exists an isomorphism of elliptic curves $\varphi : E \rightarrow E'$ with $\varphi(P) = P'$ and $\varphi(Q) = Q'$. This moduli interpretation can be made explicit by viewing $Y(N)$ as the quotient $\Gamma(N) \backslash \mathbf{H}$. The $\Gamma(N)$ -orbit of $\tau \in \mathbf{H}$ is then mapped to the isomorphism class of the triple

$$(E_\tau, 1/N, \tau/N),$$

consisting of the complex elliptic curve $E_\tau = \mathbf{C}/\langle 1, \tau \rangle$ and points $1/N, \tau/N \in E_\tau[N]$.

For any choice $a \in (\mathbf{Z}/N\mathbf{Z})^*$, we have a variant $Y(N)_a$ of $Y(N)$. As a complex curve $Y(N)_a$ is the same as $Y(N)$, hence equal to $\Gamma(N) \backslash \mathbf{H}$, but to $\tau \in \mathbf{H}$ we now associate the triple

$$(E_\tau, a/N, \tau/N).$$

We have $e_N(a/N, \tau/N) = \exp(2a\pi i/N)$.

The curves $Y(N)_a$ can be defined over $\mathbf{Q}(\zeta_N)$. For an extension $K/\mathbf{Q}(\zeta_N)$, we fix a primitive N -th root of unity $\zeta_N \in K$. For $N \geq 3$, the K -valued points $Y(N)_a(K)$ of $Y(N)_a$ are the K -isomorphism classes of triples (E, P, Q) , where E is an elliptic curve over K and $P, Q \in E[N]$ are N -torsion points of E that are defined over K and satisfy $e_N(P, Q) = \zeta_N^a$.

5.4 Computing the canonical lift

In this section we explain the mathematical idea underlying the algorithm to compute the canonical lift $\mu^{-1}(j)$ of an ordinary j -invariant $j \in \mathbf{F}_p$. A more algorithmic description will be given in sections 5.5–5.7. Throughout this section, we let $\bar{E}/\mathbf{F}_p$ be an ordinary elliptic curve with endomorphism ring $\mathcal{O}$. We have $j(\bar{E}) \neq 0, 1728 \in \mathbf{F}_p$. The canonical lift of $\bar{E}$ is denoted by $\hat{E}$.

Let $I \subset \mathcal{O}$ be an invertible $\mathcal{O}$ -ideal. As in section 3.2, we have a map

$$\rho_I : \text{Ell}_\Delta(\mathbf{Q}_p) \rightarrow \text{Ell}_\Delta(\mathbf{Q}_p)$$

that maps $j(\tilde{E})$ to $j(\tilde{E}^I)$. Here, the isogeny $\tilde{E} \rightarrow \tilde{E}^I$ has the group $\tilde{E}[I]$ of I -torsion points as kernel. The inverse of ρ_I is given by $\rho_{\bar{I}}$, where $\bar{I}$ is the complex conjugate of I . The map ρ_I is bijective.

For the remainder of this section we assume that the ideal $I \subset \mathcal{O}$ is coprime to p . We then obtain a bijection $\bar{\rho}_I : \text{Ell}_\Delta(\mathbf{F}_p) \rightarrow \text{Ell}_\Delta(\mathbf{F}_p)$ that sends $j(\bar{E})$ to $j(\bar{E}^I)$, and we have a commutative diagram


$$\begin{array}{ccc} j(\tilde{E}) & \xrightarrow{\rho_I} & j(\tilde{E}^I) \\ \downarrow & & \downarrow \\ j(\bar{E}) & \xrightarrow{\bar{\rho}_I} & j(\bar{E}^I). \end{array}$$

We have seen in section 3.2 that the map ρ_I induces an action of the Picard group $\text{Pic}(\mathcal{O})$ on $\text{Ell}_\Delta(\mathbf{Q}_p)$. This action is transitive and free. Similarly, we have an action of $\text{Pic}(\mathcal{O})$ on $\text{Ell}_\Delta(\mathbf{F}_p)$. The action of $\text{Pic}(\mathcal{O})$ on $\text{Ell}_\Delta(\mathbf{Q}_p)$ and $\text{Ell}_\Delta(\mathbf{F}_p)$ is compatible with reducing modulo p .

Let $\mathbf{C}_p$ be the completion of an algebraic closure of $\mathbf{Q}_p$. It is well known that $\mathbf{C}_p$ is itself algebraically closed. Define

$$X_\Delta(\mathbf{C}_p) = \{j \in \mathbf{C}_p \mid \bar{j} \in \text{Ell}_\Delta(\mathbf{F}_p)\} \subset \mathbf{C}_p.$$

The set $X_\Delta(\mathbf{C}_p)$ consists of $h(\Delta)$ open discs of p -adic radius 1 around the CM-points $\text{Ell}_\Delta(\mathbf{Q}_p)$. Every disc contains exactly one element of $\text{Ell}_\Delta(\mathbf{Q}_p)$ and it is this subset $\text{Ell}_\Delta(\mathbf{Q}_p)$ of $\mathbf{C}_p$ that we want to compute.



The picture visualises the situation. The elements of the set $\text{Ell}_\Delta(\mathbf{F}_p)$ are denoted by thick points. The set $X_\Delta(\mathbf{C}_p)$ is denoted by a series of open discs, one above each point in $\text{Ell}_\Delta(\mathbf{F}_p)$. Just as we denoted the canonical lift of a curve $\bar{E}/\mathbf{F}_p$ by $\tilde{E}$, we place a tilde above a thick point to denote the elements of $\text{Ell}_\Delta(\mathbf{Q}_p)$.

The fundamental idea in [12] is that the map $\rho_I : \text{Ell}_\Delta(\mathbf{Q}_p) \rightarrow \text{Ell}_\Delta(\mathbf{Q}_p)$ has a natural extension to a map $\rho_I : X_\Delta(\mathbf{C}_p) \rightarrow X_\Delta(\mathbf{C}_p)$, which we proceed to define. We let $N \in \mathbf{Z}_{>0}$ be the norm of I . Recall our assumption $p \nmid N$. Now take an arbitrary element $j \in X_\Delta(\mathbf{C}_p)$, and write $\bar{j} \in \mathbf{F}_p$ for its reduction modulo p . Pick a curve $\bar{E}/\mathbf{F}_p$ with j -invariant $j(\bar{E}) = \bar{j}$, and take any curve $E/\mathbf{C}_p$ with $j(E) = j \in \mathbf{C}_p$ that reduces to $\bar{E}/\mathbf{F}_p$. We have a natural isomorphism

$$\varphi : E[N] \xrightarrow{\sim} \bar{E}[N]$$

by the assumption that N is coprime to p . The subgroup $\bar{E}[I] \subset \bar{E}[N]$ has a well defined inverse image under φ . We denote $\varphi^{-1}(\bar{E}[I])$ by $E[I]$. The group $E[I]$ is a subgroup of order N of $E[N]$. Using fancy language, this provides a lift of $\bar{E}[I]$ to a *group scheme* over the p -adic disc in $X_\Delta(\mathbf{C}_p)$ lying over $\bar{j} \in \text{Ell}_\Delta(\mathbf{F}_p)$. We define $\rho_I(j) = j(E^I)$. The j -invariant $j(E^I)$ is independent of the choice of E and therefore, ρ_I is well-defined.

For $j \in \text{Ell}_\Delta(\mathbf{Q}_p)$ we now have two definitions of ρ_I : one in terms of a Galois action and one in terms of a group scheme. A moment's reflection shows however that these two definitions coincide. In sections 5.5–5.7 we will show how we can explicitly compute $\rho_I(j)$ for $j \in X_\Delta(\mathbf{C}_p)$.

REMARK. For two invertible $\mathcal{O}$ -ideals I, J that are coprime to p , we have $\rho_{IJ} = \rho_I \rho_J$. Furthermore, if J is contained in $\mathbf{Z}$, we have $\rho_J = \text{id}$.

The map ρ_I has a geometric interpretation. After possibly multiplying with a principal fractional ideal, we assume that $\mathcal{O}/I$ is cyclic. Again, let N be the norm of I . We map $Y_0(N)_{\mathbf{C}_p}$ inside $\mathbf{A}_{\mathbf{C}_p}^1 \times \mathbf{A}_{\mathbf{C}_p}^1$:

$$\begin{array}{ccccc} Y_0(N)_{\mathbf{C}_p} & \longrightarrow & C(\mathbf{C}_p) & \xrightarrow{\quad} & \mathbf{A}^1(\mathbf{C}_p) \times \mathbf{A}^1(\mathbf{C}_p) \xrightarrow{p_1} \mathbf{A}^1(\mathbf{C}_p) \\ \downarrow \Psi & & & & \downarrow p_2 \\ (E, G) & \longmapsto & (j(E), j(E/G)) & & \mathbf{A}^1(\mathbf{C}_p). \end{array}$$

The maps p_1, p_2 are the normal projection maps. The curve C is defined by $\Phi_N = 0$, with Φ_N the classical modular polynomial. Take a j -invariant $j(E) \in X_\Delta(\mathbf{C}_p)$. The fiber $p_1^{-1}(j(E)) \subset C(\mathbf{C}_p)$ above $j(E)$ consists of the points $(j(E), j(E/G_i))$, with G_i ranging over the $\psi(N)$ cyclic subgroups of order N of $E[N]$. We have $\rho_I(j(E)) = j(E^I) = p_2((j(E), j(E^I)))$.

In other words, we have chosen two functions $j_1, j_2 : Y_0(N)_{\mathbf{C}_p} \rightarrow \mathbf{A}_{\mathbf{C}_p}^1$. They are defined by $j_1((E, G)) = j(E)$ and $j_2((E, G)) = j(E/G)$. For $j(E) \in X_\Delta(\mathbf{C}_p)$,

we have $\rho_I(j(E)) = j_2((E, E[I]))$. We will often write $j_1(E)$ instead of $j_1((E, G))$ if there cannot be any confusion about which subgroup $E[G] \subset E[N]$ we mean. Likewise for j_2 .

Let now $I \subset \mathcal{O}$ be a *principal* ideal, and let $\alpha \in \mathcal{O}$ be a generator. We keep the assumption that p does not divide the norm of I . We write ρ_α to denote the map $\rho_{(\alpha)}$.

THEOREM 5.4. *Let $(\alpha) \subset \mathcal{O}$ be a principal ideal such that $\mathcal{O}/(\alpha)$ is cyclic as abelian group. Assume that (α) is coprime to p . Then the map $\rho_\alpha : X_\Delta(\mathbf{C}_p) \rightarrow X_\Delta(\mathbf{C}_p)$ is analytic, i.e., it can locally be given by a power series.*

PROOF. Take an elliptic curve $E/\mathbf{Q}_p$ with $j(E) \in \text{Ell}_\Delta(\mathbf{Q}_p)$, and assume that E has good reduction modulo p . We define $P = (E, E[(\alpha)]) \in Y_0(N)(\mathbf{C}_p)$. Note that P lies on the diagonal if we map $Y_0(N)_{\mathbf{C}_p}$ into $\mathbf{A}_{\mathbf{C}_p}^1 \times \mathbf{A}_{\mathbf{C}_p}^1$, i.e., we have

$$j_1(E) = j_2(E).$$

Since $\overline{j(E)}$ is the j -invariant of an ordinary curve over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$ of discriminant $\Delta < -4$, we see that $\overline{j(E)}$ is not equal to 0, $1728 \in \mathbf{F}_p$. It follows that $j(E) \in \mathbf{C}_p$ has positive p -adic distance to 0, $1728 \in \mathbf{C}_p$.

First we show that $P = (E, E[(\alpha)])$ may be defined over $\mathbf{Q}_p$, i.e., that both E and $E[(\alpha)]$ may be defined over $\mathbf{Q}_p$. By assumption, the curve E is defined over $\mathbf{Q}_p$. The prime p splits in $\mathcal{O}$, and we have $\alpha \in \mathbf{Q}_p$. This shows that $E[(\alpha)]$ is defined over $\mathbf{Q}_p$ and we have

$$P = (E, E[(\alpha)]) \in Y_0(N)(\mathbf{Q}_p).$$

Now consider the local ring $\mathcal{O}_{Y_0(N)_{\mathbf{Q}_p}, P}$ and its completion $\widehat{\mathcal{O}}_{Y_0(N)_{\mathbf{Q}_p}, P}$ at the point P . Since $Y_0(N)_{\mathbf{Q}_p}$ is a smooth curve, $\widehat{\mathcal{O}}_{Y_0(N)_{\mathbf{Q}_p}, P}$ is a complete discrete valuation ring over $\mathbf{Q}_p$. Since $j_1(E)$ and $j_2(E)$ are not equal to one of the ramification points $j = 0, 1728$ of the cover $Y_0(N)_{\mathbf{Q}_p}/\mathbf{A}_{\mathbf{Q}_p}^1$, the functions $j_1 - j_1(E)$ and $j_2 - j_2(E)$ are uniformising parameters for $\widehat{\mathcal{O}}_{Y_0(N)_{\mathbf{Q}_p}, P}$.

The isomorphism $\widehat{\mathcal{O}}_{Y_0(N)_{\mathbf{Q}_p}, P} \cong \mathbf{Q}_p[[j_1 - j_1(E)]]$ shows that we can express $j_2 - j_2(E)$ as a formal power series in $j_1 - j_1(E)$:

$$j_2 - j_2(E) = \sum_{i \geq 1} c_i (j_1 - j_1(E))^i \quad \text{with } c_i \in \mathbf{Q}_p.$$

The theorem follows if we prove that the coefficients c_i of the power series lie in $\mathbf{Z}_p$. We prove this in lemma 5.5. □

LEMMA 5.5. *The coefficients c_i in the power series above lie in $\mathbf{Z}_p$.*

PROOF. The proof of this lemma is more difficult, since it requires working with the modular curve $X_0(N)$ over $\mathbf{Z}_p$. Working with curves over rings requires more geometric theory than we have given, and is not easily explained in a few lines. Hence, we assume more background in this proof.

As in [17, Section 9.3], we consider the modular curve $X_0(N)_{\mathbf{Q}_p}$ as a scheme over $\mathrm{Spec}(\mathbf{Q}_p)$. The diagram

$$\begin{array}{ccccc} X_0(N)_{\mathbf{Q}_p} & \xrightarrow{\quad} & X_0(N)_{\mathbf{Z}_p} & \xleftarrow{\quad} & X_0(N)_{\mathbf{F}_p} \\ \downarrow \bigg) \scriptstyle P & & \downarrow \bigg) \scriptstyle P' & & \downarrow \bigg) \scriptstyle P \\ \mathrm{Spec}(\mathbf{Q}_p) & \xrightarrow{\quad} & \mathrm{Spec}(\mathbf{Z}_p) & \xleftarrow{\quad} & \mathrm{Spec}(\mathbf{F}_p) \end{array}$$

explains the situation. We view the point P as a section $\mathrm{Spec}(\mathbf{Q}_p) \rightarrow X_0(N)_{\mathbf{Q}_p}$. As $X_0(N)_{\mathbf{Z}_p}$ is proper over $\mathrm{Spec}(\mathbf{Z}_p)$, there exists a unique section $P' : \mathrm{Spec}(\mathbf{Z}_p) \rightarrow X_0(N)_{\mathbf{Z}_p}$ making the left square commutative. The existence of $\bar{P} : \mathrm{Spec}(\mathbf{F}_p) \rightarrow X_0(N)_{\mathbf{F}_p}$ is automatic from the existence of the section P' .

Since we assumed $p \nmid N$, the curve $X_0(N)_{\mathbf{F}_p}$ is smooth over $\mathrm{Spec}(\mathbf{F}_p)$. We have $j_1(\bar{E}) = j_2(\bar{E}) \neq 0, 1728 \in \mathbf{F}_p$, and the functions $j_1 - j_1(E)$ and $j_2 - j_2(E)$ remain uniformising parameters for the complete discrete valuation ring $\hat{\mathcal{O}}_{X_0(N)_{\mathbf{F}_p}, \bar{P}}$ over $\mathbf{F}_p$. We get $(p, j_1 - j_1(E))$ and $(p, j_2 - j_2(E))$ as parameters for $\mathcal{O}_{X_0(N)_{\mathbf{Z}_p}, \bar{P}}$, and the ring $\mathcal{O}_{X_0(N)_{\mathbf{Z}_p}, \bar{P}}$ is a 2-dimensional regular local ring. Exactly as in the proof of [43, Theorem 29.7], we get an isomorphism

$$\hat{\mathcal{O}}_{X_0(N)_{\mathbf{Z}_p}, \bar{P}} \cong \mathbf{Z}_p[[j_1 - j_1(E)]]. \quad \square$$

The map ρ_α fixes the CM-points $\mathrm{Ell}_\Delta(\mathbf{Q}_p)$ and therefore stabilizes every disc. We have constructed an analytic map that has the CM-points as *fixed points*. We will use a kind of Newton iteration to converge to the j -invariant of the canonical lift $\tilde{E}/\mathbf{Q}_p$ of $\bar{E}/\mathbf{F}_p$ starting from a curve $E_1/\mathbf{C}_p$ that reduces to $\bar{E}$ modulo p . The following lemma gives the derivative of ρ_α in a CM-point, i.e., the first coefficient c_1 in the power series on the previous page.

LEMMA 5.6. *Let $(\alpha) \subset \mathcal{O}$ be a principal ideal such that $\mathcal{O}/(\alpha)$ is cyclic as abelian group. Assume that (α) is coprime to p . Then the derivative of ρ_α in $j(\tilde{E}) \in \mathrm{Ell}_\Delta(\mathbf{Q}_p)$ is given by $\alpha\bar{\alpha}^{-1}$, where $\bar{\alpha}$ is the complex conjugate of α .*

PROOF. This is lemma 1 in [12]. The proof there is rooted in a complex analytic setting. After attending a talk by Couveignes on this topic, Edixhoven observed

that this lemma can be proven completely geometrically, see proposition 3.3.2 of his thesis [18]. For convenience, we give the (slightly modified) proof from [12]. The main difference with the proof in [12] is that we have removed the explicit computation of normal forms of ideals.

Let N be the norm of the principal $\mathcal{O}$ -ideal (α) . We take a curve $E_{\overline{\mathbf{Q}}}$ defined over $\overline{\mathbf{Q}}$ with $j(E_{\overline{\mathbf{Q}}}) = j(E) \in \text{Ell}_{\Delta}(\mathbf{Q}_p)$. This gives a point $P_{\overline{\mathbf{Q}}} = (E_{\overline{\mathbf{Q}}}, E_{\overline{\mathbf{Q}}}[(\alpha)]) \in Y_0(N)(\overline{\mathbf{Q}})$. After a base change to $\mathbf{C}$, we get a point $P \in Y_0(N)(\mathbf{C})$. We will work with the modular curve $Y_0(N)_{\mathbf{C}}$ over $\mathbf{C}$. The diagram

$$\begin{array}{ccccccc}
 & & & \xrightarrow{j_1} & & & \\
 Y_0(N)(\mathbf{C}) & \xrightarrow{\quad} & C(\mathbf{C}) & \xrightarrow{\quad} & \mathbf{A}^1(\mathbf{C}) \times \mathbf{A}^1(\mathbf{C}) & \longrightarrow & \mathbf{A}^1(\mathbf{C}) \\
 & \searrow j_2 & & & \downarrow & & \\
 & & & & \mathbf{A}^1(\mathbf{C}) & &
 \end{array}$$

gives the two functions $j_1, j_2 : Y_0(N)_{\mathbf{C}} \rightarrow \mathbf{A}_{\mathbf{C}}^1$ that we have chosen, i.e., we have $j_1((E, G)) = j(E)$ and $j_2((E, G)) = j(E/G)$. For $(E, G) = (E_{\tau}, \langle 1/N \rangle)$ we have $j_1(E) = j(\tau)$ and $j_2((E_{\tau}), \langle 1/N \rangle) = j(N\tau)$. Let $F = \mathbf{C}(j_1, j_2)$ be the function field of $Y_0(N)_{\mathbf{C}}$ and let $\Omega_{F/\mathbf{C}}$ be its module of Kähler differentials. The module $\Omega_{F/\mathbf{C}}$ has dimension 1 as a vector space over F . Hence, there is an element $\sigma \in F$ with $\sigma dj_1 = dj_2$. We map $Y_0(N)_{\mathbf{C}}$ to the curve C inside $\mathbf{A}_{\mathbf{C}}^1 \times \mathbf{A}_{\mathbf{C}}^1$. The function value $\sigma((E, E[(\alpha)]) \in \mathbf{C}$ is the slope of the tangent line at $(j_1(E), j_2(E)) \in C$ at the branch of (E, G) . We have $c_1 = \sigma(P)$.

View $Y_0(N)_{\mathbf{C}}$ as the quotient $\Gamma_0(N) \backslash \mathbf{H}$ and choose a representative $\tau \in \mathbf{H}$ of $P \in Y_0(N)_{\mathbf{C}}$. Defining $j_N(z) = j(Nz)$, we can compute c_1 as

$$c_1 = \frac{dj_N}{dj}(\tau).$$

Let $j' = \frac{dj}{d\tau}$ be the derivative of the j -function and let $G_i(\tau)$ be the i -th Eisenstein series attached to the lattice $\langle 1, \tau \rangle$.

Claim. There exists a constant $c \in \mathbf{C}$ with

$$\frac{j'}{j} = c \frac{G_6}{G_4}.$$

Proof of claim. The j -function has a triple zero at ζ_3 , and has no other zeroes in the standard fundamental domain of $\text{SL}_2(\mathbf{Z}) \backslash \mathbf{H}$. The quotient j'/j is a rational modular form of weight 2, with a simple pole at ζ_3 .

The quotient G_6/G_4 is a modular form of weight 2 and has a simple pole at ζ_3 . There exists a constant $c \in \mathbf{C}$ such that $j'/j - cG_6/G_4$ has no poles on the upper half plane $\mathbf{H}$. We see that $j'/j - cG_6/G_4$ is a modular function of weight 2 which is everywhere holomorphic, including infinity. It is therefore equal to zero, which proves our claim.

We derive $\frac{dj}{d\tau}(\tau) = cG_6(\tau)j(\tau)/G_4(\tau)$ and $\frac{dj_N}{d\tau}(\tau) = N\frac{dj}{d\tau}(N\tau)$, i.e., we have

$$c_1 = \frac{dj_N}{d\tau} \frac{d\tau}{dj}(\tau) = N \frac{j(N\tau)}{j(\tau)} \cdot \frac{(G_6/G_4)(N\tau)}{(G_6/G_4)(\tau)}.$$

The curve $E_\tau = \mathbf{C}/\langle 1, \tau \rangle$ has endomorphism ring $\mathcal{O}$ and we have a commutative diagram

$$\begin{array}{ccc} \mathbf{C}/\langle 1, \tau \rangle & \xrightarrow{\times N} & \mathbf{C}/\langle 1, N\tau \rangle \\ & \searrow \times \alpha & \swarrow \sim \times \frac{\alpha}{N} \\ & \mathbf{C}/\langle 1, \tau \rangle, & \end{array}$$

since α is an endomorphism of E_τ . We see that we have $\frac{\alpha}{N}\langle 1, N\tau \rangle = \langle 1, \tau \rangle$, i.e., we get $\alpha\tau = a\tau + b$, $\alpha/N = c\tau + d$ with $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z})$. Using the relation $N\tau = (a\tau + b)/(c\tau + d)$, we compute

$$c_1 = N \frac{(G_6/G_4)(N\tau)}{(G_6/G_4)(\tau)} = N(c\tau + d)^2 = \frac{\alpha^2}{N} = \frac{\alpha}{\bar{\alpha}}. \quad \square$$

We return to the problem of computing the j -invariant of the canonical lift $\tilde{E}/\mathbf{Q}_p$ of an ordinary curve $\bar{E}/\mathbf{F}_p$ with endomorphism ring $\mathcal{O}$. We are looking for a fixed point of ρ_α , i.e., for a zero of the function $\rho_\alpha - \mathrm{id}$. We use a Newton iteration process to converge to a zero of $\rho_\alpha - \mathrm{id}$. First we pick an elliptic curve $E_1/\mathbf{C}_p$ that reduces to $\bar{E}/\mathbf{F}_p$ modulo p . Assume that we have $\alpha/\bar{\alpha} - 1 \in \mathbf{Z}_p^*$ and consider the following iteration process

$$j(E_{k+1}) = j(E_k) - \frac{\rho_\alpha(j(E_k)) - j(E_k)}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1}.$$

This computation is carried out with 2 digits precision for $k = 1$ and the precision is doubled in each iteration step. This process is a modified version of Newton iteration. For classical Newton iteration we would need $\rho'_\alpha(j(E_k)) - 1$ in the denominator instead of $(\alpha/\bar{\alpha}) - 1 = \rho'_\alpha(j(\tilde{E})) - 1$. We are working with bounded precision in each step and we have to check that

$$\frac{\rho_\alpha(j(E_k)) - j(E_k)}{\rho'_\alpha(j(\tilde{E})) - 1} = \frac{\rho_\alpha(j(E_k)) - j(E_k)}{\rho'_\alpha(j(E_k)) - 1} \in \mathbf{Z}_p/(p^{2^k}) \quad (*)$$

holds in the k -th iteration step. For $k = 1$, we have $j(E_1) = j(\tilde{E}) \bmod p$, and therefore also $\rho'_\alpha(j(E_1)) = \rho'_\alpha(j(\tilde{E})) \bmod p$. As $\rho_\alpha(j(E_1)) - j(E_1)$ is divisible by p , we see that $(*)$ holds for $k = 1$, i.e., modulo p^2 . Now suppose $k > 1$. With induction we see that $j(E_k) = j(\tilde{E}) \bmod p^{2^{k-1}}$ holds and $\rho_\alpha(j(E_k)) - j(E_k)$ is divisible by $p^{2^{k-1}}$. We conclude that equality $(*)$ holds for all $k \in \mathbf{Z}_{\geq 1}$, and that for $\alpha/\bar{\alpha} - 1 \in \mathbf{Z}_p^*$, the process above converges to the j -invariant $j(\tilde{E})$ of the canonical lift.

5.5 Isogenous curves with isomorphic endomorphism rings

In order to compute $\rho_\alpha(j)$ for $j \in \text{Ell}_\Delta(\mathbf{C}_p)$, we need a more algorithmic description of the map ρ_α . Pick any principal $\mathcal{O}$ -ideal (α) that is coprime to p . Without loss of generality we assume that (α) is *primitive*, i.e., (α) is not divisible by elements of $\mathbf{Z}$.

Fix $j \in \text{Ell}_\Delta(\mathbf{C}_p)$ and write $\bar{j} \in \mathbf{F}_p$ for its reduction modulo p . Take an elliptic curve $\bar{E}/\mathbf{F}_p$ with j -invariant $\bar{j} \in \mathbf{F}_p$. Note that the curve $\bar{E}$ has endomorphism ring $\mathcal{O}$, and j -invariant $j(\bar{E}) \neq 0, 1728 \in \mathbf{F}_p$.

Let $(\alpha) = \prod_i \mathfrak{l}_i^{e_i}$ be the prime factorization of (α) . If we can compute $\rho_{\mathfrak{l}_i}(j)$, we can compute $\rho_\alpha(j)$ since we have $\rho_{\mathfrak{l}_i \mathfrak{l}_j} = \rho_{\mathfrak{l}_i} \rho_{\mathfrak{l}_j}$. We fix a prime ideal $\mathfrak{l} | (\alpha)$ of degree 1. We let $l \in \mathbf{Z}$ be the norm of $\mathfrak{l}$. Note that l is prime, and we have $l \neq p$.

We want to compute the group of $\mathfrak{l}$ -torsion points $\bar{E}[\mathfrak{l}] \subset \bar{E}[l]$. Algorithmically, we will ‘code’ the x -coordinates of the points in $\bar{E}[\mathfrak{l}]$ as roots of a polynomial $\bar{f}_\mathfrak{l} \in \mathbf{F}_p[X]$. More precisely, the polynomial $\bar{f}_\mathfrak{l}$ will have the property that $x \in \bar{\mathbf{F}}_p$ is a root of $\bar{f}_\mathfrak{l}$ if and only if there exists a point $P \in \bar{E}[\mathfrak{l}]$ with x -coordinate x .

LEMMA 5.7. *The polynomial $\bar{f}_\mathfrak{l} \in \mathbf{F}_p[X]$ that vanishes on the x -coordinates of the points in $\bar{E}[\mathfrak{l}]$ has degree $(l-1)/2$ for $l > 2$. For $l = 2$, the degree equals 1.*

PROOF. The group $\bar{E}[\mathfrak{l}]$ is a subgroup of $\bar{E}[l] \cong \mathbf{Z}/l\mathbf{Z} \times \mathbf{Z}/l\mathbf{Z}$. We see that $\#E[\mathfrak{l}]$ equals $1, l$ or l^2 . By the assumption that $\mathfrak{l}$ has degree 1, we see that we also have a subgroup $E[\bar{\mathfrak{l}}] \subset E[l]$, belonging to the conjugate ideal $\bar{\mathfrak{l}} \subset \mathcal{O}$. We see that $\bar{E}[\mathfrak{l}]$ has order l and hence for $l > 2$ the polynomial $\bar{f}_\mathfrak{l}$ will have degree $(l-1)/2$. For $l = 2$, the degree equals 1. $\square$

The inclusion $\bar{E}[\mathfrak{l}] \subset \bar{E}[l]$ yields that $\bar{f}_\mathfrak{l}$ is a divisor of the l -th division polynomial $\Psi_l \in \mathbf{F}_p[X]$. We recall that the l -th division polynomial vanishes exactly on the x -coordinates on the l -torsion points, cf. [36, Theorem 2.1]. The degree of Ψ_l therefore equals $(l^2 - 1)/2$ for odd primes l and 3 for $l = 2$. As an example, the 3-rd division

polynomial for the curve defined by $Y^2 = X^3 + aX + b$ is

$$\Psi_3(X) = 3X^4 + 6aX^2 + 12bX - a^2.$$

Before we can compute the polynomial $\tilde{f}_l \in \mathbf{F}_p[X]$, we need to know the j -invariant of the l -isogenous curve $\bar{E}^l$. In this section we explain that we can ‘almost’ compute this j -invariant.

THEOREM 5.8. *Let $\bar{E}/\mathbf{F}_p$ and l be as above. Then we have*

$$\Phi_l(j(\bar{E}), j(\bar{E}^l)) = 0,$$

where Φ_l denotes the classical l -th modular polynomial. Let t be the trace of Frobenius of $\bar{E}$. If the order $\tilde{\mathcal{O}}$ of discriminant $t^2 - 4p$ is maximal at l , then the polynomial

$$\Phi_l(j(\bar{E}), X) \in \mathbf{F}_p[X]$$

has exactly 2 roots in $\mathbf{F}_p$.

The first statement in the theorem follows immediately from the properties of the modular polynomial. Indeed, for any algebraically closed field $k = \bar{k}$ of characteristic $\text{char}(k) \neq l$ and for any $j \in k$, the roots of $\Phi_l(j, X) \in k[X]$ are exactly the j -invariants of curves that are l -isogenous to a curve with j -invariant j . We know that $j(\bar{E}^l)$ is contained in $\mathbf{F}_p$, and the first result follows.

Proving the second statement requires more work, and we postpone the proof to the end of this section. First we give some lemma’s that help us determine the number of roots in $\mathbf{F}_p$ of $\Phi_l(j(\bar{E}), X) \in \mathbf{F}_p[X]$.

LEMMA 5.9. *Let $E/\mathbf{F}_p$ be an elliptic curve and let $\varphi : E \rightarrow E'$ be a non-zero isogeny. Then the endomorphism algebras of E and E' are isomorphic.*

PROOF. The isogeny $\varphi : E \rightarrow E'$ induces an injective ring homomorphism

$$\begin{aligned} f_\varphi : \text{End}(E') &\rightarrow \text{End}(E) \otimes_{\mathbf{Z}} \mathbf{Q} \\ \psi &\mapsto \hat{\varphi}\psi\varphi \otimes m^{-1}, \end{aligned}$$

where $m \in \mathbf{Z}_{\geq 1}$ is the degree of φ , and $\hat{\varphi} : E' \rightarrow E$ is the dual isogeny of φ . The homomorphism corresponding to $\hat{\varphi}$ is the inverse of f_φ . □

By the Deuring lifting theorem, we can lift $\bar{E}/\mathbf{F}_p$ together with its endomorphism ring to a curve $\tilde{E}$ defined over the ring class field $H_{\mathcal{O}}$. We have $\text{End}(\tilde{E}) \cong \mathcal{O} \cong \text{End}(E)$. We have seen in section 3.2 that $\tilde{E}$ is isomorphic to $\mathbf{C}/I$, with I an

invertible $\mathcal{O}$ -ideal. The curves $\mathbf{C}/I$ can be given by Weierstraß equations over $H_{\mathcal{O}}$, and as such, they are all Galois conjugate to the curve $\mathbf{C}/\mathcal{O}$. There exists a prime $\mathfrak{P}$ of $H_{\mathcal{O}}$ such that $\mathbf{C}/\mathcal{O}$ reduces to $\bar{E}/\mathbf{F}_p$. We may therefore write $\tilde{E} \cong \mathbf{C}/\mathcal{O}$.

As l -isogenous curve to $\tilde{E}$ is isomorphic to $\mathbf{C}/\Lambda$, with Λ a lattice in $\mathbf{C}$ that contains $\mathcal{O}$ with index $[\Lambda : \mathcal{O}] = l$. After replacing Λ by $\frac{1}{l}\Lambda$, we see that an l -isogenous curve is isomorphic to $\mathbf{C}/\Lambda'$ with Λ' a lattice in $\mathbf{C}$ with $[\mathcal{O} : \Lambda'] = l$. There are $l + 1$ curves that are l -isogenous curves to $\tilde{E}$, corresponding to the $l + 1$ subgroups of index l of $\mathbf{Z}/l\mathbf{Z} \times \mathbf{Z}/l\mathbf{Z}$. The l -isogenous curves are defined over $\bar{\mathbf{Q}}$, and, for $l \neq p$, their reductions modulo p are the curves that are l -isogenous to $\bar{E}/\bar{\mathbf{F}}_p$.

LEMMA 5.10. *Let $E/\mathbf{F}_p$ be an ordinary elliptic curve and let $\varphi : E \rightarrow E'$ be an isogeny of prime degree $l \neq p$. Then $\mathcal{O} = \text{End}(E)$ contains $\mathcal{O}' = \text{End}(E')$ or $\mathcal{O}'$ contains $\mathcal{O}$ (inside the endomorphism algebra) and the index of the one in the other divides l .*

PROOF. We lift $E/\mathbf{F}_p$ together with its endomorphism ring to a curve $\tilde{E}/H_{\mathcal{O}}$, and write $\tilde{E} = \mathbf{C}/\mathcal{O}$. Likewise, we lift E' to $\tilde{E}'/H_{\mathcal{O}'}$. The curve $\tilde{E}'$ is then isomorphic to $\mathbf{C}/\Lambda$, with Λ a lattice in $\mathbf{C}$ that contains $\mathcal{O}$ with index l . The multiplier ring $\mathcal{O}$ of the lattice $\mathcal{O}$ contains $l\mathcal{O}'$ since $\mathcal{O}'$ is the multiplier ring of the lattice Λ . Furthermore, the lattice Λ has index l in $\frac{1}{l}\mathcal{O}$, and $\mathcal{O}'$ contains $l\mathcal{O}$. $\square$

From these two lemmas we derive that the endomorphism ring $\mathcal{O}'$ of an l -isogenous curve E' is an order in the imaginary quadratic field $K = \mathcal{O} \otimes_{\mathbf{Z}} \mathbf{Q}$. Furthermore, if $\mathcal{O}$ and $\mathcal{O}'$ are not equal, then $\mathcal{O}'$ is contained in $\mathcal{O}$ with index l (or vice versa). As we know that $\bar{E}^l$ has endomorphism ring $\mathcal{O}$, we want to know how many roots in $\mathbf{F}_p$ of $\Phi_l(j(\bar{E}), X) \in \mathbf{F}_p[X]$ are j -invariants of curves with endomorphism ring $\mathcal{O}$. The following lemma is proposition 23 in [34]. The proof there uses Tate modules. Our proof is based on the Deuring lifting theorem. We write $\mathcal{O}_l = \mathcal{O} \otimes \mathbf{Z}_l$.

LEMMA 5.11. *Let $E/\mathbf{F}_p$ be an ordinary elliptic curve with endomorphism ring $\mathcal{O}$ of discriminant Δ and with $j(E) \neq 0, 1728 \in \mathbf{F}_p$. Let $l \neq p$ be prime.*

- (i) *If $\mathcal{O}_l$ is maximal, there are exactly $(\frac{\Delta}{l}) + 1$ isogenies of degree l to curves with endomorphism ring $\mathcal{O}$. These isogenies are defined over $\mathbf{F}_p$.*
- (ii) *If $\mathcal{O}_l$ is non-maximal, there are no isogenies of degree l to curves with endomorphism ring $\mathcal{O}$.*
- (iii) *Assume that $\mathcal{O}_l$ is maximal. If there are more than $(\frac{\Delta}{l}) + 1$ isogenies of degree l over $\mathbf{F}_p$, then all $l + 1$ isogenies of degree l are defined over $\mathbf{F}_p$. This happens exactly when the index $[\mathcal{O} : \mathbf{Z}[F_p]]$ is divisible by l .*

PROOF. As before, we lift $\bar{E}/\mathbf{F}_p$ together with its endomorphism ring to a curve $\tilde{E}/H_{\mathcal{O}}$ with $\text{End}(\tilde{E}) = \mathcal{O}$, and we write $\tilde{E} = \mathbf{C}/\mathcal{O}$.

If an l -isogenous curve E' to $\tilde{E}$ has endomorphism ring $\mathcal{O}$, we can write $E' = \mathbf{C}/I$, with I an invertible $\mathcal{O}$ -ideal of norm $[\mathcal{O} : I] = l$. Part (i) of the lemma follows. If $\mathcal{O}_l$ is not maximal, there are no invertible $\mathcal{O}$ -ideals of norm l , and part (ii) also follows.

For part (iii), assume that $\mathcal{O}_l$ is maximal and suppose that there are more than $(\frac{\Delta}{l}) + 1$ isogenies of degree l over $\mathbf{F}_p$. There are $(\frac{\Delta}{l}) + 1$ curves with endomorphism ring $\mathcal{O}$ that are l -isogenous to E by part (i). The other $l - (\frac{\Delta}{l})$ curves have endomorphism ring $\mathcal{O}'$, where $\mathcal{O}'$ has index l in $\mathcal{O}$ by lemma 5.10 and the assumption that $\mathcal{O}_l$ is maximal. As $\mathbf{Z}[F_p]$ is contained in $\mathcal{O}'$, all the l -isogenous curves are defined over $\mathbf{F}_p$. $\square$

COROLLARY 5.12. *Let $l \neq p$ be prime. Let $E/\mathbf{F}_p$ be an ordinary elliptic curve with endomorphism ring $\mathcal{O}$ and with $j(E) \neq 0, 1728$. Assume that $\mathcal{O}_l$ is maximal. Then:*

- (i) *if $[\mathcal{O} : \mathbf{Z}[F_p]]$ is divisible by l , the polynomial $\Phi_l(j(E), X) \in \mathbf{F}_p[X]$ splits completely over $\mathbf{F}_p$*
- (ii) *otherwise, $\Phi_l(j(E), X) \in \mathbf{F}_p[X]$ has exactly $(\frac{\Delta}{l}) + 1$ roots in $\mathbf{F}_p$.*

PROOF OF THEOREM 5.8. It remains to prove the second statement in the theorem. This is now easy. Indeed, if l splits in $\mathcal{O}$ and if $\tilde{\mathcal{O}} \cong \mathbf{Z}[F_p]$ is maximal at l , then the polynomial $\Phi_l(j(\bar{E}), X) \in \mathbf{F}_p[X]$ has exactly 2 roots in $\mathbf{F}_p$ by corollary 5.12. $\square$

REMARK. The polynomial $\Phi_l(X, Y) \in \mathbf{Z}[X, Y]$ can easily be computed for relatively small values of l . An algorithm is given in [37, Chapter 5]. As an example, for $l = 2$ we get

$$\begin{aligned} \Phi_2 = & X^3 + Y^3 - X^2Y^2 + 1488(X^2Y + XY^2) - 162000(X^2 + Y^2) \\ & + 40773375XY + 8748000000(X + Y) - 15746400000000 \in \mathbf{Z}[X, Y]. \end{aligned}$$

5.6 Computing the kernel polynomial

Letting the notation be as in the previous section, we continue with the algorithmic description of the map $\rho_l : X_{\Delta}(\mathbf{C}_p) \rightarrow X_{\Delta}(\mathbf{C}_p)$. Write $\tilde{\mathcal{O}}$ for the imaginary quadratic order of discriminant $t^2 - 4p$. Here, t denotes the trace of Frobenius of the fixed elliptic curve $\bar{E}/\mathbf{F}_p$ with endomorphism ring $\mathcal{O} = \mathcal{O}_{\Delta}$.

We will have to make some assumptions on l . As we will see in section 5.7, these assumptions are harmless for our algorithm to compute the j -invariant of the

canonical lift of $\bar{E}$. First we assume that $\tilde{\mathcal{O}}$ is maximal at l . This assumption implies that the ring $\mathcal{O}$ is also maximal at l .

Theorem 5.8 tells us that the polynomial $\Phi_l(j(\bar{E}), X) \in \mathbf{F}_p[X]$ then has 2 roots in $\mathbf{F}_p$. Fix a root $h \in \mathbf{F}_p$ of $\Phi_l(j(\bar{E}), X)$. Let $\bar{E}/C$ have j -invariant h , corresponding to a cyclic subgroup $C \subset \bar{E}[l]$ of order l , i.e., C is the kernel of the isogeny $\bar{E} \rightarrow \bar{E}/C$. Note that $\bar{E}/C$ has endomorphism ring $\mathcal{O}$.

There are only two possibilities for C : we either have $C = E[l]$ or we have $C = E[\tilde{l}]$. The techniques that Elkies used to improve Schoof's original point counting algorithm [50, Sections 7, 8] allow us to compute, given h , a polynomial $f_C \in \mathbf{F}_p[X]$ that vanishes exactly on the x -coordinates of the points in C . We will see that this also enables us to determine whether we have $C = \bar{E}[l]$ or not. Since our point of view is rather different from that of Elkies, we summarize sections 7 and 8 of [50].

Let $Y^2 = X^3 + aX + b$ be a Weierstraß equation for $\bar{E}/\mathbf{F}_p$. We want to know a Weierstraß equation $Y^2 = X^3 + a'X + b'$ for $\bar{E}/C$. The idea is to lift the isogeny $\varphi : \bar{E} \rightarrow \bar{E}/C$ to characteristic 0 using the Deuring lifting theorem and use analytic functions to derive formulas for a' and b' . For reasons to become clear, it is convenient to use *Tate curves* for the lifted curves.

► Formulas for a' and b' using Tate curves

Recall that a complex elliptic curve $E/\mathbf{C}$ is isomorphic as Riemann surface to $\mathbf{C}/L$, with L a lattice of rank 2. For $\tau \in \mathbf{H}$, define $q = \exp(2\pi i\tau)$. We get a complex analytic isomorphism

$$\mathbf{C}/2\pi i(\mathbf{Z} + \tau\mathbf{Z}) \xrightarrow[\cong]{\exp} \mathbf{C}^*/q^{\mathbf{Z}}.$$

The curve $\mathbf{C}^*/q^{\mathbf{Z}}$ is called a Tate curve. It admits a Weierstraß equation

$$Y^2 = X^3 - \frac{E_4(q)}{48}X + \frac{E_6(q)}{864}, \quad (5.2)$$

where $E_4(q), E_6(q) \in \mathbf{Z}[[q]]$ are the power series

$$\begin{aligned} E_4(q) &= 1 + 240 \sum_{n=1}^{\infty} \frac{n^3 q^n}{1 - q^n} \\ E_6(q) &= 1 - 504 \sum_{n=1}^{\infty} \frac{n^5 q^n}{1 - q^n}. \end{aligned}$$

Interpreting q as $\exp(2\pi i\tau)$ with $\tau \in \mathbf{H}$, the power series E_k are the Fourier expansions of the *normalized* Eisenstein series $G_k/(2\zeta(k))$, cf. [55, Section 7.4].

Using Deuring's lifting theorem, we can lift the isogeny

$$0 \longrightarrow C \longrightarrow \bar{E} \xrightarrow{\varphi} \bar{E}/C \longrightarrow 0$$

to characteristic 0. We claim that there exists $q = \exp(2\pi i\tau) \in \mathbf{C}$ such that

- ◇ the Eisenstein series $E_4(q)$ and $E_6(q)$ are elements of the ring class field $H_{\mathcal{O}}$
- ◇ there exists a prime ideal $\mathfrak{P}$ of $H_{\mathcal{O}}$ lying over p such that the 'reduction' mod $\mathfrak{P}$ of the isogeny

$$0 \longrightarrow \mu_l \longrightarrow \mathbf{C}^*/q^{\mathbf{Z}} \xrightarrow{l} \mathbf{C}^*/q^{l\mathbf{Z}} \longrightarrow 0$$

yields the isogeny $\varphi : \bar{E} \rightarrow \bar{E}/C$ over $\mathbf{F}_p$. Here, by reducing we mean that we take the Weierstraß equation (5.2) for the curve $\mathbf{C}^*/q^{\mathbf{Z}}$ and reduce that modulo $\mathfrak{P}$.

To prove our claim, we first lift $\bar{E}/\mathbf{F}_p$ with its endomorphism ring to a curve $\tilde{E}$ defined over the ring class field $H_{\mathcal{O}}$. The curve defined by the Weierstraß equation

$$Y^2 + XY = X^3 - \frac{36}{j(\tilde{E}) - 1728}X - \frac{1}{j(\tilde{E}) - 1728}$$

has j -invariant $j(\tilde{E})$. For this specific equation, we have $c_4 = c_6 = j(\tilde{E})/(j(\tilde{E}) - 1728)$. Here, c_4, c_6 are the usual quantities associated to a Weierstraß equation. Hence, there exists $q \in \mathbf{C}$ such that $E_4(q)$ and $E_6(q)$ are *both* contained in $H_{\mathcal{O}}$. This proves our claim.

The isogenous curve $\mathbf{C}^*/q^{l\mathbf{Z}}$ admits the equation

$$Y^2 = X^3 - \frac{E_4(q^l)}{48}X + \frac{E_6(q^l)}{864},$$

and we have $a' \equiv -E_4(q^l)/48 \pmod{\mathfrak{P}}$ and $b' \equiv -E_6(q^l)/864 \pmod{\mathfrak{P}}$. Furthermore, with $\Delta(q), j(q) \in \mathbf{Z}[[q]]$ given by

$$\begin{aligned} \Delta(q) &= \frac{E_4(q)^3 - E_6(q)^2}{1728} = q \prod_{n=1}^{\infty} (1 - q^n)^{24} \\ j(q) &= \frac{E_4(q)^3}{\Delta(q)} = \frac{1}{q} + 744 + 196884q + \dots \end{aligned}$$

we have $j(q^l) \equiv h = j(\bar{E}/C) \pmod{\mathfrak{P}}$. We will give 'formulas' for $E_4(q^l)$ and $E_6(q^l)$ in terms of j and a suitable 'derivative' j' .

For a Laurent series $f(q) = \sum_n a_n q^n \in \mathbf{Z}[[q]]$ we denote by $f'(q)$ the Laurent series

$$f'(q) = q \frac{df}{dq} = \sum_n n a_n q^n.$$

If we interpret q as $\exp(2\pi i\tau)$ with $\tau \in \mathbf{H}$, and view the power series as a Fourier expansion, the differential operator $f \mapsto q \frac{df}{dq}$ is just the usual differentiation of $f(q)$ with respect to the variable $2\pi i\tau$.

LEMMA 5.13. *The following equalities hold in $\mathbf{Z}[[q]]$:*

$$\begin{aligned} (i) \quad & \frac{j'}{j} = -\frac{E_6}{E_4} \\ (ii) \quad & \frac{j'}{j-1728} = -\frac{E_4^2}{E_6}. \end{aligned}$$

PROOF. In the proof of lemma 5.6 we proved that the logarithmic derivative $(\frac{d}{d\tau}j)/j$ and G_6/G_4 differ by a constant. The same is therefore true for the normalisations j'/j and E_6/E_4 , viewed as complex functions on the upper half plane $\mathbf{H}$.

The constant term of the Fourier expansion of j'/j is -1 , and it is 1 for E_6/E_4 . This proves part (i) of the lemma. The proof of part (ii) is similar. $\square$

COROLLARY 5.14. *We have the following equalities of power series in $\mathbf{Z}[[q]]$:*

$$\begin{aligned} (i) \quad & E_4(q) = \frac{j'^2}{j(j-1728)} \\ (ii) \quad & E_6(q) = -\frac{j'^3}{j^2(j-1728)}. \end{aligned}$$

PROOF. Multiply 5.13(ii) once and twice by 5.13(i). $\square$

To find $E_4(q^l)$ and $E_6(q^l)$, and therefore $a', b' \in \mathbf{F}_p$, it remains to derive an expression for $\tilde{j}'(q)$, where we write $\tilde{j}(q) = j(q^l)$. We compute $\tilde{j}' \in \mathbf{Z}[[q]]$ using the identity

$$\Phi_l(j, \tilde{j}) = 0 \in \mathbf{Z}[[q]],$$

where Φ_l denotes the classical modular polynomial. Applying the differential operator $f \mapsto q \frac{df}{dq} = f'$ to this identity yields the following identity of Laurent series:

$$j'\Phi_X(j, \tilde{j}) + l\tilde{j}'\Phi_Y(j, \tilde{j}) = 0. \quad (5.3)$$

Here, Φ_X and Φ_Y denote the partial derivatives $\partial\Phi_l/\partial X$, $\partial\Phi_l/\partial Y$.

Suppose that $\Phi_Y(j(q), \tilde{j}(q)) \bmod \mathfrak{P}$ is non-zero. Using relation (5.3) and lemma 5.13, we derive

$$\tilde{j}'(q) \equiv s \stackrel{\text{def}}{=} -\frac{18}{l} \frac{b}{a} \frac{\Phi_X(j(\overline{E}), h)}{\Phi_Y(j(\overline{E}), h)} j(\overline{E}) \in \mathbf{F}_p.$$

Together with corollary 5.14, this enables us to compute the coefficients of the Weierstraß equation for $\bar{E}/C$. We obtain the following formulas for a' and b' :

$$a' = -\frac{1}{48} \frac{s^2}{h(h-1728)} \in \mathbf{F}_p$$

$$b' = -\frac{1}{864} \frac{s^3}{h^2(h-1728)} \in \mathbf{F}_p.$$

Note that although we use the complex analytic theory to justify our computations, all computations take place in $\mathbf{F}_p$.

These formulas clearly require $\Phi_Y(j(q), \tilde{j}(q)) \neq 0 \in \mathbf{F}_p$. From relation (5.3), we see that if we have $\Phi_Y(j(q), \tilde{j}(q)) = 0$, then we also have $\Phi_X(j(q), \tilde{j}(q)) = 0$. For $\Phi_X(j(q), \tilde{j}(q)) = \Phi_Y(j(q), \tilde{j}(q)) = 0 \in \mathbf{F}_p$, the point $(j, \tilde{j})$ is a singular point for the curve defined by $\Phi_l(X, Y) = 0$ over $\mathbf{F}_p$. Using some algebraic geometry, one can prove [50, Section 7] that we must necessarily have

$$|\Delta| \leq 4l^2,$$

where Δ is the discriminant of the endomorphism ring $\text{End}(\bar{E}) = \mathcal{O}$. The approach presented in this section will therefore not work for all $\alpha \in \mathcal{O}$. However, in the algorithm for computing the j -invariant of the canonical lift of $\bar{E}/\mathbf{F}_p$, we have the freedom to choose the element α , and consequently the prime l , ourselves. If we pick a *smooth* $\alpha \in \mathcal{O}$, the condition $|\Delta| > 4l^2$ is automatically fulfilled. For the remainder of this section, we assume $|\Delta| > 4l^2$.

► The coefficients of the kernel polynomial

Knowing a Weierstraß equation

$$Y^2 = X^3 + a'X + b'$$

for $\bar{E}/\mathbf{F}_p$, we proceed to compute the polynomial $f_C \in \mathbf{F}_p[X]$ that vanishes exactly on the x -coordinates of the points in $C \subset E[l]$. We introduce the following power series in $\mathbf{Z}[\frac{1}{6}, \zeta, \frac{1}{\zeta(1-\zeta)}][[q]]$:

$$x(\zeta; q) = \frac{1}{12} - 2 \sum_{n=1}^{\infty} \frac{q^n}{(1-q^n)^2} + \sum_{n \in \mathbf{Z}} \frac{\zeta q^n}{(1-\zeta q^n)^2},$$

$$y(\zeta; q) = \frac{1}{2} \sum_{n \in \mathbf{Z}} \frac{\zeta q^n (1 + \zeta q^n)}{(1-\zeta q^n)^3}.$$

If we interpret q as $\exp(2\pi i\tau)$ for some $\tau \in \mathbf{H}$, and write $\zeta = \exp(2\pi iz)$, then $x(\zeta; q)$ is just $(2\pi i)^2 \wp(z; \tau)$ and we have $y(\zeta; q) = (2\pi i)^3 \wp'(z; \tau)/2$. The importance of these power series lies in the following lemma.

LEMMA 5.15. *We have the following equalities of power series:*

(i)

$$y(\zeta; q)^2 = x(\zeta; q)^3 - \frac{E_4(q^l)}{48}x(\zeta; q) + \frac{E_6(q^l)}{864}.$$

(ii)

$$\sum_{\zeta \in \mu_l, \zeta \neq 1} x(\zeta; q) = \frac{1}{12}l(E_2(q) - lE_2(q^l)).$$

Here, $E_2(q)$ is given by $E_2(q) = 1 - 24 \sum_{n=1}^{\infty} \frac{nq^n}{1 - q^n} \in \mathbf{Z}[[q]]$.

PROOF. This is proposition 7.2 in [50]. □

The map $\zeta \mapsto (x(\zeta; q), y(\zeta; q))$ gives an isomorphism between $\mathbf{C}^*/q^{\mathbf{Z}}$ and the $\mathbf{C}$ -valued points of $Y^2 = X^3 - (E_4(q)/48)X + E_6(q)/864$. The points in C are exactly the ones that correspond to $\zeta \in \mu_l$. The hardest part in computing the coefficients of f_C is to compute the sum p_1 of the x -coordinates of the points in C . From lemma 5.15(ii) we see that p_1 equals

$$\frac{1}{12}l(E_2(q) - lE_2(q^l))$$

for the $q \in \mathbf{C}$ belonging to our lifted curve $\mathbf{C}^*/q^{\mathbf{Z}}$.

LEMMA 5.16. *The following equalities hold in $\mathbf{Z}[[q]]$:*

$$\begin{aligned} \text{(i)} \quad & \frac{j''}{j'} = \frac{1}{6}E_2 - \frac{1}{2}\frac{E_4^2}{E_6} - \frac{2}{3}\frac{E_6}{E_4} \\ \text{(ii)} \quad & \frac{j''}{j'} - l\frac{\tilde{j}''}{\tilde{j}'} = -\frac{j'^2\Phi_{XX}(j, \tilde{j}) + 2lj'\tilde{j}'\Phi_{XY}(j, \tilde{j}) + l^2\tilde{j}'^2\Phi_{YY}(j, \tilde{j})}{j'\Phi_X(j, \tilde{j})}. \end{aligned}$$

PROOF. For part (i), see [50, Proposition 7.1(iii)]. Part (ii) follows from differentiating the identity $\Phi_l(j, \tilde{j}) = 0$ twice. □

Formula manipulation yields $p_1 =$

$$\frac{l}{6} \left[-3 \frac{h^2\Phi_{XX}(j, h) + 2ljs\Phi_{XY}(j, h) + l^2s^2\Phi_{YY}(j, h)}{h\Phi_X(j, h)} - 4l\frac{a'^2}{b'} + 36l\frac{b'}{a'} + 4\frac{a^2}{b} - 36\frac{b}{a} \right],$$

where we temporarily write $j = j(\bar{E})$. To compute the other coefficients of f_C , it is more convenient to change the isogeny a bit. Let $\tilde{E}/H_O$ be complex analytically isomorphic to $\mathbf{C}/(\omega_1\mathbf{Z} + \omega_2\mathbf{Z})$. We also write $Y^2 = X^3 + aX + b$ for the Weierstraß equation for $\tilde{E}$ that reduces to the chosen equation $Y^2 = X^3 + aX + b$ for $\bar{E}/\mathbf{F}_p$. We lifted the isogeny

$$\bar{E} \rightarrow \bar{E}/C$$

over $\mathbf{F}_p$ to an isogeny

$$\mathbf{C}/(\omega_1\mathbf{Z} + \omega_2\mathbf{Z}) \rightarrow \mathbf{C}/(\omega_1\mathbf{Z} + l\omega_2\mathbf{Z})$$

given by $z \mapsto lz$ in characteristic 0. Write $Y^2 = X^3 + a'X + b'$ for the isogenous curve. To compute the coefficients of f_C , it is more convenient to work with the isogeny

$$\mathbf{C}/(\omega_1\mathbf{Z} + \omega_2\mathbf{Z}) \rightarrow \mathbf{C}/(\frac{1}{l}\omega_1\mathbf{Z} + \omega_2\mathbf{Z})$$

given by $z \mapsto z$. The kernel of this isogeny is the same as before, but the Weierstraß equation for the isogenous curve is given by

$$Y^2 = X^3 + \frac{a'}{l^4}X + \frac{b'}{l^6}.$$

We introduce a notation for the coefficients of the Weierstraß $\wp$ -function. Let $\wp(z)$ be the Weierstraß $\wp$ -function for the lattice $\omega_1\mathbf{Z} + \omega_2\mathbf{Z}$ and let

$$\wp(z) = \frac{1}{z^2} + \sum_{k=1}^{\infty} c_k z^{2k}$$

be the corresponding Laurent series. Explicitly, one has

$$c_1 = \frac{-a}{5}, \quad c_2 = -b/7$$

$$c_k = \frac{3}{(k-2)(2k+3)} \sum_{j=1}^{k-2} c_k c_{k-i-j}, \quad (k \geq 3).$$

The coefficients c'_k , corresponding to the lattice $\frac{\omega_1}{l}\mathbf{Z} + \omega_2\mathbf{Z}$, are defined similarly.

LEMMA 5.17. *Let l be prime and let f be the polynomial that vanishes on the x -coordinates of the points in the kernel of the isogeny*

$$\mathbf{C}/(\omega_1\mathbf{Z} + \omega_2\mathbf{Z}) \rightarrow \mathbf{C}/(\frac{1}{l}\omega_1\mathbf{Z} + \omega_2\mathbf{Z}).$$

Then

$$z^{l-1}f(\wp(z)) = \exp \left(-\frac{1}{2}p_1z^2 - \sum_{k=1}^{\infty} \frac{c'_k - lc_k}{(2k+1)(2k+2)} z^{2k+2} \right).$$

PROOF. This is theorem 8.3 in [50]. $\square$

By equating powers of z on the left hand side and on the right hand side, we can now compute the coefficients of $f = x^{(l-1)/2} + a_{\frac{l-3}{2}}x^{(l-3)/2} + \dots + a_0$. As an example:

$$\begin{aligned} a_{\frac{l-3}{2}} &= -\frac{p_1}{2} \\ a_{\frac{l-5}{2}} &= \frac{1}{8}p_1^2 - \frac{c'_1 - lc_1}{12} - \frac{l-1}{2}c_1. \end{aligned}$$

These formulas are valid modulo p for $l < p$. As we have $|\Delta| < 4p$, our assumption $l < \sqrt{-\Delta/4}$ implies that we have $l < p$. We can therefore compute the polynomial $f_C \in \mathbf{F}_p[X]$.

► Checking the eigenspace

The techniques presented in this section enable us to compute, on input a zero $h \in \mathbf{F}_p$ of $\Phi(j(\overline{E}), X) \in \mathbf{F}_p[X]$, the polynomial $f_C \in \mathbf{F}_p[X]$ corresponding to the isogeny $E \rightarrow E/C$. Here E/C has j -invariant $h \in \mathbf{F}_p$. We either have $f_C = \bar{f}_l$ or $f_C = \bar{f}_l$, and we now explain how we can check, for $l \subseteq \tilde{\mathcal{O}}$, in which case we are.

Write $l = (l, c + d\pi_p)$, with $l \nmid c$. Here, $\pi_p \in \mathcal{O}$ is the image of the Frobenius $F_p \in \text{End}(\overline{E})$ under the fixed isomorphism $\text{End}(\overline{E}) \xrightarrow{\sim} \mathcal{O}$.

The Frobenius acts on $l \subset \overline{E}[l]$ as multiplying by $-c/d \in \mathbf{F}_l$. We test if

$$(X^p, Y^p) = (-c/d) \cdot (X, Y)$$

holds for the points in C , i.e., we compute both (X^p, Y^p) and $(-c/d) \cdot (X, Y)$ in the ring

$$\mathbf{F}_p[X, Y]/(f_C(X), Y^2 - X^3 - aX - b).$$

Note that the $\cdot$ means repeated adding *on the curve* and $(-c/d) \cdot (X, Y)$ can be computed by employing division polynomials.

If we find that $\bar{f}_C$ does not equal $\bar{f}_l$ we know that the unique other zero $h_2 \in \mathbf{F}_p$ of

$$\gcd(X^p - X, \Phi_l(j(\overline{E}), X)) \in \mathbf{F}_p[X]$$

must be the j -invariant of $\overline{E}^l$ and we repeat the computation from the beginning of this section with h replaced by h_2 to find the polynomial $f_C = \bar{f}_l \in \mathbf{F}_p[X]$.

5.7 Algorithm for computing the canonical lift

In this section we give the algorithm to compute the j -invariant of a canonical lift of an ordinary elliptic curve $\bar{E}/\mathbf{F}_p$ with j -invariant $j(\bar{E}) \neq 0, 1728$ and endomorphism ring $\text{End}(\bar{E}) = \mathcal{O} = \mathcal{O}_\Delta$. We can choose the element $\alpha \in \mathcal{O}$ that we use for the map $\rho_\alpha : X_\Delta(\mathbf{C}_p) \rightarrow X_\Delta(\mathbf{C}_p)$ ourselves. We recall the conditions that α should satisfy.

1. α is contained in $\mathbf{Z}[F_p] \cong \tilde{\mathcal{O}}$
2. α is primitive
3. $\alpha/\bar{\alpha} - 1$ is a p -adic unit
4. for any prime divisor l of $N(\alpha)$, we have $l \nmid [\mathcal{O} : \tilde{\mathcal{O}}]$
5. for any prime divisor l of $N(\alpha)$, we have $l < \sqrt{-\Delta/4}$

The input of the algorithm below consists of an ordinary curve $\bar{E}/\mathbf{F}_p$ with j -invariant $j(\bar{E}) \neq 0, 1728 \in \mathbf{F}_p$, an element $\alpha \in \mathcal{O} \setminus \mathbf{Z}$ as above, together with the factorization $(\alpha) = \prod_i \mathfrak{l}_i$ into prime ideals, and a positive integer k . The output is the j -invariant of the canonical lift $j(\tilde{E})$ in k digits accuracy.

Step 1. As in sections 5.5–5.6, compute the polynomial $\bar{f}_{l_1} \in \mathbf{F}_p[X]$ corresponding to the subgroup $\bar{E}[l_1] \subset \bar{E}[l_1]$. In the same way, we compute a cycle of isogenies

$$\bar{E} \xrightarrow{l_1} \bar{E}^{l_1} \longrightarrow \dots \xrightarrow{l_n} \bar{E}^{(\alpha)} \cong \bar{E}. \quad (*)$$

over $\mathbf{F}_p$. The isomorphism $\bar{E}^{(\alpha)} \cong \bar{E}$ follows from the fact that principal ideals act trivially. This is a good check for our computations so far.

Step 2. Choose an arbitrary lift $E_1/\mathbf{Q}_p$, in two p -adic digits precision, of $\bar{E}/\mathbf{F}_p$.

Step 3. Lift the cycle $(*)$ of isogenies over $\mathbf{F}_p$ to a ‘cycle’ of isogenies over $\mathbf{Q}_p$ in the following way. For $l = l_1$, we lift $j(\bar{E}^{l_1}) \in \mathbf{F}_p$ to $\mathbf{Z}_p$ as a root of $\Phi_l(j(E_1), X) \in \mathbf{Z}_p[X]$ to $h \in \mathbf{Z}_p$. We use Hensel’s lemma for this lifting process. Hensel requires that $\frac{d}{dX}\Phi_l(j(\bar{E}), X)$ is non-zero when evaluated in $X = j(\bar{E}^{l_1}) \in \mathbf{F}_p$. This requirement is satisfied by assumption 5.

We have $h = \rho_l(j(E_1))$. In the same way, we compute the ‘cycle’ of j -invariants

$$j(E_1) \xrightarrow{\rho_{l_1}} j(E_1)^{l_1} \longrightarrow \dots \xrightarrow{\rho_{l_n}} j(E_1)^{(\alpha)},$$

over $\mathbf{Q}_p$. This computation is carried out with two p -adic digits precision.

Step 4. Update $j(E_1)$ according to the Newton formula

$$j(E_{k+1}) = j(E_k) - \frac{\rho_\alpha(j(E_k)) - j(E_k)}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1}$$

to find $j(E_2) \in \mathbf{Z}_p$. The value $j(E_2)$ is the two digit approximation of the canonical lift $j(\tilde{E})$.

Step 5. Lift the cycle $(*)$ to a ‘cycle’ of j -invariants

$$j(E_2) \xrightarrow{\rho_{l_1}} j(E_2)^{l_1} \longrightarrow \dots \xrightarrow{\rho_{l_n}} j(E_2)^{(\alpha)},$$

over $\mathbf{Q}_p$. This computation is carried out with four p -adic digits precision. Update $j(E_2)$ according to the Newton formula above. The value $j(E_3)$ is the four digit approximation of the canonical lift $j(\tilde{E})$.

Step 6. Repeat step 5 with E_2 replaced by E_k with $k = 3, 4$, etc. until we have computed $j(\tilde{E}) \in \mathbf{Z}_p$ with the desired precision. The precision is doubled in each iteration step.

REMARK. There is a different way to lift the cycle $(*)$ of isogenies in step 3. The polynomial $\bar{f}_l$ has a unique Hensel lift to a factor $f_l \in \mathbf{Z}_p[X]$ of the l -th division polynomial Ψ_l of E_1 . This lift is the algorithmic version of the group scheme from section 5.4: every choice of E_1 gives us a subgroup $E_1[l] \subset E_1[l]$. Computing the isogenous curve E_1^l is now easy, since we can apply ‘Vélu’s formulas’ [64]. This approach has the disadvantage that lifting $\bar{f}_l \in \mathbf{F}_p[X]$ to $f_l \in \mathbf{Z}_p[X]$ is rather ‘expensive’. Indeed, the polynomial $\bar{f}_l$ has degree $(l-1)/2$ for $l > 2$. In our approach in step 3, we only perform a simple Hensel lift of a zero of a polynomial of degree $l+1$.

The run time of this algorithm depends heavily on the primes l_i , i.e., on the smoothness properties of (α) . In computing the canonical lift of $\bar{E}/\mathbf{F}_p$, we have the freedom to choose $\alpha \in \mathcal{O}$ ourselves. Subject to the 5 conditions from the beginning of this section, we want (α) to be *smooth*, i.e., the norm $N(\alpha)$ should be smooth.

Write $\alpha = c + d\pi_p$, with $\gcd(c, d) = 1$ and with $d \neq 0$. Here $\pi_p \in \mathbf{Z}[\pi_p] = \tilde{\mathcal{O}}$ is an element of norm p . Condition 3 is satisfied precisely when p does not divide $2d\pi_p$. We conclude that $\alpha/\bar{\alpha} - 1$ will be a p -adic unit for $p > d$.

The following lemma guarantees that there are enough smooth elements α satisfying our conditions.

LEMMA 5.18. *Let $\varepsilon \in (0, 1/2)$ be a real number and let π_p be imaginary quadratic with minimal polynomial $\pi_p^2 - t\pi_p + p = 0$. Put $t^2 - 4p = \tilde{\Delta}$ and $B = \lfloor \exp(\sqrt{\log |\tilde{\Delta}|}) \rfloor$. Let A_ε be the set of $c + d\pi_p \in \mathbf{Z}[\pi_p]$ with $c \in \mathbf{Z}$ and $1 \leq d \leq 2 \exp((\log |\tilde{\Delta}|)^{1/2+\varepsilon})$ satisfying the properties*

- ◊ $|c + \frac{1}{2}dt| \leq |\tilde{\Delta}|^{1/2} \exp((\log |\tilde{\Delta}|)^{1/2+\varepsilon})$
- ◊ c and d are coprime
- ◊ $c + d\pi_p$ and $p\tilde{\Delta}$ are coprime.

If GRH holds true, the fraction of B -smooth elements in A_ε is at least

$$\exp(-2(\log |\tilde{\Delta}|)^{1/2} \log \log |\tilde{\Delta}|)$$

for $|\tilde{\Delta}|$ large enough, depending on ε .

PROOF. This is lemma 2 in [12]. □

An element $\alpha \in \mathbf{Z}[\pi_p]$ satisfying the conditions of lemma 5.18 automatically satisfies the 5 conditions from the beginning of this section.

We find a suitable α by sieving in the set

$$S = \{c + d\pi_p : c, d \in \mathbf{Z}, d \neq 0, (c, d) = 1, c + d\pi_p \text{ and } p\tilde{\Delta} \text{ are coprime}\},$$

where $\tilde{\Delta}$ is the discriminant of $\tilde{\mathcal{O}}$. This is the main probabilistic step in the algorithm. Note that α has norm bounded by $|\tilde{\Delta}|^{1/2} \exp(\log(|\tilde{\Delta}|)^{1/2+\varepsilon})$.

PROOF OF THEOREM 5.3. Fix a real number $\varepsilon \in (0, 1/2)$. We sieve for a smooth element $\alpha \in \mathbf{Z}[\pi_p] \setminus \mathbf{Z}$. Next we apply the algorithm from this section, with this principal ideal (α) , to compute the j -invariant of the canonical lift in k digits accuracy.

It remains to analyse the run time of the algorithm. Searching in

$$\{c + d\pi_p : c, d \in \mathbf{Z}, d \neq 0, (c, d) = 1, c + d\pi_p \text{ and } p\tilde{\Delta} \text{ are coprime}\}$$

for a suitable B -smooth element α takes probabilistic time $O(\exp(\sqrt{\log p} \log \log p)^4)$ by lemma 5.18, with $B = \lfloor \exp(\sqrt{\log p}) \rfloor$. Here, we used the estimate

$$|\tilde{\Delta}| \leq 4p.$$

In the first step of the algorithm we compute the cycle of isogenies over $\mathbf{F}_p$ corresponding to the map ρ_α . We only have to perform ‘simple’ tasks in this step,

like computing a modular polynomial, Euclidean division, computing a root of $\Phi_l(j(\bar{E}), X) \in \mathbf{F}_p[X]$, etc. We compute the cycle in time $O((B^2(\log p)^3)^{1+o(1)})$.

In steps 3, 5 and 6 we lift the cycle to a ‘cycle’ of isogenies over $\mathbf{Z}_p$. Fix an integer $n \in \mathbf{Z}_{>0}$ and assume that we have computed the cycle over $\mathbf{Z}_p/(p^{2^n})$. Lifting the cycle to $\mathbf{Z}_p/(p^{2^{n+1}})$ boils down to evaluating the modular polynomial and some Hensel lifts. We can lift the cycle in time $O(B^2 2^{n+1}(\log p)^2)$.

Combining the sieving step, the computation of the cycle over $\mathbf{F}_p$ and the lifting process, we see that the expected run time is $O(B^{4+\varepsilon}(\log p)^{3+\varepsilon} k \log k)$, which proves the theorem. $\square$

REMARK. Instead of sieving for a smooth element $\alpha \in S$, we can also pick the smallest prime l that splits in the order $\tilde{\mathcal{O}}$ of discriminant $\tilde{\Delta}$. Write $(l) = \tilde{l}\bar{l}$ and let n be the order of $\tilde{l} \in \text{Pic}(\tilde{\mathcal{O}})$. Now let α be a generator of the principal ideal $\tilde{l}^n$. If GRH holds true, the Bach bound yields that l is of size $O((\log |\tilde{\Delta}|)^2)$, but unfortunately we do not have any guarantee that $\alpha/\bar{\alpha} - 1$ is a p -adic unit. In practice this condition never poses a problem. Computing the canonical lift $\tilde{E}$ may take a lot more time however. Indeed, as class groups are ‘often’ cyclic it might very well be that $[\tilde{l}] \in \text{Pic}(\mathcal{O})$ generates the Picard group. The length of the cycle of isogenies over $\mathbf{F}_p$ then becomes $O(|\Delta|^{1/2+o(1)})$, instead of $O((\log |\tilde{\Delta}|)^{1+o(1)})$ for the sieving method.

5.8 Computing the Hilbert class polynomial

Once we have computed one element $j \in \text{Ell}_\Delta(\mathbf{Q}_p)$ with high enough accuracy, it is an easy matter to compute its conjugates under the action of the Picard group $\text{Pic}(\mathcal{O})$. Namely, let $l = \tilde{l}\bar{l}$ be a prime that splits in $\mathcal{O}$. The conjugates of $j \in \text{Ell}_\Delta(\mathbf{Q}_p)$ under the action of $[\tilde{l}], [\bar{l}] \in \text{Pic}(\mathcal{O})$ are the 2 roots of $\Phi_l(j, X) \in \mathbf{Z}_p[X]$. If GRH holds true, we can compute a set of primes S generating $\text{Pic}(\mathcal{O})$ with the property that the largest element of S does not exceed the Bach bound $O((\log |\Delta|)^2)$. We get the following algorithm for computing the Hilbert class polynomial P_Δ .

Algorithm. (*Non-archimedean algorithm*)

Input: a negative discriminant $\Delta < -4$. Output: the Hilbert class polynomial P_Δ .

1. Apply the endomorphism ring algorithm to find a prime p and an ordinary elliptic curve $\bar{E}/\mathbf{F}_p$ with $\text{End}(\bar{E}) = \mathcal{O}_\Delta$.
2. Put $k \leftarrow \left(\frac{\pi \sqrt{|\Delta|}}{\log p} \sum_{[a,b,c] \in \mathcal{F}_\Delta^+} \frac{1}{a} \right) + \log_p \binom{h}{\lfloor h/2 \rfloor}$, with $h = h(\Delta)$.
3. Compute $\mu^{-1}(j(\bar{E})) \in \mathbf{Q}_p$ up to k p -adic digits accuracy using the algorithm in section 5.7.

4. Compute a set C of conjugates of $\mu^{-1}(j(\overline{E}))$ under the action of $\text{Pic}(\mathcal{O})$ in k digits accuracy.
5. Put $P_\Delta = \prod_{j \in C} (X - j) \in (\mathbf{Z}_p/(p^k))[X]$.
6. Lift the coefficients of P_Δ from $\mathbf{Z}_p/(p^k) = \mathbf{Z}/(p^k)$ to $\mathbf{Z}$, where we take the representative between $-p^k/2$ and $p^k/2$. Return $P_\Delta \in \mathbf{Z}[X]$.

THEOREM 5.19. *If GRH holds true, then the non-archimedean algorithm has an expected run time $O(|\Delta|^{1+\varepsilon})$ for every $\varepsilon > 0$.*

PROOF. The run time of step 1 is $O(|\Delta|^{1/2+o(1)})$. To estimate the run time of step 3, we apply theorem 5.3 with the k from step 2. We can compute the j -invariant of the canonical lift of $\overline{E}$ with high enough accuracy in time $O(|\Delta|^{1/2+o(1)})$. Computing the conjugates in step 4 takes time $O(|\Delta|^{1+o(1)})$. Finally, we compute the polynomial in step 5 by employing a ‘divide-and-conquer’ algorithm as in [24, Section 10.2]. $\square$

REMARK. The run time of the non-archimedean approach is the same as the run time for the complex analytic approach from section 3.3. Both run times are in a sense best possible, since just writing down the polynomial $P_\Delta \in \mathbf{Z}[X]$ already takes time $\widetilde{O}(|\Delta|)$. Computer experiments have shown that both methods are equally fast in practice.

► Implementation details

We give some tricks to speed up an implementation of the non-archimedean algorithm. First of all, it is a good idea to precompute a reasonable amount of modular polynomials. Experience has shown that computing the first 25 polynomials, i.e., for primes up to 100, suffices for discriminants down to -10^{12} .

One can save some time in computing the cycles of isogenies over $\mathbf{F}_p$. Let $\overline{E}/\mathbf{F}_p$ be an elliptic curve with $\text{End}(\overline{E}) = \mathcal{O}$ and let $\mathfrak{l} \subset \mathcal{O}$ be of norm $l \neq p$. After we have computed a root $h \in \mathbf{F}_p$ of $\Phi_l(j(\overline{E}), X) \in \mathbf{F}_p[X]$ we have to check if h is the j -invariant of $\overline{E}^{\mathfrak{l}}$, and not the j -invariant of $\overline{E}^{\overline{\mathfrak{l}}}$. In many cases this check can be performed very easily. Namely, suppose that $\mathfrak{l}^2$ divides (α) , i.e., we have to compute the map $\bar{\rho}_{\mathfrak{l}}$ twice. The first time we apply the check proposed at the end of section 5.6 and compute the isogenous curve $\overline{E}^{\mathfrak{l}}$. The j -invariant of $\overline{E}^{\mathfrak{l}^2}$ can now easily be computed. Namely, we compute the 2 roots in $\mathbf{F}_p$ of $\Phi_l(j(\overline{E}^{\mathfrak{l}}), X) \in \mathbf{F}_p[X]$ and note that one of these roots has to be the j -invariant of $\overline{E}^{\overline{\mathfrak{l}}} \cong \overline{E}^{(\mathfrak{l})} = \overline{E}$ and hence we know right away which root is the j -invariant of $\overline{E}^{\mathfrak{l}^2}$.

Finally, the upper bound

$$k = \left(\frac{\pi\sqrt{|\Delta|}}{\log p} \sum_{[a,b,c] \in \mathcal{F}_\Delta^+} \frac{1}{a} \right) + \log_p \binom{h}{\lfloor h/2 \rfloor}$$

for the required precision is somewhat pessimistic. Practical experience has shown that for discriminants down to -10^{12} it suffices to work with

$$k = \frac{\pi\sqrt{|\Delta|}}{\log p} \sum_{[a,b,c] \in \mathcal{F}_\Delta^+} \frac{1}{a} + 10$$

p -adic digits. This is a significant speed up in the practical performance of the algorithm.

5.9 Example

We illustrate the non-archimedean algorithm by computing the Hilbert class polynomial P_Δ for $\Delta = -639 = -3^2 \cdot 71$. First we find a finite field $\mathbf{F}_p$ and an elliptic curve $\overline{E}/\mathbf{F}_p$ with endomorphism ring $\mathcal{O} = \mathcal{O}_\Delta$.

We apply the algorithm from section 5.1. As we have $\Delta \equiv 1 \pmod{8}$, the equation $4p = t^2 - \Delta$ has no solutions with p prime. The smallest integer $t > 0$ for which $(t^2 - 4\Delta)/4$ is prime is $t = 4$, leading to $p = 643$. We fix p for the rest of this section. We apply the naïve algorithm and look for a curve with $p + 1 \pm t$ points. We find that the curve $\overline{E}/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + 89X - 89$$

of j -invariant $j(\overline{E}) = 295 \in \mathbf{F}_p$ has trace of Frobenius 4.

Let $\mathcal{O}_K$ be the maximal order of $K = \mathbf{Q}(\sqrt{\Delta})$. We have inclusions

$$\mathbf{Z}[F_p] \stackrel{2}{\subset} \mathcal{O} \stackrel{3}{\subset} \mathcal{O}_K,$$

and we have to compute the endomorphism ring of $\overline{E}$. The 2-division polynomial $X^3 + 89X - 89 \in \mathbf{F}_p[X]$ splits completely, showing that $\overline{E}$ has CM by $\mathcal{O}$. The prime 3 splits in $\mathcal{O}_K$. If $\overline{E}$ has CM by $\mathcal{O}_K$, the modular polynomial $\Phi_3(j(\overline{E}), X) \in \mathbf{F}_p[X]$ has 4 roots, cf. corollary 5.12. We compute $\gcd(\Phi_3(j(\overline{E}), X), X^p - X) = X - 429 \in \mathbf{F}_p[X]$. We conclude that $\overline{E}$ does not have CM by $\mathcal{O}_K$ and hence has endomorphism ring $\mathcal{O}$.

We need to compute the canonical lift $\tilde{E}/\mathbf{Q}_p$ up to k p -adic digits accuracy, with

$$k = \frac{\pi\sqrt{|\Delta|}}{\log p} \sum_{[a,b,c] \in \mathcal{F}_\Delta^+} \frac{1}{a}.$$

The Picard group $\text{Pic}(\mathcal{O})$ has order 14 and representing the elements as binary quadratic forms as in section 3.3, we find $k \approx 44$. We will compute $j(\tilde{E}) \in \mathbf{Q}_p$ up to 45 p -adic digits precision.

As smooth element $\alpha \in \mathcal{O} \setminus \mathbf{Z}$ for the map $\rho_\alpha : X_\Delta(\mathbf{C}_p) \rightarrow X_\Delta(\mathbf{C}_p)$ we will use $\alpha = \pi_p - 108$ of norm $11875 = 5^4 \cdot 19$. Here, $\pi_p = \frac{4+\sqrt{\Delta}}{2}$ is an element of norm p . We factor

$$(\alpha) = \mathfrak{p}_5^4 \cdot \mathfrak{p}_{19} = (5, \pi_p - 3)^4 \cdot (19, \pi_p - 13).$$

We compute the action of the prime ideal $\mathfrak{p}_5$ on $j(\bar{E}) \in \mathbf{F}_p$. If we evaluate the modular polynomial $\Phi_5(X, Y) \in \mathbf{F}_p[X, Y]$ in $X = j(\bar{E}) = 295$ we get a polynomial that has 2 roots in $\mathbf{F}_p$, namely 449 and 532. From this we deduce that $\mathfrak{p}_5$ sends $j(\bar{E})$ to one of these roots. We do not know which one yet.

Let $\bar{E}/C$ have j -invariant $449 \in \mathbf{F}_p$, corresponding to a cyclic subgroup $C \subset \bar{E}[5]$. We either have $C = \bar{E}[\mathfrak{p}_5]$ or $C = \bar{E}[\bar{\mathfrak{p}}_5]$. As in section 5.6, we compute the Weierstraß equation

$$Y^2 = X^3 + 390X + 466$$

for $\bar{E}/C$. We get the x -coordinates of the points in C as zeroes of

$$\bar{f}_C = X^2 + 614X + 471 \in \mathbf{F}_p[X].$$

The eigenvalue for the action of Frobenius on the torsion $\bar{E}[\mathfrak{p}_5]$ is $3 \in \mathbf{F}_5$. We now check whether

$$(X^p, Y^p) = 3 \cdot (X, Y)$$

holds for the points in C , i.e., we compute both (X^p, Y^p) and $3 \cdot (X, Y)$ in the ring

$$\mathbf{F}_p[X, Y]/(\bar{f}_C, Y^2 - X^3 - 89X + 89).$$

Here, the $\cdot$ means adding *on the curve*. It turns out that (X^p, Y^p) and $3 \cdot (X, Y)$ are the same. We deduce that we have $j(\bar{E})^{\mathfrak{p}_5} = 449 \in \mathbf{F}_p$.

The action of $\mathfrak{p}_5$ on the j -invariant $449 \in \mathbf{F}_p$ is now easier to compute. The polynomial $\Phi_5(449, X) \in \mathbf{F}_p[X]$ has 2 roots in $\mathbf{F}_p$, but one of these roots corresponds to the action of $\bar{\mathfrak{p}}_5$ and is therefore equal to $j(\bar{E})$. We pick the other root $73 \in \mathbf{F}_p$. If we compute the entire cycle of j -invariants over $\mathbf{F}_p$, we get

$$295 \xrightarrow{\mathfrak{p}_5} 449 \xrightarrow{\mathfrak{p}_5} 73 \xrightarrow{\mathfrak{p}_5} 55 \xrightarrow{\mathfrak{p}_5} 328 \xrightarrow{\mathfrak{p}_{19}} 295.$$

We knew beforehand that this cycle is closed, since we know that (α) acts trivially on $j(\bar{E})$.

We now lift $\bar{E}/\mathbf{F}_p$ to $E_1/\mathbf{Q}_p$ by lifting the coefficients of its Weierstraß equation arbitrarily. The polynomial $\Phi_5(j(E_1), X) \in \mathbf{Z}_p[X]$ has exactly 2 roots, one of which reduces to $449 \in \mathbf{F}_p$. Taking the lift $E_1/\mathbf{Q}_p$ defined by $Y^2 = X^3 + 89X - 89$ of j -invariant $295 - 233p + O(p^2) \in \mathbf{Q}_p$, we compute the ‘cycle’

$$295 - 233p \xrightarrow{\mathfrak{p}_5} 449 + 296p \xrightarrow{\mathfrak{p}_5} 73 - 236p \xrightarrow{\mathfrak{p}_5} 55 + 155p \xrightarrow{\mathfrak{p}_5} 328 + 131p \xrightarrow{\mathfrak{p}_{19}} 295 - 236p$$

over $\mathbf{Q}_p$. We update $j(E_1)$ according to the ‘Newton formula’

$$j(E_{k+1}) = j(E_k) - \frac{\rho_\alpha(j(E_k)) - j(E_k)}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1}$$

and find that $j(E_2) = 295 - 155p$ is the two digit approximation of the j -invariant of the canonical lift $\tilde{E}/\mathbf{Q}_p$.

Starting from $j(E_2)$, we now lift the cycle to four p -adic digits precision, compute $j(E_3)$ from this, and so on. We obtain

$$\begin{aligned} j(\tilde{E}) &= 295 + O(p) \\ &= 295 - 155p + O(p^2) \\ &= 295 - 155p + 195p^2 + 287p^3 + O(p^4) \\ &= 295 - 155p + 195p^2 + 287p^3 - 153p^4 + 245p^5 + 272p^6 + 298p^7 + O(p^8). \\ &= 295 - 155p + 195p^2 + 287p^3 - 153p^4 + 245p^5 + 272p^6 + 298p^7 - 277p^8 \\ &\quad + 170p^9 - 123p^{10} - 86p^{11} - 165p^{12} - 115p^{13} + 195p^{14} + 56p^{15} + O(p^{16}). \end{aligned}$$

We continue this process until we have computed the canonical lift in 45 p -adic digits accuracy.

Next, we compute the conjugates of $j(\tilde{E})$ under $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$. The Picard group $\text{Pic}(\mathcal{O})$ is cyclic of order 14 and is generated by a prime of norm 5. We compute the conjugates of $j(\tilde{E})$ by employing the modular polynomial Φ_5 : the roots of $\Phi_5(j(\tilde{E}), X) \in \mathbf{Z}_p[X]$ give us the conjugates $j(\tilde{E})^{\mathfrak{p}_5}$ and $j(\tilde{E})^{\bar{\mathfrak{p}}_5}$, etc. In the end we expand the degree 14 polynomial

$$P_{-639} = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - j(\tilde{E})^I) \in \mathbf{Z}[X].$$

The polynomial P_Δ has coefficients up to 126 decimal digits.

6

Class invariants

6.1 Introduction

In section 3.3, we briefly discussed the classical algorithm to compute, on input of a negative discriminant $\Delta < 0$, the Hilbert class polynomial P_Δ for the ring class field $H_{\mathcal{O}}$ corresponding to the order $\mathcal{O}$ of discriminant Δ . The run time is $O(|\Delta|^{1+\varepsilon})$ for every $\varepsilon > 0$. In chapter 5 we gave a non-archimedean algorithm to compute P_Δ that has the same asymptotic run time.

As noted in chapter 3, *any* algorithm that computes P_Δ will have a run time that is exponential in $\log |\Delta|$. Indeed, the final step of any algorithm will be writing down the answer. The degree of P_Δ equals the class number $h(\Delta)$ and by the Brauer-Siegel theorem, $h(\Delta)$ grows asymptotically like $|\Delta|^{1/2+o(1)}$. Hence any algorithm has to write down roughly $|\Delta|^{1/2}$ coefficients.

A serious drawback of computing the Hilbert class polynomial P_Δ is that the coefficients are *huge*. Not only do they grow exponentially in size for $|\Delta| \rightarrow \infty$, but also for moderately small discriminants, the coefficients are massive. As an example, consider the polynomial for $\Delta = -71$:

$$\begin{aligned} P_{-71} = & X^7 + 313645809715X^6 - 3091990138604570X^5 \\ & + 98394038810047812049302X^4 \\ & - 823534263439730779968091389X^3 \\ & + 5138800366453976780323726329446X^2 \\ & - 425319473946139603274605151187659X \\ & + 737707086760731113357714241006081263 \in \mathbf{Z}[X]. \end{aligned}$$

With modern computers we can only compute P_Δ for $|\Delta|$ at most 10^7 in a reasonable amount of time. This is quite unsatisfactory.

History tells us that we should be able to do better. In his *Lehrbuch der Algebra* (1908) [65], Weber explains that function values of ‘smaller’ functions than the j -

function sometimes also generate the ring class field $H_{\mathcal{O}}$. For example, the j -function has a holomorphic cube root $\gamma_2 : \mathbf{H} \rightarrow \mathbf{C}$ with integral Fourier expansion. In §125 we read:

$\gamma_2(\omega)$ ist eine Klasseninvariante, wenn ω die Wurzel einer quadratischen Form ist, deren Diskriminante und erster Koeffizient durch 3 nicht teilbar sind, während der mittlere Koeffizient durch 3 teilbar ist.

That $\gamma_2(\omega)$ is a class invariant means that $\gamma_2(\omega)$ generates the ring class field of the order $\mathbf{Z}[\omega]$. A more general definition of class invariant will be given in section 6.3.

From Weber's remark we see that for discriminants Δ with $3 \nmid \Delta$, we can also use the function γ_2 to generate the ring class field. The minimal polynomial of $\gamma_2(\frac{-3+\sqrt{-71}}{2})$ is

$$P_{-71}^{\gamma_2} = X^7 + 6745X^6 - 327467X^5 + 51857115X^4 + 2319299751X^3 \\ + 41264582513X^2 - 307873876442X + 903568991567 \in \mathbf{Z}[X].$$

The logarithmic height of the coefficients of $P_{-71}^{\gamma_2}$ is only one third of the height of P_{-71} . This should come as no surprise, since we have $\gamma_2^3 = j$. We can even do better in this case however. In §127, Weber introduces a function $\mathfrak{f} : \mathbf{H} \rightarrow \mathbf{C}$ with the property that $\mathfrak{f}(\omega)$ generates the ring class field of $\mathbf{Z}[\omega]$ for all quadratic orders $\mathbf{Z}[\omega]$ in which 3 is unramified and 2 splits completely. For $\Delta = -71$, a root of the polynomial

$$P_{-71}^{\mathfrak{f}} = X^7 + X^6 - X^5 - X^4 - X^3 + X^2 + 2X - 1 \in \mathbf{Z}[X]$$

generates the Hilbert class field of $\mathbf{Q}(\sqrt{-71})$. Reading Weber, one finds many theorems, properties and questions on class invariants. It is not always clear whether his statements are actually theorems or just observations.

We need more theory than described in Weber's Lehrbuch to be able to treat class invariants in a systematic and algorithmic way. Our main tool will be Shimura's reciprocity law from 1971, which we explain in sections 6.4 and 6.5. Just as for the algorithms to compute the Hilbert class polynomial P_{Δ} , we distinguish two cases: the complex analytic and the non-archimedean setting.

In the complex analytic setting, we follow Gee [26] and Stevenhagen [59]. We explain how Shimura reciprocity allows us to compute 'small' polynomials, like $P_{-71}^{\mathfrak{f}}$, that generate the ring class field.

If we work over the non-archimedean field $\mathbf{Q}_p$, Shimura reciprocity does not suffice. In sections 6.6–6.9 we explain how one can work with class invariants in this case.

6.2 The modular function field

The modular group $\mathrm{SL}_2(\mathbf{Z})$ acts in a natural way on the complex upper half plane $\mathbf{H}$ and its completion $\bar{\mathbf{H}} = \mathbf{H} \cup \mathbf{P}^1(\mathbf{Q})$. A matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z})$ acts via

$$Az = \frac{az + b}{cz + d} \quad (z \in \bar{\mathbf{H}}).$$

The quotient $\mathrm{SL}_2(\mathbf{Z}) \backslash \bar{\mathbf{H}}$ has the structure of a compact Riemann surface and the modular j -function gives an isomorphism

$$j : \mathrm{SL}_2(\mathbf{Z}) \backslash \bar{\mathbf{H}} \xrightarrow{\sim} \mathbf{P}^1(\mathbf{C}).$$

Thinking geometrically, we interpret $\mathrm{SL}_2(\mathbf{Z}) \backslash \bar{\mathbf{H}}$ as the modular curve $X(1)$ as in section 5.3. The elements of the function field $F_{1,\mathbf{C}}$ of the curve $X(1)$ over $\mathbf{C}$ are called *modular functions* over $\mathbf{C}$ of level 1. It is well known [37, Theorem 6.1] that we have $F_{1,\mathbf{C}} = \mathbf{C}(j)$.

In order to define modular functions of higher level, we first define the principal congruence subgroup

$$\Gamma(N) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z}) \mid a \equiv d \equiv 1 \pmod{N}, b \equiv c \equiv 0 \pmod{N} \right\}$$

of $\mathrm{SL}_2(\mathbf{Z})$ for integers $N \in \mathbf{Z}_{>0}$. Another way of phrasing this, is saying that $\Gamma(N)$ is the kernel of the natural reduction map $\mathrm{SL}_2(\mathbf{Z}) \rightarrow \mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})$. Here, $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})$ is the group of matrices with coefficients in the ring $\mathbf{Z}/N\mathbf{Z}$ and with determinant $1 \in \mathbf{Z}/N\mathbf{Z}$. A standard argument as in [37, Section 6.1] shows that we have an exact sequence

$$0 \longrightarrow \Gamma(N) \longrightarrow \mathrm{SL}_2(\mathbf{Z}) \longrightarrow \mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z}) \longrightarrow 0.$$

The quotient $\Gamma(N) \backslash \bar{\mathbf{H}}$ has the structure of a compact Riemann surface, and as such, it is isomorphic to the modular curve $X(N)$ over $\mathbf{C}$, cf. section 5.3. We let $F_{N,\mathbf{C}}$ be the function field of the curve $X(N)$. The elements of $F_{N,\mathbf{C}}$ are called *modular functions* over $\mathbf{C}$ of level N . Explicitly, a modular function of level N is a meromorphic function $f : \bar{\mathbf{H}} \rightarrow \mathbf{P}^1(\mathbf{C})$ that is invariant under $\Gamma(N) \subseteq \mathrm{SL}_2(\mathbf{Z})$ for some $N \geq 1$. Since we have $\begin{pmatrix} 1 & N \\ 0 & 1 \end{pmatrix} \in \Gamma(N)$, the function f is invariant under $\tau \mapsto \tau + N$. Hence, f is periodic and has a Fourier expansion in $q^{1/N}$ with $q = \exp(2\pi i\tau)$.

The natural map $X(N) \rightarrow X(1)$ is a Galois cover with group

$$\mathrm{SL}_2(\mathbf{Z})/(\pm 1 \cdot \Gamma(N)) = \mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})/\pm 1.$$

Equivalently, the extension $F_{N,\mathbf{C}}/F_{1,\mathbf{C}}$ is Galois with group $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})/\pm 1$, cf. [37, Theorem 6.2].

We want modular functions to yield algebraic values, so we make the step from $\mathbf{C}$ to a number field as base field. The modular curve $X(N)$ can be defined over the cyclotomic field $\mathbf{Q}(\zeta_N)$, where ζ_N is a primitive N -th root of unity. We define F_N to be the function field of $X(N)$ over $\mathbf{Q}(\zeta_N)$. We have $F_1 = \mathbf{Q}(j)$. Elements of F_N are called *modular functions* of level N . Hence, if we use the term ‘modular function’ without any condition, we always mean a modular function over $\mathbf{Q}$ (and not over $\mathbf{C}$). Modular functions are functions on $X(N)_{\mathbf{Q}(\zeta_N)}$, i.e., functions from $F_{N,\mathbf{C}}$ having Fourier coefficients in $\mathbf{Q}(\zeta_N)$.

We describe the Galois group of the extension F_N/F_1 . For $d \in (\mathbf{Z}/N\mathbf{Z})^*$, let $\sigma_d \in \mathrm{Gal}(\mathbf{Q}(\zeta_N)/\mathbf{Q})$ be the field automorphism that raises ζ_N to the d -th power. The algebraic closure of $\mathbf{Q}$ inside F_1 equals $\mathbf{Q}$, and we have a natural isomorphism

$$\mathrm{Gal}(F_1(\zeta_N)/F_1) \cong \mathrm{Gal}(\mathbf{Q}(\zeta_N)/\mathbf{Q}) \cong (\mathbf{Z}/N\mathbf{Z})^*,$$

which we can lift to F_N in the following way. For $f \in F_N$ with Fourier expansion $\sum_k c_k \cdot q^{k/N} \in \mathbf{Q}(\zeta_N)((q^{1/N}))$ we define $f^{\sigma_d} = \sum_k \sigma_d(c_k) \cdot q^{k/N}$. The function f^{σ_d} is again contained in F_N . We get a group action of $(\mathbf{Z}/N\mathbf{Z})^*$ on the field F_N . We have $\mathrm{Gal}(F_N/F_1(\zeta_N)) \cong \mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})/\pm 1$, and we can describe $\mathrm{Gal}(F_N/F_1)$ as a semidirect product

$$(\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})/\pm 1) \rtimes (\mathbf{Z}/N\mathbf{Z})^* \cong \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})/\pm 1.$$

Here we embed $(\mathbf{Z}/N\mathbf{Z})^*$ in $\mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$ as the subgroup

$$\left\{ \begin{pmatrix} 1 & 0 \\ 0 & d \end{pmatrix} \mid d \in (\mathbf{Z}/N\mathbf{Z})^* \right\} \subset \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z}).$$

Define the *modular function field* $\mathcal{F}$ as the union $\mathcal{F} = \bigcup_{N \geq 1} F_N$. The extension $\mathcal{F}/F_1$ is an infinite Galois extension. The discussion above shows that we have an exact sequence

$$1 \longrightarrow \{\pm 1\} \longrightarrow \mathrm{GL}_2(\widehat{\mathbf{Z}}) \longrightarrow \mathrm{Gal}(\mathcal{F}/F_1) \longrightarrow 1.$$

REMARK. We can also give generators for the fields $F_{N,\mathbf{C}}$ and F_N . Define the function f by

$$f(w, \tau) = -2^7 3^5 \cdot \frac{g_2(\tau)g_3(\tau)}{\Delta(\tau)} \wp(w; \langle 1, \tau \rangle)$$

for $w \in \mathbf{C}$ and $\tau \in \mathbf{H}$. Here, $\wp(\cdot; \langle 1, \tau \rangle)$ is the Weierstraß $\wp$ -function associated to the lattice $\mathbf{Z} + \mathbf{Z} \cdot \tau$. The function f is called the *first Weber function*. Fix an integer $N > 1$ and for $r, s \in \frac{1}{N}\mathbf{Z}/\mathbf{Z}$, not both 0, define the *Fricke function* $f_{r,s}$ of level N by

$$f_{r,s}(\tau) = f(rN + s, \tau).$$

The Fourier coefficients of $f_{r,s}$ are contained in $\mathbf{Q}(\zeta_N)$. If we fix τ and let r, s vary over $\frac{1}{N}\mathbf{Z}/\mathbf{Z}$, not both equal to 0, we get the normalized x -coordinates of the $N^2 - 1$ non-trivial points of order N of the complex elliptic curve $\mathbf{C}/(\mathbf{Z} + \mathbf{Z} \cdot \tau)$.

THEOREM 6.1. *We have*

$$F_{N,\mathbf{C}} = \mathbf{C}(j, f_{r,s} \mid r, s \in \frac{1}{N}\mathbf{Z}/\mathbf{Z}, \text{ not both } 0)$$

and

$$F_N = \mathbf{Q}(j, f_{r,s} \mid r, s \in \frac{1}{N}\mathbf{Z}/\mathbf{Z}, \text{ not both } 0)$$

PROOF. [37, Theorem 6.2] and [37, beginning of section 6.3].

6.3 Class invariants

Let K be an imaginary quadratic number field. We define the ring of *finite* K -adeles $\widehat{K} = K \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$. Let $\widehat{K}^* = (K \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}})^*$ be the unit group of $\widehat{K}$. The group $\widehat{K}^*$ is called the group of *finite* K -ideles. We have

$$\widehat{K}^* = \prod'_{\mathfrak{p} \text{ finite}} K_{\mathfrak{p}}^*,$$

where the restricted product is taken with respect to the unit groups of the maximal order of the completion $K_{\mathfrak{p}}$. Explicitly, for an element $\alpha = (\alpha_{\mathfrak{p}})_{\mathfrak{p}} \in \widehat{K}^*$, with $\mathfrak{p}$ ranging over the finite primes of K , nearly all components $\alpha_{\mathfrak{p}}$ are units.

Let K_{ab} be the maximal abelian extension of K . Class field theory tells us that $\text{Gal}(K_{\text{ab}}/K)$ can be described by an exact sequence

$$1 \longrightarrow K^* \longrightarrow \widehat{K}^* \xrightarrow{\text{Artin}} \text{Gal}(K_{\text{ab}}/K) \longrightarrow 1.$$

Let

$$\widehat{\mathcal{O}} = \varprojlim_N (\mathcal{O}/N\mathcal{O}) = \mathcal{O} \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$$

be the profinite completion of a not necessarily maximal order $\mathcal{O}$ of K . We have an inclusion $\widehat{\mathcal{O}}^* \subset \widehat{K}^*$. Let $H_{\mathcal{O}}$ be the ring class field corresponding to $\mathcal{O}$. The

unit group $\widehat{\mathcal{O}}^*$ maps onto $\text{Gal}(K_{\text{ab}}/H_{\mathcal{O}})$ under the Artin map. We therefore have an exact sequence

$$1 \longrightarrow \mathcal{O}^* \longrightarrow \widehat{\mathcal{O}}^* \longrightarrow \text{Gal}(K_{\text{ab}}/H_{\mathcal{O}}) \longrightarrow 1.$$

We obtain K_{ab} as the union of finite extensions $H_{N,\mathcal{O}}$ corresponding to the finite quotients

$$\widehat{\mathcal{O}}^* \twoheadrightarrow (\widehat{\mathcal{O}}/N\widehat{\mathcal{O}})^* = (\mathcal{O}/N\mathcal{O})^*.$$

The field $H_{N,\mathcal{O}}$ is called the *ray class field of conductor N for the order $\mathcal{O}$* . The Artin map gives an isomorphism

$$(\mathcal{O}/N\mathcal{O})^*/\text{im}[\mathcal{O}^*] \xrightarrow{\sim} \text{Gal}(H_{N,\mathcal{O}}/H_{\mathcal{O}}).$$

If $\mathcal{O}$ is the maximal order of K , the field $H_{N,\mathcal{O}}$ is the ray class field of conductor N of K .

We use the unique infinite prime of K to view $\mathbf{C}$ as the archimedean completion of K . The following theorem, known as the second main theorem of complex multiplication, gives the link between modular functions and the ray class fields of an order.

THEOREM 6.2. *Let f be a modular function of level $N \geq 1$ and let $\mathcal{O}$ be an imaginary quadratic order. Write $\mathcal{O} = \mathbf{Z}[\tau]$ with $\tau \in \mathbf{H}$. Then we have*

$$f(\tau) \in H_{N,\mathcal{O}}$$

and

$$H_{N,\mathcal{O}} = K(g(\tau) \mid g \in F_N, g(\tau) \neq \infty).$$

PROOF. The first statement follows directly from the second. The second statement can be found for instance in [37, Chapter 10]. $\square$

The first main theorem of complex multiplication, theorem 3.4, is a direct consequence of theorem 6.2. Indeed, the j -function is modular of level 1 and without poles on $\mathbf{H}$, and we have an equality $H_{1,\mathcal{O}} = H_{\mathcal{O}}$.

Theorem 6.2 tells us that if we evaluate a modular function f of level N in a generator τ for an order $\mathcal{O}$, we end up in the ray class field of conductor N for this order. It may of course happen that $f(\tau)$ already lives in a smaller field than $H_{N,\mathcal{O}}$. Following Weber, we call $f(\tau)$ a *class invariant* if we have

$$K(f(\tau)) = K(j(\tau)).$$

The functions γ_2 and $\mathfrak{f}$ mentioned in the introduction are examples of functions that yield class invariants when evaluated in appropriate points $\tau \in \mathbf{H}$. The logarithmic heights of the Fourier coefficients of these functions are a constant factor smaller than the height of the Fourier coefficients of j . The minimal polynomial of e.g. a class invariant $\mathfrak{f}(\tau)$ has smaller coefficients than the Hilbert class polynomial. We are therefore led to consider the following problem:

PROBLEM. *Given an imaginary quadratic order $\mathcal{O} = \mathbf{Z}[\tau]$, find a modular function f with the property that $f(\tau)$ is a class invariant and with the property that the minimal polynomial of $f(\tau)$ lives in $\mathbf{Z}[X]$ and has ‘smaller’ coefficients than the Hilbert class polynomial. Moreover, if $f(\tau)$ is a class invariant, compute the conjugates of $f(\tau)$ under the Galois group $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$.*

This problem is not completely well-posed, since there is no definition of the word ‘smaller’. In the remainder of this chapter we will show how to use various modular functions, for which we gain a *constant factor* in the logarithmic height of the coefficients. As an example, for γ_2 we gain a constant factor 3, and for Weber’s function $\mathfrak{f}$ we gain a factor 72.

6.4 Shimura reciprocity over the ring class field

Although Weber had already partially ‘solved’ the problem from section 6.3, his method consisted mostly of clever tricks and can hardly be considered to be a ‘systematic way’. For more than 50 years, all theory regarding class invariants relied on Weber’s *Lehrbuch*. This situation changed with the appearance of Shimura’s textbook [56], and more specifically with his reciprocity law.

Shimura reciprocity provides a link between the exact sequences describing the Galois groups $\text{Gal}(\mathcal{F}/F_1)$ and $\text{Gal}(K_{\text{ab}}/H_{\mathcal{O}})$. Write $\mathcal{O} = \mathbf{Z}[\tau]$ with τ an algebraic integer. We will define a map g_{τ} connecting the two exact rows in the following diagram.

$$\begin{array}{ccccccc} 1 & \longrightarrow & \mathcal{O}^* & \longrightarrow & \widehat{\mathcal{O}}^* & \xrightarrow{\text{Artin}} & \text{Gal}(K_{\text{ab}}/H_{\mathcal{O}}) \longrightarrow 1 \\ & & & & \downarrow g_{\tau} & & \\ 1 & \longrightarrow & \{\pm 1\} & \longrightarrow & \text{GL}_2(\widehat{\mathbf{Z}}) & \longrightarrow & \text{Gal}(\mathcal{F}/\mathbf{Q}(j)) \longrightarrow 1 \end{array} \quad (6.1)$$

The map $g = g_{\tau} : \widehat{\mathcal{O}}^* \rightarrow \text{GL}_2(\widehat{\mathbf{Z}})$ sends an idele $x \in \widehat{\mathcal{O}}^*$ to the *transpose* of the matrix representing multiplication by x on the free $\widehat{\mathbf{Z}}$ -module $\widehat{\mathcal{O}} = \widehat{\mathbf{Z}} \cdot \tau + \widehat{\mathbf{Z}}$ with

respect to the basis $[\tau, 1]$. If τ has minimal polynomial $X^2 + BX + C \in \mathbf{Z}[X]$, we have

$$g_\tau : x = s\tau + t \mapsto \begin{pmatrix} t - Bs & -Cs \\ s & t \end{pmatrix},$$

as may be checked easily. We get an action of the group $\widehat{\mathcal{O}}^*$ on the modular function field $\mathcal{F}$.

Let $f \in \mathcal{F}$ be a modular function of level $N \geq 1$. Shimura's reciprocity law states that the Galois conjugate $(f(\tau))^x$ of $f(\tau)$ under the Artin symbol $x \in \text{Gal}(K_{\text{ab}}/H_{\mathcal{O}})$ can be computed via the reciprocity relation

$$(f(\tau))^x = (f^{g_\tau(x^{-1})})(\tau),$$

cf. [56, Theorem 6.31]. If the extension $\mathcal{F}/\mathbf{Q}(f)$ is Galois, we have the fundamental equivalence

$$(f(\tau))^x = f(\tau) \iff f^{g_\tau(x)} = f.$$

The implication $\Leftarrow$ is immediate from the reciprocity relation. The other implication requires the hypothesis and an additional argument [56, Proposition 6.33].

Suppose that $\mathcal{F}/\mathbf{Q}(f)$ is Galois. If we want to know whether a given value $f(\tau)$ is a class invariant, we need to check whether all $x \in \widehat{\mathcal{O}}^*$ act trivially on $f(\tau)$. Shimura reciprocity law tells us that this is equivalent with checking whether all $g_\tau(x) \in \text{GL}_2(\widehat{\mathbf{Z}})$ fix the function f . The infinite groups $\widehat{\mathcal{O}}^*$ and $\text{GL}_2(\widehat{\mathbf{Z}})$ are not directly suited for explicit computations. However, theorem 6.2 tells us that for f of level $N \geq 1$, the function value $f(\tau)$ lives in the ray class field $H_{N,\mathcal{O}}$ of conductor N for the order $\mathcal{O} = \mathbf{Z}[\tau]$. Hence, the action of $\widehat{\mathcal{O}}^*$ on $f(\tau)$ can be computed via the *finite* quotient $(\widehat{\mathcal{O}}/N\widehat{\mathcal{O}})^* = (\mathcal{O}/N\mathcal{O})^*$. We obtain a diagram with finite groups and exact rows.

$$\begin{array}{ccccccc} \mathcal{O}^* & \longrightarrow & (\mathcal{O}/N\mathcal{O})^* & \xrightarrow{\text{Artin}} & \text{Gal}(H_{N,\mathcal{O}}/H_{\mathcal{O}}) & \longrightarrow & 1 \\ & & \downarrow \bar{g}_\tau & & & & \\ \{\pm 1\} & \longrightarrow & \text{GL}_2(\mathbf{Z}/N\mathbf{Z}) & \longrightarrow & \text{Gal}(F_N/\mathbf{Q}(j)) & \longrightarrow & 1 \end{array}$$

We compute generators $x_1, \dots, x_k$ for $(\mathcal{O}/N\mathcal{O})^*$ and map them to $\text{GL}_2(\mathbf{Z}/N\mathbf{Z})$ using the map $\bar{g}_\tau$. The value $f(\tau)$ is contained in the ring class field $H_{\mathcal{O}}$ if and only if $\bar{g}_\tau(x_1), \dots, \bar{g}_\tau(x_k)$ act trivially on f . If we for instance also know that there is an inclusion $\mathbf{Q}(j) \subseteq \mathbf{Q}(f)$, then $f(\tau)$ is also a class invariant if $\bar{g}_\tau(x_1), \dots, \bar{g}_\tau(x_k)$ act trivially on f .

All that we need to know is the explicit action of $\text{GL}_2(\mathbf{Z}/N\mathbf{Z})$ on f . It suffices to know the action of the ‘standard generators’ $S, T \in \text{SL}_2(\mathbf{Z}/N\mathbf{Z})$ on f and the action of $(\mathbf{Z}/N\mathbf{Z})^*$ on the Fourier coefficients of f . Using Shimura reciprocity, it is

a rather mechanical process to check whether $f(\tau)$ is a class invariant or not. In [26] there is a large number of examples. We give two examples of statements proven there.

The j -function has a holomorphic cube root $\gamma_2 : \mathbf{H} \rightarrow \mathbf{C}$ with integral Fourier expansion. The function γ_2 is modular of level 3 and the matrices $S, T \in \mathrm{SL}_2(\mathbf{Z})$ act via

$$\gamma_2 \circ S = \gamma_2 \qquad \gamma_2 \circ T = \zeta_3^{-1} \gamma_2.$$

Let $\mathcal{O}$ be the maximal order of $K = \mathbf{Q}(\sqrt{D})$ and let $\tau = \frac{-B+\sqrt{D}}{2}$ be a generator of $\mathcal{O}$ as a $\mathbf{Z}$ -algebra. Theorem 1.10 of [26] states:

$$3 \nmid D \implies \zeta_3^B \gamma_2(\tau) \text{ is a class invariant and } P_D^{\gamma_2} \in \mathbf{Z}[X].$$

This is exactly Weber's result that we quoted in section 6.1.

For the second example, define the Dedekind η -function by the product expansion

$$\eta(z) = q^{1/24} \prod_{n=1}^{\infty} (1 - q^n), \quad \text{in } q = \exp(2\pi iz).$$

The η -function is holomorphic and non-zero for $z \in \mathbf{H}$. Define the Weber functions

$$\mathfrak{f}(z) = \zeta_{48}^{-1} \cdot \frac{\eta(\frac{z+1}{2})}{\eta(z)}, \quad \mathfrak{f}_1(z) = \frac{\eta(\frac{z}{2})}{\eta(z)}, \quad \mathfrak{f}_2(z) = \sqrt{2} \cdot \frac{\eta(2z)}{\eta(z)},$$

where $\zeta_{48} = \exp(2\pi i/48)$ is a primitive 48-th root of unity. The Weber functions are modular of level 48. The notation suggests that they are conjugates over $\mathbf{Q}(j)$, but this is *not* the case. In fact, $\mathfrak{f}$ satisfies the relation

$$(X^{24} - 16)^3 - jX^{24} \in \mathbf{Z}[j, X]$$

and both $\mathfrak{f}_1$ and $\mathfrak{f}_2$ satisfy

$$(X^{24} + 16)^3 - jX^{24} \in \mathbf{Z}[j, X].$$

The notation $\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2$ is inspired by the fact that $\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2$ have integral Fourier expansion. We stick to this perhaps confusing historical notation. The matrices $S, T \in \mathrm{SL}_2(\mathbf{Z})$ act via

$$(\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2) \circ S = (\mathfrak{f}, \mathfrak{f}_2, \mathfrak{f}_1), \quad (\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2) \circ T = (\zeta_{48}^{-1} \mathfrak{f}_1, \zeta_{48}^{-1} \mathfrak{f}, \zeta_{48}^2 \mathfrak{f}_2).$$

Let $\mathcal{O}$ be the maximal order of $K = \mathbf{Q}(\sqrt{D})$ and suppose that $D = \text{disc}(K)$ is congruent to 1 modulo 8. Let $\tau = \frac{-1+\sqrt{D}}{2}$ be a generator for $\mathcal{O}$ as a $\mathbf{Z}$ -algebra. Theorem 1.15 of [26] states:

$$3 \nmid D \implies \zeta_{48} f_2(\tau) \text{ is a class invariant and } P_D^{\zeta_{48} f_2} \in \mathbf{Z}[X].$$

We see that in both examples there is a condition on the discriminant of the number field. In the second example, the condition $3 \nmid D$ can be removed by replacing $\zeta_{48} f_2$ by $\zeta_{16} f_2^3$.

To further illustrate Shimura reciprocity, we prove a theorem concerning Weber functions and a discriminant D with $D \equiv 5 \pmod{8}$. First we state a lemma that will help us with the explicit computations. For a matrix $A \in \text{SL}_2(\mathbf{Z})$, we denote by $A_N \in \text{SL}_2(\mathbf{Z}/N\mathbf{Z})$ its reduction modulo N .

LEMMA 6.3. *Let N be a prime power, and let $\begin{pmatrix} a & b \\ c & d \end{pmatrix}_N \in \text{SL}_2(\mathbf{Z}/N\mathbf{Z})$ be a matrix. For $(c, N) = 1$, define $y = (a+1)c^{-1} \in \mathbf{Z}/N\mathbf{Z}$. Otherwise, for $(a, N) = 1$, define $z = (c+1)a^{-1} \in \mathbf{Z}/N\mathbf{Z}$. Then $\begin{pmatrix} a & b \\ c & d \end{pmatrix}_N$ can be written as*

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{cases} (T^y S T^c S T^{dy-b})_N & \text{for } (c, N) = 1 \\ (S T^{-z} S T^{-a} S T^{bz-d})_N & \text{for } (a, N) = 1. \end{cases}$$

PROOF. This is lemma 6 in [26]. □

THEOREM 6.4. *Let K be an imaginary quadratic number field with discriminant $D \equiv 5 \pmod{8}$. Let $\mathcal{O} = \mathbf{Z}[\sqrt{D}]$ be the order of index 2 in the maximal order of K . Then the following holds:*

$$\begin{aligned} 3 \nmid D &\implies f(\sqrt{D}) \text{ is a class invariant} \\ 3 \mid D &\implies f(\sqrt{D})^3 \text{ is a class invariant.} \end{aligned}$$

Furthermore: $P_D^f \in \mathbf{Z}[X]$.

PROOF. The inclusion $\mathbf{Q}(j) \subset \mathbf{Q}(f)$ yields that the extension $\mathcal{F}/\mathbf{Q}(f)$ is Galois, so we can apply Shimura reciprocity. We have to compute generators $x_1, \dots, x_k$ of $(\mathcal{O}/48\mathcal{O})^*$, map them to $\text{GL}_2(\mathbf{Z}/48\mathbf{Z})$ using the map $\bar{g}_{\sqrt{D}}$ and prove that the elements $\bar{g}_{\sqrt{D}}(x_1), \dots, \bar{g}_{\sqrt{D}}(x_k) \in \text{GL}_2(\mathbf{Z}/48\mathbf{Z})$ act trivially on f under the conditions of the theorem.

The Chinese remainder theorem gives us natural isomorphisms $(\mathcal{O}/48\mathcal{O})^* \cong (\mathcal{O}/3\mathcal{O})^* \times (\mathcal{O}/16\mathcal{O})^*$ and $\text{GL}_2(\mathbf{Z}/48\mathbf{Z}) \cong \text{GL}_2(\mathbf{Z}/3\mathbf{Z}) \times \text{GL}_2(\mathbf{Z}/16\mathbf{Z})$. We will deal with $(\mathcal{O}/3\mathcal{O})^*$ and $(\mathcal{O}/16\mathcal{O})^*$ separately. We use the Chinese remainder theorem to

lift the action of $\mathrm{GL}_2(\mathbf{Z}/3\mathbf{Z})$ to F_{48} , and embed $S_3, T_3 \in \mathrm{SL}_2(\mathbf{Z}/3\mathbf{Z})$ in $\mathrm{SL}_2(\mathbf{Z}/48\mathbf{Z})$ as

$$\begin{aligned} S_3 &\longmapsto \begin{pmatrix} 33 & 32 \\ 16 & 33 \end{pmatrix}_{48} = (T^2 S^3 T^{-16} S T^{14})_{48} \\ T_3 &\longmapsto \begin{pmatrix} 1 & 16 \\ 0 & 1 \end{pmatrix}_{48} = (T^{16})_{48}. \end{aligned}$$

This yields an action of S_3, T_3 on F_{48} . The action of $k \in (\mathbf{Z}/3\mathbf{Z})^*$ is given by $\zeta_3 \mapsto \zeta_3^k$ and $\zeta_{16} \mapsto \zeta_{16}$. We obtain the following action of S_3, T_3 on $\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2$:

$$(\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2) \circ S_3 = (\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2), \quad (\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2) \circ T_3 = (\zeta_3^2 \mathfrak{f}, \zeta_3^2 \mathfrak{f}_1, \zeta_3^2 \mathfrak{f}_2). \quad (6.2)$$

Similarly, we lift the action of $\mathrm{GL}_2(\mathbf{Z}/16\mathbf{Z})$ to F_{48} . We embed $S_{16}, T_{16} \in \mathrm{SL}_2(\mathbf{Z}/16\mathbf{Z})$ in $\mathrm{SL}_2(\mathbf{Z}/48\mathbf{Z})$ as

$$\begin{aligned} S_{16} &\longmapsto \begin{pmatrix} 16 & 15 \\ 33 & 16 \end{pmatrix}_{48} = (T^{15} S T^{15} S^{-1} T^{15})_{48} \\ T_{16} &\longmapsto \begin{pmatrix} 1 & 33 \\ 0 & 1 \end{pmatrix}_{48} = (T^{33})_{48}. \end{aligned}$$

This yields an action of S_{16}, T_{16} on F_{48} . The action of $k \in (\mathbf{Z}/16\mathbf{Z})^*$ is given by $\zeta_3 \mapsto \zeta_3$ and $\zeta_{16} \mapsto \zeta_{16}^k$. We obtain the following action of S_{16}, T_{16} on $\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2$:

$$(\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2) \circ S_{16} = (\mathfrak{f}, \mathfrak{f}_2, \mathfrak{f}_1), \quad (\mathfrak{f}, \mathfrak{f}_1, \mathfrak{f}_2) \circ T_{16} = (\zeta_{16}^5 \mathfrak{f}_1, \zeta_{16}^5 \mathfrak{f}, \zeta_{16}^6 \mathfrak{f}_2). \quad (6.3)$$

We compute generators for $(\mathcal{O}/3\mathcal{O})^*$. For $D \equiv 1 \pmod{3}$, the prime 3 splits in $\mathcal{O}$ and we take $-1, \sqrt{D}$ as generators for $(\mathcal{O}/3\mathcal{O})^* \cong \mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/2\mathbf{Z}$. It is easily checked that the corresponding matrices $\begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \in \mathrm{GL}_2(\mathbf{Z}/3\mathbf{Z})$ act trivially on $\mathfrak{f}$. For $D \equiv 2 \pmod{3}$, the prime 3 is inert in $\mathcal{O}$ and we take $1 + \sqrt{D}$ as generator for $(\mathcal{O}/3\mathcal{O})^* \cong (\mathbf{Z}/8\mathbf{Z})$. The corresponding matrix $\begin{pmatrix} 1 & 2 \\ 1 & 1 \end{pmatrix} \in \mathrm{GL}_2(\mathbf{Z}/3\mathbf{Z})$ acts trivially on $\mathfrak{f}$. Finally, we see from the transformation rules (6.2) that every matrix in $\mathrm{GL}_2(\mathbf{Z}/3\mathbf{Z})$ acts trivially on $\mathfrak{f}^3$.

Computing generators for $(\mathcal{O}/16\mathcal{O})^*$ is slightly more difficult. Let A be the maximal order of K . First we compute the structure of $(A/16A)^*$. Since we assumed $D \equiv 5 \pmod{8}$, the prime 2 is inert in A . We localize and complete the ring A at the prime ideal (2). We have $A_{(2)} = \mathbf{Z}_2[\frac{-1+\sqrt{D}}{2}]$, the unramified quadratic extension of $\mathbf{Z}_2$. The unit group $A_{(2)}^*$ is isomorphic to $\mu_3 \times (1 + 2A_{(2)})$, with μ_3 the group of 3-rd roots of unity, cf. [46, Proposition II.5.3]. Hence, we have to compute $(1 + 2A_{(2)})/(1 + 16A_{(2)})$. We have a natural isomorphism

$$\begin{aligned} (1 + 2^k A_{(2)})/(1 + 2^{2k} A_{(2)}) &\xrightarrow{\sim} A_{(2)}/2^k A_{(2)} \\ 1 + 2^k x &\longmapsto x. \end{aligned}$$

On the right hand side, we take the generators -1 and $\frac{-1+\sqrt{D}}{2}$ for $A_{(2)}/2^k A_{(2)}$ and get

$$(1 + 2A_{(2)})/(1 + 16A_{(2)}) \cong \langle -1 \rangle \times \langle 1 + 2\frac{-1+\sqrt{D}}{2} \rangle \times \langle 1 + 4\frac{-1+\sqrt{D}}{2} \rangle,$$

a group of type $(2) \times (8) \times (4)$.

We have $A_{(2)}^* = \mu_3 \times \mathcal{O}_{(2)}^*$. The inclusions

$$1 + 16A_{(2)} \stackrel{2}{\supset} 1 + 16\mathcal{O}_{(2)} \stackrel{2}{\supset} 1 + 32A_{(2)}$$

show that $(1 + 2\mathcal{O}_{(2)})/(1 + 16\mathcal{O}_{(2)})$ is an abelian group of type $(2) \times (8) \times (8)$. Since $-1 + 4\frac{-1+\sqrt{D}}{2} = -1 + 2\sqrt{D}$ has order 8 in $(\mathcal{O}/16\mathcal{O})^*$, we conclude

$$(\mathcal{O}/16\mathcal{O})^* \cong \langle -1 \rangle \times \langle \sqrt{D} \rangle \times \langle -1 + 2\sqrt{D} \rangle.$$

The matrix corresponding to $\sqrt{D} \in (\mathcal{O}/16\mathcal{O})^*$ is $\begin{pmatrix} 0 & D \\ 1 & 0 \end{pmatrix} \in \text{GL}_2(\mathbf{Z}/16\mathbf{Z})$. Using lemma 6.3 we write this matrix as

$$\begin{pmatrix} 0 & 5 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 11 \end{pmatrix} T^{11} S T^3 S T^{11} \in \text{GL}_2(\mathbf{Z}/16\mathbf{Z})$$

for $D \equiv 5 \pmod{16}$. Using the transformation formulas (6.3), it is easily checked that this matrix fixes $\mathfrak{f}$. The computation for $D \equiv 13 \pmod{16}$ proceeds similarly.

The matrix corresponding to $-1 + 2\sqrt{D} \in (\mathcal{O}/16\mathcal{O})^*$ is $\begin{pmatrix} -1 & 10 \\ 2 & -1 \end{pmatrix} \in \text{GL}_2(\mathbf{Z}/16\mathbf{Z})$, independent of $D \pmod{16}$. We write

$$\begin{pmatrix} -1 & 10 \\ 2 & -1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 13 \end{pmatrix} S T^{11} S T S T^7 \in \text{GL}_2(\mathbf{Z}/16\mathbf{Z}),$$

and it is easily checked that this matrix leaves $\mathfrak{f}$ invariant. We conclude that $\mathfrak{f}(\sqrt{D})$ is a class invariant for $3 \nmid D$, and that $(\mathfrak{f}(\sqrt{D}))^3$ is a class invariant for $3 \mid D$.

It remains to prove that $P_D^{\mathfrak{f}}$ has integer coefficients. Let $\sigma \in \text{Aut}(\mathbf{C})$ denote complex conjugation. The Weber function $\mathfrak{f}$ takes on real values along the imaginary axis, and hence we have $\sigma(\mathfrak{f}(\sqrt{D})) = \mathfrak{f}(\sqrt{D})$. Since $P_D^{\mathfrak{f}}$ is the minimal polynomial of $\mathfrak{f}(\sqrt{D})$ over K , we have an equality

$$P_D^{\mathfrak{f}} = (P_D^{\mathfrak{f}})^{\sigma}.$$

As $\mathfrak{f}(\sqrt{D})$ is also integral over $\mathcal{O}$, we have $P_D^{\mathfrak{f}} \in \mathbf{Z}[X]$. □

6.5 Shimura reciprocity

Knowing that $f(\tau)$ is a class invariant does not enable us yet to compute the minimal polynomial g of $f(\tau)$ over $K = \mathbf{Q}(\tau)$. If g has integral coefficients, we could compute $f(\tau) \in \mathbf{C}$ with very high accuracy and use the LLL-algorithm to compute g . This is a bad idea however, since the accuracy needed for this approach would be much too high.

In this section we describe a more general version of Shimura reciprocity that enables us to compute the conjugates of $f(\tau) \in H_{\mathcal{O}}$ under the Galois action of $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$. For this we need to consider the full automorphism group $\text{Aut}(\mathcal{F})$ of the modular function field $\mathcal{F}$, rather than the subgroup $\text{Gal}(\mathcal{F}/\mathbf{Q}(j)) \subset \text{Aut}(\mathcal{F})$ that we used in section 6.4.

Besides the action of $\text{GL}_2(\widehat{\mathbf{Z}})$, there is also an action of the group $\text{GL}_2(\mathbf{Q})^+$ of rational 2×2 -matrices with positive determinant on $\mathcal{F}$. Namely, for $A \in \text{GL}_2(\mathbf{Q})^+$ we define $f^A(\tau) = f(A\tau)$. We get a homomorphism $\text{GL}_2(\mathbf{Q})^+ \rightarrow \text{Aut}(\mathcal{F})$, the kernel of which is the subgroup of matrices

$$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$$

with $a \in \mathbf{Q}^*$. We identify the kernel with $\mathbf{Q}^*$.

Let $\widehat{\mathbf{Q}} = \mathbf{Q} \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$ be the finite adele ring. We obtain an action of the group $\text{GL}_2(\widehat{\mathbf{Q}})$ on $\mathcal{F}$ in the following way. Write an element $x \in \text{GL}_2(\widehat{\mathbf{Q}})$ as

$$x = u \cdot \alpha,$$

with $u \in \text{GL}_2(\widehat{\mathbf{Z}})$ and $\alpha \in \text{GL}_2(\mathbf{Q})^+$. The elements u and α are not uniquely determined by x , since we have

$$\text{GL}_2(\widehat{\mathbf{Z}}) \cap \text{GL}_2(\mathbf{Q})^+ = \text{SL}_2(\mathbf{Z}).$$

However, the action

$$f^{u \cdot \alpha} = (f^u)^\alpha$$

is well-defined, cf. [37, Theorem 7.4]. The exact sequence

$$1 \longrightarrow \mathbf{Q}^* \longrightarrow \text{GL}_2(\widehat{\mathbf{Q}}) \longrightarrow \text{Aut}(\mathcal{F}) \longrightarrow 1$$

describes the full automorphism group $\text{Aut}(\mathcal{F})$. The hard part in proving that this sequence is exact, is the surjectivity. See for instance [37, Theorem 7.6] or [56, Theorem 6.23].

The map g_τ from section 6.4 has a natural $\mathbf{Q}$ -linear extension

$$g_\tau : \widehat{K}^* = (\widehat{\mathcal{O}} \otimes_{\mathbf{Z}} \mathbf{Q})^* \rightarrow \mathrm{GL}_2(\widehat{\mathbf{Q}}),$$

which we also denote by g_τ . We get a diagram with exact rows.

$$\begin{array}{ccccccc} 1 & \longrightarrow & K^* & \longrightarrow & \widehat{K}^* & \xrightarrow{\text{Artin}} & \mathrm{Gal}(K_{\mathrm{ab}}/K) \longrightarrow 1 \\ & & & & \downarrow g_\tau & & \\ 1 & \longrightarrow & \mathbf{Q}^* & \longrightarrow & \mathrm{GL}_2(\widehat{\mathbf{Q}}) & \longrightarrow & \mathrm{Aut}(\mathcal{F}) \longrightarrow 1 \end{array} \quad (6.4)$$

Shimura reciprocity law states that also for this map g_τ we have the reciprocity relation

$$(f(\tau))^x = (f^{g(x^{-1})})(\tau)$$

for all $x \in \widehat{K}^*$.

Let $f(\tau)$ be a class invariant. We want to compute $f(\tau)^{\mathfrak{a}}$ for an ideal class $[\mathfrak{a}] \in \mathrm{Pic}(\mathcal{O})$. It suffices to pick an idele $x \in \widehat{K}^*$ that locally generates the $\widehat{\mathcal{O}}$ -ideal $\mathfrak{a} \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$. We then have

$$f(\tau)^{\mathfrak{a}} = f(\tau)^x.$$

Such an x exists since every *invertible* $\mathcal{O}$ -ideal is locally principal. The idele x is only determined up to multiplication by elements of $\widehat{\mathcal{O}}^*$. However, since $f(\tau)$ is a class invariant, we have $f(\tau)^u = f(\tau)$ for $u \in \widehat{\mathcal{O}}^*$.

Explicitly, we view the elements of $\mathrm{Pic}(\mathcal{O})$ as quadratic forms $[a, b, c]$ just like we did in section 3.3. Recall that the form $[a, b, c]$ corresponds to the invertible $\mathcal{O}$ -ideal with $\mathbf{Z}$ -basis $[\frac{-b+\sqrt{D}}{2}, a]$. We now fix τ to be the unique generator of $\mathcal{O}$ with $\mathrm{Tr}(\tau) \in \{0, 1\}$. We may then take the idele $x = (x_p)_p \in \widehat{K}^*$ with components

$$x_p = \begin{cases} a & \text{if } p \nmid a \\ \frac{-b+\sqrt{D}}{2} & \text{if } p \mid a \text{ and } p \nmid c \\ \frac{-b+\sqrt{D}}{2} - a & \text{if } p \mid a \text{ and } p \mid c, \end{cases}$$

cf. [26, Lemma 19]. The reciprocity relation yields

$$f(\tau)^{[a, -b, c]} = (f^{g_\tau(x)})(\tau). \quad (6.5)$$

The right hand side of (6.5) is independent of the choice of x .

Equality (6.5) can be used to compute the conjugates of $f(\tau)$ under the action of $\mathrm{Gal}(H_{\mathcal{O}}/K) \cong \mathrm{Pic}(\mathcal{O})$. One writes $g_\tau(x) \in \mathrm{GL}_2(\widehat{\mathbf{Q}})$ as $g_\tau(x) = u \cdot \alpha$ with $u \in \mathrm{GL}_2(\widehat{\mathbf{Z}})$ and $\alpha \in \mathrm{GL}_2(\mathbf{Q})^+$. In order to compute $f^{g_\tau(x)} = (f^u)^\alpha$, we reduce u modulo N

and compute the action of $\bar{u} \in \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$ on f . Computing the action of $\bar{u}$ on f proceeds exactly the same as in section 6.4.

We make this approach explicit. Recall the formula to compute the conjugates of $j(\tau)$

$$j(\tau)^{[a, -b, c]} = j\left(\frac{-b + \sqrt{D}}{2a}\right),$$

from section 3.3. Inspired by this formula, we want to find a function $\tilde{f} \in \mathcal{F}$ with

$$f(\tau)^{[a, -b, c]} = \tilde{f}\left(\frac{-b + \sqrt{D}}{2a}\right).$$

The difference with (6.5) is that the argument of the right hand side is now $\frac{-b + \sqrt{D}}{2a}$ instead of τ .

The element $g_\tau(x) \in \mathrm{GL}_2(\hat{\mathbf{Q}})$ is by definition the transpose of a $\hat{\mathbf{Q}}$ -linear map on $\hat{K}$ that maps the $\hat{\mathbf{Z}}$ -lattice $\hat{\mathcal{O}} = \hat{\mathbf{Z}} \cdot \tau + \hat{\mathbf{Z}} \cdot 1$ onto $\hat{\mathfrak{a}} = x\hat{\mathcal{O}}$. Define $M \in \mathrm{GL}_2(\mathbf{Q})^+ \subset \mathrm{GL}_2(\hat{\mathbf{Q}})$ to be the transpose of the $\hat{\mathbf{Q}}$ -linear map on $\hat{K}^* = \hat{\mathbf{Q}} \cdot \tau + \hat{\mathbf{Q}} \cdot 1$ that maps the basis $[\tau, 1]$ to the basis $[\frac{-b + \sqrt{D}}{2}, a]$. We see that $u_x = g_\tau(x) \cdot M^{-1} \in \mathrm{GL}_2(\hat{\mathbf{Q}})$ is actually already contained in $\mathrm{GL}_2(\hat{\mathbf{Z}})$. Indeed, u_x is the transpose of a $\hat{\mathbf{Q}}$ -linear map that stabilizes the $\hat{\mathbf{Z}}$ -lattice $\hat{\mathcal{O}}$ and is therefore an element of $\mathrm{GL}_2(\hat{\mathbf{Z}})$.

The action of M on the upper half plane $\mathbf{H}$ satisfies $M(\tau) = \frac{-b + \sqrt{D}}{2a}$. Hence, we can rewrite (6.5) as

$$f(\tau)^{[a, -b, c]} = f^{g_\tau(x) \cdot M^{-1}}\left(\frac{-b + \sqrt{D}}{2a}\right) = f^{u_x}\left(\frac{-b + \sqrt{D}}{2a}\right). \quad (6.6)$$

Since u_x is contained in $\mathrm{GL}_2(\hat{\mathbf{Z}})$, the function f^{u_x} is a conjugate of f over the field $\mathbf{Q}(j)$.

It is straightforward to give the components $u_p \in \mathrm{GL}_2(\mathbf{Z}_p)$ of the idele $u_x = (u_p)_p \in \mathrm{GL}_2(\hat{\mathbf{Z}})$. As in [26, Section 1.10], for $D \equiv 0 \pmod{4}$ we have

$$u_p = \begin{cases} \begin{pmatrix} a & \frac{b}{2} \\ 0 & 1 \end{pmatrix} & \text{if } p \nmid a \\ \begin{pmatrix} \frac{-b}{2} & -c \\ 1 & 0 \end{pmatrix} & \text{if } p \mid a \text{ and } p \nmid c \\ \begin{pmatrix} \frac{-b}{2} - a & \frac{-b}{2} - c \\ 1 & -1 \end{pmatrix} & \text{if } p \mid a \text{ and } p \mid c \end{cases}$$

and for $D \equiv 1 \pmod{4}$ we have

$$u_p = \begin{cases} \begin{pmatrix} a & \frac{b-1}{2} \\ 0 & 1 \end{pmatrix} & \text{if } p \nmid a \\ \begin{pmatrix} \frac{-b-1}{2} & -c \\ 1 & 0 \end{pmatrix} & \text{if } p \mid a \text{ and } p \nmid c \\ \begin{pmatrix} \frac{-b-1}{2} - a & \frac{1-b}{2} - c \\ 1 & -1 \end{pmatrix} & \text{if } p \mid a \text{ and } p \mid c. \end{cases}$$

Returning to the example from the previous section, we get the following conjugates of the class invariant $\mathfrak{f}(\tau)$.

THEOREM 6.5. *Let K be an imaginary quadratic number field of discriminant $D \equiv 5 \pmod{8}$ with $3 \nmid D$. Let $\zeta = \exp(2\pi i/48)$ be a primitive 48-th root of unity and $[a, b, c]$ a primitive quadratic form of discriminant $4D$. The Galois conjugate of $\mathfrak{f}(\sqrt{D})$ is given by*

$$(\mathfrak{f}(\sqrt{D}))^{[a, -b, c]} = \begin{cases} \zeta^{\frac{b}{2}(c-a-a^2c)} \cdot \mathfrak{f}\left(\frac{-b+2\sqrt{D}}{2a}\right) & \text{if } 2 \nmid a \text{ and } 2 \nmid c \\ -(-1)^{\frac{a^2-1}{8}} \cdot \zeta^{\frac{b}{2}(ac^2-a-2c)} \cdot \mathfrak{f}_1\left(\frac{-b+2\sqrt{D}}{2a}\right) & \text{if } 2 \nmid a \text{ and } 2 \mid c \\ -(-1)^{\frac{c^2-1}{8}} \cdot \zeta^{\frac{b}{2}(c-a-5ac^2)} \cdot \mathfrak{f}_2\left(\frac{-b+2\sqrt{D}}{2a}\right) & \text{if } 2 \mid a \text{ and } 2 \nmid c. \end{cases}$$

PROOF. Let $x \in \widehat{K}^*$ and $u_x = (u_p)_p \in \mathrm{GL}_2(\widehat{\mathbf{Z}})$ be as above. We have

$$(\mathfrak{f}(\sqrt{D}))^{[a, -b, c]} = \mathfrak{f}^{u_x}\left(\frac{-b+2\sqrt{D}}{2a}\right).$$

The residue class $\bar{u}_3 \in \mathrm{GL}_2(\mathbf{Z}/3\mathbf{Z})$ is

$$\bar{u}_3 = \begin{cases} \begin{pmatrix} 1 & 0 \\ 0 & a \end{pmatrix} ST^{-a} ST^{-a} ST^{-a(b+1)} & \text{if } 3 \nmid a \\ \begin{pmatrix} 1 & 0 \\ 0 & c \end{pmatrix} T^{(1+b)c} ST^c ST^c & \text{if } 3 \mid a, 3 \nmid c \\ \begin{pmatrix} 1 & 0 \\ 0 & b \end{pmatrix} T^{1+b} ST^b ST^{b-1} & \text{if } 3 \mid a, 3 \mid c \end{cases}$$

and the residue class $\bar{u}_2 \in \mathrm{GL}_2(\mathbf{Z}/16\mathbf{Z})$ is

$$\bar{u}_2 = \begin{cases} \begin{pmatrix} 1 & 0 \\ 0 & a \end{pmatrix} ST^{\frac{-1}{a}} ST^{-a} ST^{\frac{1}{a}(\frac{b}{2}-1)} & \text{if } 2 \nmid a \\ \begin{pmatrix} 1 & 0 \\ 0 & c \end{pmatrix} T^{(1-\frac{b}{2})c} ST^{\frac{1}{c}} ST^c & \text{if } 2 \mid a, 2 \nmid c \\ \text{the case } 2 \mid a, 2 \mid c \text{ cannot occur.} & \end{cases}$$

The last line deserves some explanation. If both a and c would be even, the determinant of $\bar{u}_2 \in \mathrm{GL}_2(\mathbf{Z}/16\mathbf{Z})$ would be $a + b + c$. The condition $b^2 - 4ac \equiv 4 \pmod{8}$ implies however that b is always even. Hence, we would have that $a + b + c$ is even, a contradiction.

We have

$$\mathfrak{f}^{[a, -b, c]} = (\mathfrak{f}^{u_3})^{u_2}$$

and the right hand side is easily computed using the transformation rules (6.2) and (6.3). A simple – although a bit laborious – check then shows that the resulting formulas are the same as the ones given in the theorem. $\square$

6.6 Class invariants in a non-archimedean setting

Shimura reciprocity enables us to work with class invariants in a systematic way. For a class invariant $f(\tau) \in H_{\mathcal{O}}$, we know which modular function $\tilde{f}$ we should take to compute the conjugate

$$f(\tau)^{[a, -b, c]} = \tilde{f}\left(\frac{-b + \sqrt{D}}{2a}\right)$$

of $f(\tau)$ under the action of the quadratic form $[a, -b, c]$. We can for instance explicitly compute $f(\tau) \in \mathbf{C}$ by using the Fourier expansion of f . If it is known that the minimal polynomial of $f(\tau)$ has integer coefficients, we can approximate $f(\tau)^{[a, -b, c]} \in \mathbf{C}$ with high enough accuracy and expand the product

$$P_{\Delta}^f = \prod_{[a, b, c] \in \mathcal{F}_{\Delta}^+} (X - f(\tau)^{[a, b, c]}) \in \mathbf{Z}[X],$$

just like we did for the j -function.

We also want to use class invariants in a non-archimedean setting. The computation of $f(\tau)$ via the Fourier expansion has no p -adic analogue. In chapter 5 we explained a p -adic algorithm to work with the j -function, where p is a prime that splits completely in $H_{\mathcal{O}}$. We viewed j as an element of the function field of the modular curve $X(1)$, and gave an algorithm to compute the finite set $\text{Ell}_{\Delta}(\mathbf{Q}_p)$ of j -invariants of elliptic curves over $\mathbf{Q}_p$ with endomorphism ring $\mathcal{O}_{\Delta}$. The Hilbert class polynomial P_{Δ} can then be computed as

$$P_{\Delta} = \prod_{j(E) \in \text{Ell}_{\Delta}(\mathbf{Q}_p)} (X - j(E)) \in \mathbf{Z}[X].$$

We extend the algorithm from chapter 5 to cope with class invariants. The extension we present in this section is not ideally suited for explicit computations yet, and serves as a stepping stone for the more practical version of the next sections.

The modular functions we will consider are *integral* over $\mathbf{Z}[j]$, so they are given as the zero of some irreducible polynomial $\Psi_f(X, j) \in \mathbf{Z}[j, X]$. Let f be such a modular function, say of level $N \geq 1$. In sections 6.4–6.5 we only needed the Fourier expansion of f , but here we also need to know the polynomial $\Psi_f(X, j)$.

For a j -value $j(\tilde{E}) \in \text{Ell}_{\Delta}(\mathbf{Q}_p)$, the roots of the polynomial $\Psi_f(X, j(\tilde{E})) \in H_{\mathcal{O}}[X]$ lie in the ray class field of conductor $H_{N, \mathcal{O}}$ of conductor N for the order $\mathcal{O}$, cf. theorem 6.2. If we know that f yields class invariants (for instance by using Shimura reciprocity), we know that some of these roots actually lie in the ring class

field $H_{\mathcal{O}}$. We need to decide which ones, and compute the action of the Galois group $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$ on such roots.

The first step will be to compute an element $j(\tilde{E}) \in \text{Ell}_{\Delta}(\mathbf{Q}_p)$ with high enough accuracy. The accuracy needed depends on the function f in the following way. Define the Mahler measure $M(g)$ of a polynomial $g = \prod_{i=1}^n (X - \alpha_i) \in \mathbf{C}[X]$ as $M(g) = \sum_{i=1}^n \log \max(1, |\alpha_i|)$. Then the following holds:

$$\frac{M(P_D)}{M(P_D^f)} = \frac{\deg_j(\Psi(j, X))}{\deg_X(\Psi(j, X))} (1 + o(1)), \quad (6.7)$$

for $|D|$ tending to infinity, cf. [21, Proposition 3]. Define $r(f) = \frac{\deg_j(\Psi(j, X))}{\deg_X(\Psi(j, X))}$, the ‘reduction factor’ of f , so for γ_2 we have $r(\gamma_2) = 1/3$ and for f we have $r(f) = 1/72$. In practice, relation (6.7) means that if d is the accuracy needed for the computation of P_D , and hence for $j(\tilde{E}) \in \text{Ell}_D(\mathbf{Q}_p)$, we may multiply d by $r(f)$ if we are working with f instead of with j .

Next we need to decide which roots of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$ lie in the ring class field $H_{\mathcal{O}}$. Of course we may be lucky, meaning that $\Psi_f(X, j(\tilde{E}))$ has only one root in $\mathbf{Q}_p$. This occurs for instance for the function γ_2 and $p \equiv 2 \pmod{3}$. Indeed, if $\Psi_{\gamma_2}(X, j(\tilde{E})) = X^3 - j(\tilde{E})$ would have two (and hence three) roots in $\mathbf{Q}_p$, then we would have $\zeta_3 \in \mathbf{Q}_p$, contradicting the assumption $p \equiv 2 \pmod{3}$. Rather than looking if there is a prime p such that $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$ has a single root that lies in the ring class field $H_{\mathcal{O}}$, we discuss how we can decide in general which roots lie in $H_{\mathcal{O}}$. We assume that p does not divide the level N . This assumption is harmless, as N is small and p is large.

The key observation is that f is an element of the function field

$$F_N = \mathbf{Q}(j, f_{r,s} \mid r, s \in \frac{1}{N}\mathbf{Z}/\mathbf{Z}, \text{ not both } 0)$$

of the modular curve $X(N)$ over $\mathbf{Q}(\zeta_N)$, cf. theorem 6.1. The Fricke functions $f_{r,s}$ are normalized x -coordinates of N -torsion of points on the elliptic curve $\mathbf{C}/(\mathbf{Z} + \mathbf{Z} \cdot \tau)$. We can write f as a $\mathbf{Q}$ -rational function in j and the functions $f_{r,s}$.

Fix a primitive N -th root of unity $\zeta_N \in \overline{\mathbf{Q}}_p$. For $a \in (\mathbf{Z}/N\mathbf{Z})^*$, let $Y(N)_a$ be the modular curve from section 5.3. Then f is an element of the function field of $Y(N)_a$, $\overline{\mathbf{Q}}_p$ for every $a \in (\mathbf{Z}/N\mathbf{Z})^*$. Let $x \in \mathbf{Q}_p$ be a root of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$. There exist $a \in (\mathbf{Z}/N\mathbf{Z})^*$ and $(\tilde{E}, P, Q) \in Y(N)_a(\overline{\mathbf{Q}}_p)$ with $f(\tilde{E}, P, Q) = x \in \mathbf{Q}_p$.

From chapter 5 we know that there is an action of the group of invertible $\mathcal{O}$ -ideals on the set $\text{Ell}_{\Delta}(\mathbf{Q}_p)$. An invertible ideal I sends $j(\tilde{E})$ to $j(\tilde{E}^I)$ and we have $j(\tilde{E})^{[I, H_{\mathcal{O}}/K]} = j(\tilde{E}^I)$. If N is coprime to the norm l of I , the isogeny

$$\varphi_I : \tilde{E} \longrightarrow \tilde{E}^I$$

extends to a natural isomorphism

$$\varphi_I : \tilde{E}[N] \longrightarrow \tilde{E}^I[N].$$

A basis $\langle P, Q \rangle$ for $\tilde{E}[N]$ gets mapped to a basis $\langle P^I, Q^I \rangle$ for $\tilde{E}^I[N]$. We compute

$$e_N(P^I, Q^I) = e_N(P, \hat{\varphi}_I(Q^I)) = e_N(P, lQ) = \zeta_N^l$$

and conclude that we have $(\tilde{E}^I, P^I, Q^I) \in Y(N)_{la}(\bar{\mathbf{Q}}_p)$. We have the fundamental equality

$$f(\tilde{E}, P, Q)^{[I, H_{N, \mathcal{O}}/K]} = f(\tilde{E}^I, P^I, Q^I).$$

We can explicitly compute the isogeny φ_I : first we compute the kernel polynomial $g_I \in \mathbf{Q}_p[X]$ corresponding to I as in chapter 5 and then we compute the isogeny using Vélú's formulas [64]. Hence, we have a way of computing $f(\tilde{E}, P, Q)^{[I, H_{N, \mathcal{O}}/H_{\mathcal{O}}]}$.

A root $x \in \mathbf{Q}_p$ of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$ lies in $H_{\mathcal{O}}$ if and only if it is invariant under

$$\text{Gal}(H_{N, \mathcal{O}}/H_{\mathcal{O}}) \cong (\mathcal{O}/N\mathcal{O})^*/\mathcal{O}^*.$$

We write $x = f(\tilde{E}, P, Q)$ for some choice of basis $P, Q \in \tilde{E}[N]$ and test whether $x^{[(y), H_{N, \mathcal{O}}/H_{\mathcal{O}}]} = x$ holds for all generators y of $(\mathcal{O}/N\mathcal{O})^*/\mathcal{O}^*$. See below for an example how to write $x = f(\tilde{E}, P, Q)$ and section 6.7 for remarks on smoothness bounds for a set of generators of $(\mathcal{O}/N\mathcal{O})^*/\mathcal{O}^*$.

Once we have found that a certain root $x \in \mathbf{Q}_p$ lies in the ring class field $H_{\mathcal{O}}$, we need to compute its conjugates under $\text{Gal}(H_{\mathcal{O}}/K) \cong \text{Pic}(\mathcal{O})$. This proceeds exactly as before, since we have

$$x^{[I, H_{\mathcal{O}}/K]} = f(\tilde{E}, P, Q)^{[I, H_{\mathcal{O}}/K]} = f(\tilde{E}^I, P^I, Q^I) \in \mathbf{Q}_p.$$

All we require is that the norm $N(I)$ of I is coprime to the level N of f . If the minimal polynomial of x has integer coefficients, we may expand the product

$$\prod_{[I] \in \text{Pic}(\mathcal{O})} (X - f(\tilde{E}, P, Q)^{[I, H_{\mathcal{O}}/K]}) \in \mathbf{Z}_p[X]$$

and round the coefficients to integers, just as we did in chapter 5.

EXAMPLE 1. Let $\gamma_2 : \mathbf{H} \rightarrow \mathbf{C}$ be the holomorphic cube root of j with integral Fourier expansion. It is a classical fact that γ_2 is modular of level 3. Let $\mathcal{O} = \mathbf{Z}[\tau]$ have discriminant D with $3 \nmid D$ and suppose that we have $\tau + \bar{\tau} = 0 \pmod{3}$. We have seen in section 6.4 that $\gamma_2(\tau)$ is then a class invariant.

Let $E : Y^2 = X^3 + aX + b$ be an elliptic curve over $\mathbf{Q}_p$, with $p > 3$. Let $c_1, \dots, c_4 \in \overline{\mathbf{Q}}_p$ be the roots of the 3-division polynomial of degree $(3^2 - 1)/2 = 4$. Then

$$\frac{-48a}{2a - 3(c_1c_2 + c_3c_4)} \quad (6.8)$$

is a cube root of $j(E)$, as may be checked by using for instance the Fourier expansion of the Fricke functions. Expression 6.8 nicely illustrates that γ_2 is not a function of an elliptic curve alone: also some ordering on the 3-torsion is required. We indeed get three distinct cube roots of $j(E)$. From a geometric point of view, there is no way to single out a root ‘corresponding’ to γ_2 .

Next we illustrate how we can use this ‘geometric γ_2 ’ to compute the polynomial $P_{-31}^{\gamma_2} \in \mathbf{Z}[X]$ for the order $\mathcal{O}$ of discriminant -31 using p -adic methods. The primes $47 = 4^2 + 31$ and $67 = 6^2 + 31$ both split completely in the Hilbert class field H of $K = \mathbf{Q}(\sqrt{-31})$. For primes p with $p \equiv 1 \pmod{3}$, a j -invariant $j(\tilde{E}) \in \text{Ell}_{-31}(\mathbf{Q}_p)$ has 3 roots in $\mathbf{Q}_p$. Since this is the most difficult case, we take $p = 67$.

First we compute a curve $\tilde{E}/\mathbf{Q}_p$ with $\text{End}(\tilde{E}) \cong \mathcal{O}$. The accuracy needed is only one third of the required accuracy for the computation of the Hilbert class polynomial P_{-31} . Using the algorithm from chapter 5 we find that we may take

$$j(\tilde{E}) = 3 + 33p - 16p^2 + O(p^3) \in \mathbf{Q}_p$$

as j -invariant. The three cube roots of $j(\tilde{E})$ are

$$\eta_1 = 18 + O(p)$$

$$\eta_2 = 53 + O(p)$$

$$\eta_3 = 63 + O(p).$$

Only one of them lies in the Hilbert class field H . Indeed, if 2 roots would lie in H , then ζ_3 would be contained in H as well. This means that $\mathbf{Q}(\zeta_3)$ would be a subfield of H and hence 3 would divide -31 , which it does not.

We fix a Weierstraß equation

$$Y^2 = X^3 + aX + b$$

for $\tilde{E}/\mathbf{Q}_p$. Let $c_1, \dots, c_4 \in \overline{\mathbf{Q}}_p$ be the 4 roots of the 3-division polynomial for $\tilde{E}$. We compute 3-torsion points $P_i/\overline{\mathbf{Q}}_p$ with x -coordinate c_i . The points P_i are defined over the unramified extension of degree 4 over $\mathbf{Q}_p$.

Let I be an $\mathcal{O}$ -ideal with $3 \nmid N(I)$. The isogeny $\varphi_I : \tilde{E} \rightarrow \tilde{E}^I$ extends to a natural isomorphism

$$\varphi_I : \tilde{E}[3] \xrightarrow{\sim} \tilde{E}^I[3].$$

Hence, we get a natural bijection

$$\varphi_I : \{\eta_1, \eta_2, \eta_3\} \xrightarrow{\sim} \{\text{cube roots of } j(\tilde{E}^I)\}.$$

For a cube root

$$\eta = \frac{-48a}{2a - 3(c_1c_2 + c_3c_4)}$$

we have

$$\eta^{[I, H_3/H]} = \frac{-48a'}{2a' - 3(c'_1c'_2 + c'_3c'_4)}.$$

Here, c'_i is the x -coordinate of $\varphi_I(P_i) \in \tilde{E}^I[3]$ and $\tilde{E}^I$ has Weierstraß equation $Y^2 = X^3 + a'X + b'$.

The group $(\mathcal{O}/3\mathcal{O})^*/\mathcal{O}^* \cong \mathbf{Z}/4\mathbf{Z}$ is generated by $\alpha = \frac{-1+\sqrt{-31}}{2}$ of norm 8. We compute $\eta_i^{[I, H_3/H]}$ for $I = (\alpha) = \mathfrak{p}_2^3$ and get

$$\begin{aligned} \eta_1 &\xrightarrow{\varphi_I} \eta_1 \\ \eta_2 &\xrightarrow{\varphi_I} \eta_3 \\ \eta_3 &\xrightarrow{\varphi_I} \eta_2. \end{aligned}$$

Hence, $\eta_1 = 18 + O(p)$ is a class invariant. Note that $\varphi_{\mathfrak{p}_2}$ is just a 2-isogeny, so we do not actually need the ‘Atkin-Elkies’ techniques from chapter 5.

Computing the conjugates of $\eta_1 \in H$ under $\text{Gal}(H/K) \cong \text{Pic}(\mathcal{O})$ proceeds similarly. We have $\text{Pic}(\mathcal{O}) \cong \mathbf{Z}/3\mathbf{Z} \cong \langle [\mathfrak{p}_2] \rangle$ and

$$\eta_1^{[\mathfrak{p}_2, H/K]} = \varphi_{\mathfrak{p}_2}(\eta_1).$$

We compute the conjugates of η_1 under $\text{Pic}(\mathcal{O})$ and expand

$$P_{-31}^{\gamma_2} = \prod_{i=1}^3 (X - \varphi_{\mathfrak{p}_2}^i(\eta_1)) = X^3 + 342X^2 + 837X + 116127 \in \mathbf{Z}[X].$$

EXAMPLE 2. The second example concerns Δ -quotients. Let Δ be the classical modular form of weight 12 with Fourier expansion $\Delta(\tau) = q \prod_{n=1}^{\infty} (1 - q^n)^{24}$. Define

$$\begin{aligned} g_1(\tau) &= \frac{\Delta((\tau+1)/2)}{\Delta(\tau)} \\ g_2(\tau) &= \frac{\Delta(\tau/2)}{\Delta(\tau)} \\ g_3(\tau) &= 2^{12} \frac{\Delta(2\tau)}{\Delta(\tau)}. \end{aligned}$$

The functions g_i are modular of level 2 and have rational Fourier coefficients. They are the three roots of

$$(X + 16)^3 - jX \in \mathbf{Z}[j, X].$$

Let $\mathcal{O} = \mathbf{Z}[\tau]$ have discriminant D with $D \equiv 1 \pmod{8}$. Since 2 splits in $\mathcal{O}$, the ray class field H_2 of $K = \mathbf{Q}(\sqrt{D})$ equals the Hilbert class field H . Hence, $g_i(\tau)$ is a class invariant for $i = 1, 2, 3$.

The geometric interpretation is the following. Let $E : Y^2 = X^3 + aX + b$ be an elliptic curve over $\mathbf{Q}_p$, with $p > 3$, and let P_1, P_2, P_3 be the three 2-torsion points. If E has endomorphism ring $\mathcal{O}$, then the 2-torsion is defined over $\mathbf{Q}_p$, cf. section 5.2. Let $E_i = E^{(P_i)} : Y^2 = X^3 + a_iX + b_i$ be the 2-isogenous curve obtained by applying the normalized isogeny with kernel $\langle P_i \rangle$. Then

$$\frac{\Delta(X^3 + a_iX + b_i)}{\Delta(X^3 + aX + b)} \tag{6.9}$$

is a root of $(X + 16)^3 - j(E_i)X$. Here, $\Delta(f)$ means the discriminant of the polynomial f . In particular, (6.9) is independent of the choice of a Weierstraß equation for E . We write $\Delta(f(E))$ if we have *fixed* a Weierstraß equation $Y^2 = f(X)$ for E .

We again focus on the order $\mathcal{O} = \mathbf{Z}[\tau]$ of discriminant -31 . We have a choice of functions to work with now, since now all $g_i(\tau)$ are class invariants. However, only for the function g_2 the value at $\tau = \frac{-1+\sqrt{-31}}{2}$ is real, i.e., only for g_2 the minimal polynomial has integer coefficients. Also in the p -adic setting, we want to ensure that the minimal polynomial we compute has integer coefficients.

First we compute a curve $\tilde{E}/\mathbf{Q}_p$ with endomorphism ring $\text{End}(\tilde{E}) \cong \mathcal{O}$. The required precision is the same as in the previous example, so we again take

$$j(\tilde{E}) = 3 + 33p - 16p^2 + O(p^3) \in \mathbf{Q}_p$$

as j -invariant. We fix

$$Y^2 = X^3 + (11 - 2p + 20p^2)X - 22 + 4p + 27p^2$$

as Weierstraß equation for $\tilde{E}$. The three 2-torsion points of $\tilde{E}$ are

$$\begin{aligned} P_1 &= (32 + O(p), 0) \\ P_2 &= (46 + O(p), 0) \\ P_3 &= (56 + O(p), 0). \end{aligned}$$

We compute the 2-torsion of $\tilde{E}^{(P_i)}$ for $i = 1, 2, 3$. It turns out that for $\tilde{E}^{(P_1)}$ and for $\tilde{E}^{(P_3)}$ the complete 2-torsion is defined over $\mathbf{Q}_p$, and for $\tilde{E}^{(P_2)}$ it is not. Consequently, the curves $\tilde{E}^{(P_1)}$ and $\tilde{E}^{(P_3)}$ have endomorphism ring $\mathcal{O}$. The curve $\tilde{E}^{(P_2)}$ has endomorphism ring $\mathcal{O}_{-4.31} \subsetneq \mathcal{O}$. This means that

$$\frac{\Delta(f(\tilde{E}^{(P_1)}))}{\Delta(f(\tilde{E}))} \in H \quad \text{and} \quad \frac{\Delta(f(\tilde{E}^{(P_3)}))}{\Delta(f(\tilde{E}))} \in H$$

are complex conjugates. If we want integer coefficients, we have to compute the minimal polynomial of $\Delta(f(\tilde{E}^{(P_2)}))/\Delta(f(\tilde{E}))$.

We compute $\Delta(\tilde{E}^{(P_2)})/\Delta(\tilde{E}) = 5 + 2p + 32p^2 + O(p^3) \in \mathbf{Q}_p$. The class group $\text{Pic}(\mathcal{O}) \cong \mathbf{Z}/3\mathbf{Z}$ is generated by a prime $\mathfrak{p}_5$ of norm 5. We cannot use the generator $\mathfrak{p}_2$ any more, since 2 divides the level $N = 2$. We compute $\varphi_{\mathfrak{p}_5}(P_2) \in \tilde{E}^{\mathfrak{p}_5}[2]$ and use this 2-torsion point to compute

$$\frac{\Delta(f(\tilde{E}^{(P_2)}))}{\Delta(f(\tilde{E}))} \stackrel{[\mathfrak{p}_5, H/K]}{=} -22 - 32p - 28p^2 + O(p^3) \in \mathbf{Q}_p.$$

Finally, we compute the third conjugate and expand

$$P_{-31}^{g_2} = \prod_{i=1}^3 \left(X - \varphi_{\mathfrak{p}_5}^i \left(\frac{\Delta(f(\tilde{E}^{(P_2)}))}{\Delta(f(\tilde{E}))} \right) \right) = X^3 + 165X^2 + 9642X + 1 \in \mathbf{Z}[X].$$

6.7 Finding a class invariant

As in the previous section, let f be a modular function of level $N \geq 1$ that is integral over $\mathbf{Z}[j]$. In order to check which roots of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Q}_p[X]$ are in fact class invariants, we want generators of $(\mathcal{O}/N\mathcal{O})^*/\mathcal{O}^*$ that are *smooth*. Lemma 5.18 tells us that, if GRH is true, we can find a B -smooth element $a + b\pi_p \in \mathcal{O}$ with $B = \lfloor \exp(\sqrt{\log|\Delta|}) \rfloor$. We would like to extend this result, i.e., we would like that there also exists a B_N -smooth element $a + b\pi_p \in \mathcal{O}$ that lies in a prescribed residue class in $(\mathcal{O}/N\mathcal{O})^*$. Here, we put $B_N = g_N B$ with g_N some function depending only on N .

Although being B -smooth and lying in a prescribed residue class in $(\mathcal{O}/N\mathcal{O})^*$ are quite unrelated, there seems to be no hope in proving a result as in the previous paragraph. We do not know enough on the distribution of smooth elements in $\mathcal{O}$. The situation is similar to that in [39], where a conjecture from analytic number theory is needed to prove the run time for the elliptic curve factoring method. The problem lies in a sense in our understanding of analytic number theory, and not so much in our understanding of elliptic curves.

In practice there is no problem. One just sieves in the set

$$S = \{a + b\pi_p : a, b \in \mathbf{Z}, b \neq 0, (a, b) = 1, a + b\pi_p \text{ and } p\Delta \text{ are coprime}\}$$

for enough smooth elements that generate $(\mathcal{O}/N\mathcal{O})^*$. We expect that the smooth elements are equidistributed over $(\mathcal{O}/N\mathcal{O})^*$.

Let $E/\mathbf{F}_p$ be a curve with endomorphism ring $\mathcal{O}$. We can use Shimura reciprocity to compute the number of roots in $\mathbf{F}_p$ of $\Psi_f(X, j(E)) \in \mathbf{F}_p[X]$ in advance. Assume that $\Psi_f(X, j(E)) \in \mathbf{F}_p[X]$ is separable and that p does not divide the level N . These assumptions are usually fulfilled in practice, as p is not too small, of size $|\Delta|$.

Write $\mathcal{O} = \mathbf{Z}[\tau]$ for some $\tau \in \mathbf{H}$. Consider the root $x = f(\tau)$ of $\Psi_f(X, j(\tau)) \in \mathbf{C}[X]$. We know that x is an element of the ray class field $H_{\mathcal{O}, N}$ of conductor N for the order $\mathcal{O}$, cf. theorem 6.2. There is a prime $\mathfrak{p}$ of $H_{\mathcal{O}, N}$ lying over p such that the reduction $j(\tau) \bmod \mathfrak{p}$ equals $j(E) \in \mathbf{F}_p$. The reduction $\bar{x} \in \bar{\mathbf{F}}_p$ of x modulo $\mathfrak{p}$ is a root of $\Psi_f(X, j(E)) \in \mathbf{F}_p[X]$. Write $x = f(\tilde{E}, \tilde{P}, \tilde{Q})$ with $(\tilde{E}, \tilde{P}, \tilde{Q}) \in Y(N)_a(H_{\mathcal{O}, N})$. Let $F_p : E \rightarrow E$ be the Frobenius of $E/\mathbf{F}_p$ and let $\varphi \in \text{End}(\tilde{E})$ have reduction $F_p \in \text{End}(E)$. The endomorphism $\varphi : \tilde{E} \rightarrow \tilde{E}$ induces an isomorphism $\tilde{E}[N] \xrightarrow{\sim} \tilde{E}[N]$. Writing $\pi_p \in \mathcal{O}$ for the image of φ under the normalized isomorphism $\text{End}(\tilde{E}) \xrightarrow{\sim} \mathcal{O}$, we have

$$\bar{x} \in \mathbf{F}_p \iff f(\tilde{E}, \tilde{P}, \tilde{Q})^{[(\pi_p), H_{\mathcal{O}, N}/H]} = f(\tilde{E}, \tilde{P}, \tilde{Q}).$$

Shimura reciprocity tells us:

$$\bar{x} \in \mathbf{F}_p \iff f^{g(\pi_p)} = f,$$

where $g = g_\tau$ is the connecting homomorphism from section 6.4.

Denote the zeroes of $\Psi_f(j, X) \in \mathbf{Z}[j, X]$ by f_i . The functions f_i are elements of $\mathbf{Q}(\zeta_N)((q^{1/N}))$ and our fixed f is one of the f_i 's. We have proved the following equality:

$$\#\{x \in \mathbf{F}_p \mid \Psi_f(x, j(E)) = 0\} = \#\{f_i : f_i^{g(\pi_p)} = f_i\}. \quad (6.10)$$

EXAMPLE. We illustrate relation (6.10) by computing the number of roots in $\mathbf{F}_p$ of $\Psi_{\mathfrak{f}}(X, j(E)) \in \mathbf{F}_p[X]$, with $\mathfrak{f}$ the classical Weber function. The minimal polynomial of $\mathfrak{f}$ over $\mathbf{Z}[j]$ is

$$\Psi_{\mathfrak{f}} = (X^{24} - 16)^3 - jX^{24} \in \mathbf{Z}[j, X].$$

The roots of $\Psi_{\mathfrak{f}}$ are

$$\zeta_{48}^k \mathfrak{f}, \quad \zeta_{48}^l \mathfrak{f}_1, \quad \zeta_{48}^l \mathfrak{f}_2,$$

where k is an even integer and l is odd. If we take k odd and l even, we get the roots of $(X^{24} + 16)^3 - jX^{24} \in \mathbf{Z}[j, X]$.

Take an order $\mathcal{O}$ of discriminant Δ , with $\Delta \equiv 1 \pmod{8}$ and with $3 \nmid \Delta$. The equation $t^2 - 4p = \Delta$ has no solutions with p an odd prime, hence we look for a solution to $t^2 - 4p = 4\Delta$. In practice we always find a solution to this equation and we take any solution (t, p) . We see that for an elliptic curve $E/\mathbf{F}_p$ with $\text{End}(E) \cong \mathcal{O}$, the order $\mathbf{Z}[F_p]$ is contained in $\mathcal{O}$ with index 2.

We have seen that $\zeta_{48} \mathfrak{f}_2 \left(\frac{-1+\sqrt{\Delta}}{2} \right)$ is a class invariant, hence some of the roots of $\Psi_{\mathfrak{f}}(X, j(E)) \in \mathbf{F}_p[X]$ are reductions of class invariants. Using equality 6.10, we can compute the number of roots of $\Psi_{\mathfrak{f}}(X, j(E)) \in \mathbf{F}_p[X]$. Write $\mathcal{O} = \mathbf{Z}[\tau]$ with $\tau = \frac{-1+\sqrt{\Delta}}{2}$, and $\pi_p = 2\tau + 1 + t/2$.

The matrix

$$A = g_{\tau}(\pi_p) = \begin{pmatrix} \frac{t-2}{2} & \frac{\Delta-1}{2} \\ 2 & \frac{t+2}{2} \end{pmatrix} \in \text{GL}_2(\widehat{\mathbf{Z}})$$

corresponding to π_p has norm p and trace t . Modulo 3, we have

$$A = \begin{pmatrix} 1 & 0 \\ 0 & p \end{pmatrix} T^{tp} S T^{2p} S T^{t-1-p} \in \text{GL}_2(\mathbf{Z}/3\mathbf{Z}),$$

and we compute $(\zeta_3^k \mathfrak{f})^A, (\zeta_3^k \mathfrak{f}_1)^A, (\zeta_3^k \mathfrak{f}_2)^A$. Using the transformation rules (6.2) for the Weber functions, we see that for $p \equiv 1 \pmod{3}$, all 9 functions $\zeta_3^k \mathfrak{f}_i$ are invariant under the action of A . For $p \equiv 2 \pmod{3}$, only the 3 functions $\zeta_3 \mathfrak{f}_i$ are invariant under A .

Modulo 16 we get

$$A = \begin{pmatrix} 1 & 0 \\ 0 & p \end{pmatrix} S T^{(\frac{2}{p}+1)\frac{-2}{t-2}} S T^{\frac{2-t}{2}} S T^{(\frac{2}{p}+1)\frac{\Delta-1}{t-2} - \frac{t+2}{2p}} \in \text{GL}_2(\mathbf{Z}/16\mathbf{Z}).$$

Computing $(\zeta_{16}^k \mathfrak{f}_i)^A$ is slightly more cumbersome. It is perhaps easiest to just distinguish cases for p, t and Δ and to write a small computer program to compute the action of A on $\zeta_{16}^k \mathfrak{f}_i$. For p we have the possibilities 3, 7, 11 and 15 mod 16. Since we

need to know $t/2$, we have to know t modulo 32. We have the cases $t = 4, 8, \dots, 32$. Finally, Δ can be congruent to 1 or 9 modulo 16.

For instance, for $p \equiv 3 \pmod{16}$, $t \equiv 8 \pmod{32}$ and $\Delta \equiv 9 \pmod{16}$ we get

$$(\zeta_{16}^k \mathfrak{f})^A = \zeta_{16}^{3k+14} \mathfrak{f}, \quad (\zeta_{16}^k \mathfrak{f}_1)^A = \zeta_{16}^{3k+12} \mathfrak{f}_1, \quad (\zeta_{16}^k \mathfrak{f}_2)^A = \zeta_{16}^{3k+6} \mathfrak{f}_2.$$

We find that A leaves $\zeta_{16} \mathfrak{f}$ invariant for $k = 1, 9$. It leaves $\zeta_{16}^k \mathfrak{f}_1$ invariant for $k = 2, 10$ and it leaves $\zeta_{16}^k \mathfrak{f}_2$ invariant for $k = 5, 13$. The other cases proceed similarly and in all cases we find 2 *odd* k for $\mathfrak{f}$, 2 *even* k for $\mathfrak{f}_1$ and 2 *odd* k for $\mathfrak{f}_2$.

We summarize this computation in the following theorem.

THEOREM 6.6. *Let $p \geq 5$ and $E/\mathbf{F}_p$ be as above. Then the following holds:*

$$\begin{aligned} p \equiv 1 \pmod{3} &\implies (X^{24} - 16)^3 - j(E)X^{24} \in \mathbf{F}_p[X] \quad \text{has exactly 6 roots} \\ &\quad (X^{24} + 16)^3 - j(E)X^{24} \in \mathbf{F}_p[X] \quad \text{has exactly 12 roots;} \\ p \equiv 2 \pmod{3} &\implies (X^{24} - 16)^3 - j(E)X^{24} \in \mathbf{F}_p[X] \quad \text{has exactly 2 roots} \\ &\quad (X^{24} + 16)^3 - j(E)X^{24} \in \mathbf{F}_p[X] \quad \text{has exactly 4 roots.} \end{aligned}$$

REMARK. For simplicity we only considered the case where $\mathbf{Z}[F_p]$ has index 2 in $\mathcal{O}$. If we look at how the index enters the formulas, we see that we need to know the index modulo 32. Hence, it is a finite computation to resolve the other cases.

REMARK. This theorem shows that for $p \equiv 2 \pmod{3}$, both roots of $\Psi_{\mathfrak{f}}(X, j(E)) \in \mathbf{F}_p[X]$ are reductions of class invariants. For $p \equiv 1 \pmod{3}$ we get 6 roots, 2 of which are reductions of class invariants. In this case, it suffices to check which roots are invariant under the action of $(\mathcal{O}/3\mathcal{O})^*/\mathcal{O}^*$.

EXAMPLE. Let $\mathcal{O}$ be the order of discriminant -31 . Take $p = 47$. We take $16 \in \mathbf{F}_p$ as j -invariant of a curve with endomorphism ring $\mathcal{O}$. The roots of $(X^{24} - 16)^3 - 16X^{24}$ are $22, 25 = -22 \in \mathbf{F}_p$. Both roots are reductions of class invariants. The roots of $(X^{24} + 16)^3 - 16X^{24}$ are $\pm 10, \pm 13 \in \mathbf{F}_p$.

For $p = 67$ we take $3 \in \mathbf{F}_p$ as j -invariant. The roots of $(X^{24} - 16)^3 - 3X^{24}$ are $\pm 4, \pm 14, \pm 18 \in \mathbf{F}_p$, and the roots of $(X^{24} + 16)^3 - 3X^{24}$ are $\pm 6, \pm 12, \pm 13, \pm 21, \pm 25$ and $\pm 27 \in \mathbf{F}_p$.

6.8 Using modular polynomials

The theory developed in section 6.6 is not ideally suited for explicit computations yet. If we are given a modular function f of level $N \geq 1$ in terms of its Fourier expansion, we must find the moduli interpretation of this function. Concretely this means that we have to write f as a rational function in j and the x -coordinates of N -torsion points of an elliptic curve with j -invariant j . We have done this for γ_2 and Δ -quotients in section 6.6, but there is no ‘systematic way’ of approaching this problem.

Even worse: if we do find the moduli interpretation of f , it might be that we have to (partially) factor the N -division polynomial of an elliptic curve. The degree of the N -th division polynomial is $O(N^2)$ for $N \rightarrow \infty$, so it might take a lot of time to factor such a polynomial. This would destroy the advantage of working with ‘smaller’ functions than j . As an example: we have not given the geometric interpretation of the Weber functions. This is partially due to the fact that we just do not know it, but if we were to work with a geometric description, we might have to factor a 48-division polynomial of degree 1153. This annihilates the speed improvements gained by working with f instead of with j .

The second problem is that working with explicit isogenies is quite slow in practice. In the examples in section 6.6, we only needed a 2- and a 5-isogeny and these are quite easy to write down. For larger discriminants than $|\Delta| = 31$ in the example, we have to work with isogenies of larger degree. For $\Delta \approx -10^{10}$ for instance, one typically needs isogenies of degree ≈ 100 . Although the ‘Atkin-Elkies’ techniques combined with Vélú also work fine in this case, this is not something one wants to do in practice.

The answer to these problems lies in observation that it suffices to compute

$$x^I,$$

where $x \in \mathbf{Q}_p$ is a root of $\Psi_f(X, j(E)) \in \mathbf{Q}_p[X]$ and I is an invertible $\mathcal{O}$ -ideal of norm coprime to N . Indeed, if we want to know which root x of $\Psi_f(X, j(E)) \in \mathbf{Q}_p[X]$ is a class invariant, we need to check which root is invariant under $(\mathcal{O}/N\mathcal{O})^*/\mathcal{O}^*$. This amounts to computing x^I , for I the principal ideal generated by one of the generators of $(\mathcal{O}/N\mathcal{O})^*$. Once we know that $x \in \mathbf{Q}_p$ is a class invariant, we need to compute $x^I \in \mathbf{Q}_p$, with I one of the generators of $\text{Pic}(\mathcal{O})$. In this section we give a method to compute x^I that also works fast in practice. We have to make some mild assumptions on f . It does not require the moduli interpretation of the function f .

Write $\Gamma(f)$ for the stabilizer of f inside $\mathrm{SL}_2(\mathbf{Z})$. We have

$$\Gamma(N) \subseteq \Gamma(f) \subseteq \mathrm{SL}_2(\mathbf{Z}),$$

by the assumption that f is modular of level N . Write $X(f)$ for the modular curve corresponding to the congruence subgroup $\Gamma(f)$, cf. section 5.3. The complex points of this curve are $\Gamma(f) \backslash \overline{\mathbf{H}}$. The curve $X(f)$ is a quotient of the modular curve $X(N)$ by a subgroup of $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})$, and hence can be defined over $\mathbf{Q}(\zeta_N)$. We have a commutative diagram

$$\begin{array}{ccc} X(N) & \xrightarrow{f} & \mathbf{P}^1 \\ & \searrow & \nearrow f \\ & X(f) & \end{array}$$

and $f : X(N) \rightarrow \mathbf{P}_{\mathbf{C}}^1$ factors through the quotient $X(f)$.

Likewise, there exists a curve $X(f)_a$ such that we have a commutative diagram

$$\begin{array}{ccc} X(N)_a & \xrightarrow{f} & \mathbf{P}^1 \\ & \searrow & \nearrow f \\ & X(f)_a & \end{array}$$

As a complex curve, we have $X(f)_a = X(f)$.

We have fixed a primitive N -th root of unity $\zeta_N \in \overline{\mathbf{Q}}_p$. For ease of notation, we simply denote a point on $X(f)_a$ by a triple (E, P, Q) instead of $(\overline{E}, \overline{P}, \overline{Q})$. Here, P, Q form a basis for the N -torsion $E[N]$ of E with $e_N(P, Q) = \zeta_N^a$.

Let l be a prime not dividing the level N . Write $\Gamma(f; l) = \Gamma(f) \cap \Gamma_0(l)$. We have

$$\Gamma(lN) \subseteq \Gamma(f; l) \subseteq \Gamma(f).$$

Let $X(f; l)$ be the modular curve corresponding to $\Gamma(f; l)$. It can be defined over $\mathbf{Q}(\zeta_{lN})$. Just as we have curves $X(f)_a$, we also have curves $X(f; l)_a$. Points on $X(f; l)_a$ are quadruples (E, P, Q, G) , with $(E, P, Q) \in X(f)_a$ and $G \subseteq E[l]$ a subgroup of order l .

There is a natural map $s : X(f; l)_a \rightarrow X(f)_a$ and a natural map $t : X(f; l)_a \rightarrow X(f)_{la}$. The map s sends $(E, P, Q, G) \in X(f; l)_a$ to $(E, P, Q) \in X(f)_a$. The map t sends $(E, P, Q, G) \in X(f; l)_a$ to $(E/G, \varphi(P), \varphi(Q))$, where $\varphi : G \rightarrow E/G$ has kernel G . The situation is as follows.

$$\begin{array}{ccccc}
 & & X(f; l)_a & & \\
 & \swarrow s & & \searrow t & \\
 X(f)_a & & & & X(f)_{la} \\
 \downarrow f & \nearrow F & & \nwarrow F' & \downarrow f \\
 \mathbf{P}^1 & & & & \mathbf{P}^1
 \end{array}$$

Here, F and F' are the composed maps.

LEMMA 6.7. *The maps s, t in the diagram above both have degree $l + 1$.*

PROOF. We will show that the diagram

$$\begin{array}{ccc}
 X(f)_a & \xleftarrow{s} & X(f; l)_a \\
 \downarrow f & \square & \downarrow \\
 X(1) & \xleftarrow{\quad} & X_0(l)
 \end{array}$$

is cartesian in the category of smooth projective curves with surjective maps. As the cover $X_0(l)/X(1)$ has degree $l + 1$, this implies that s and t have degree $l + 1$. Here, the maps on the ‘lower right part’ of the square are the forgetful maps. Instead of working over $\mathbf{Q}(\zeta_N)$, we will work over $\mathbf{C}$; the same result then holds over $\mathbf{Q}(\zeta_N)$. We may then omit the subscript a in the diagram. Moreover, it is easier to work with $X(N)/X(1)$ and $X(l)/X(1)$ instead of $X(f)/X(1)$ and $X_0(l)/X(1)$, since in this case we explicitly know the Galois groups. The Galois group of $X(N)/X(1)$ is $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})/\{\pm 1\}$ and it is $\mathrm{SL}_2(\mathbf{Z}/l\mathbf{Z})/\{\pm 1\}$ for $X(l)/X(1)$.

The fibred product of $X(l)$ and $X(N)$ is almost equal to $X(Nl)$. Indeed, writing $d(k) = \#(\mathrm{SL}_2(\mathbf{Z}/k\mathbf{Z})/\{\pm 1\})$, the degree of $X(Nl)/X(1)$ is

$$\#\mathrm{SL}_2(\mathbf{Z}/Nl\mathbf{Z})/\{\pm 1\} = 2 \cdot d(N)d(l),$$

and we obtain the following diagram

$$\begin{array}{ccccc}
 & & & & X(Nl) \\
 & & & \nearrow 2 & \\
 X(N) & \xleftarrow{d(l)} & X(Nl)/H & & \\
 \downarrow d(N) & \square & \downarrow d(N) & & \\
 X(1) & \xleftarrow{d(l)} & X(l) & &
 \end{array}$$

where H is the subgroup $\{1\} \times \{\pm 1\} \subseteq \mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z}) \times \mathrm{SL}_2(\mathbf{Z}/l\mathbf{Z})$. Since we are working over $\mathbf{C}$, we know that the degrees on parallel sides of the square are equal. Hence, the square is cartesian.

Since we have $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \in \Gamma_0(l)$, the curve $X(Nl)/H$ is a cover of $X(f; l)$. Hence, the diagram

$$\begin{array}{ccc} X(f) & \xleftarrow{s} & X(f; l) \\ \downarrow f & \square & \downarrow \\ X(1) & \xleftarrow{t} & X_0(l) \end{array}$$

‘fits inside’ the bigger cartesian diagram for $X(Nl)/H$. In particular, it is cartesian. Since the degree of $X_0(l)/X(1)$ is $l + 1$, the same must hold for $X(f; l)/X(f)$, i.e., both s and t have degree $l + 1$. $\square$

Reformulating this result in terms of function fields, we see that the tensor product of the function fields of $X(f)_a$ and $X_0(l)$ is again a field, and equals the function field of $X(f; l)_a$.

REMARK. The curve $X_0(l)$ can be defined over $\mathbf{Q}$. The cartesian diagram shows that $X(f; l)_a$ can be defined over $\mathbf{Q}(\zeta_N)$.

We map $Y(f; l)_a$ to a curve C inside $\mathbf{A}^1 \times \mathbf{A}^1$ as in the diagram below. The map b is defined by $b(x) = (s(x), t(x))$. In practice, b will often have degree 1 and then $Y(f; l)_a$ is birational to C . The maps p_1, p_2 are the two projection maps.

$$\begin{array}{ccccc} & & F & & \\ & \nearrow & & \searrow & \\ Y(f; l)_a & \xrightarrow{b} & C & \xrightarrow{\quad} & \mathbf{A}^1 \times \mathbf{A}^1 & \xrightarrow{p_1} & \mathbf{A}^1 \\ & \searrow & F' & & \downarrow p_2 & & \\ & & & & \mathbf{A}^1 & & \end{array}$$

We are interested in the curve C , since the function field of C is generated by f and f_l . Here, f_l may be defined by $f_l(\tau) = f^{\sigma_l}(l\tau)$ in complex analytic terms, where the notation f^{σ_l} is the same as in section 6.2. If f has rational Fourier coefficients, we have $f_l(\tau) = f(l\tau)$. There exists a polynomial $\Phi \in \mathbf{Q}(\zeta_N)[X, Y]$ such that C is the curve defined by $\Phi = 0$. We will see that knowing Φ basically solves the problem of computing x^I .

As we have $\deg(F) = \deg(F')$, the previous diagram immediately tells us that we have

$$\deg_X(\Phi) = \deg_Y(\Phi) = \frac{(l+1)\deg(f)}{\deg(b)}.$$

Let Φ_{f,f_l} be the minimal polynomial of f_l over $\mathbf{Q}(\zeta_N)(f)$. The coefficients of Φ_{f,f_l} need not be polynomials in f yet, but after multiplying the coefficients by the common denominator, we get a polynomial in $\mathbf{Q}(\zeta_N)[X, Y]$. Since the function field of C is $\mathbf{Q}(\zeta_N)(f, f_l)$, this polynomial is a model for the curve C and hence we may take $\Phi = \Phi_{f,f_l}$.

We call Φ_{f,f_l} the *modular polynomial* of level l for f . For $f = j$ this definition coincides with the ‘classical’ modular equation from chapter 5. The term modular polynomial or modular equation appears already in Weber [65], see for instance §73, *Die Schlaefischen Modulargleichungen*. Weber works with specific functions, and hence there is no general definition of a modular polynomial. Our definition coincides with Weber’s definition for the functions he considers.

LEMMA 6.8. *Suppose that f generates the function field of $X(f)_C$ over $\mathbf{C}(j)$ and suppose $f \in \mathbf{Q}((q^{1/N}))$. Then we have $\Phi \in \mathbf{Q}[X, Y]$.*


PROOF. It suffices to show that $X(f)$ can be defined over $\mathbf{Q}$. Since the algebraic closure of $\mathbf{Q}$ inside $\mathbf{Q}(f, j)$ is $\mathbf{Q}$ itself, the minimal polynomial Ψ_f of f over $\mathbf{Q}(j)$ is absolutely irreducible. The curve defined by $\Psi_f = 0$ is absolutely irreducible and has $\mathbf{Q}(f, j)$ as function field. Since f generates the function field of $X(f)/\mathbf{C}$, the curve $X(f)$ can be defined over $\mathbf{Q}$. $\square$

REMARK. Computing Φ is relatively easy if we know the Fourier expansion of f . We have an upper bound

$$\deg(f)(l+1)$$

for the degrees $\deg_X(\Phi)$ and $\deg_Y(\Phi)$. In most cases this upper bound will in fact be an equality. By comparing the Fourier coefficients of f and f_l , we can recursively find the coefficients of Φ . See section 7.1 for examples.

Let $x \in \mathbf{Q}_p$ be a root of $\Psi_f(X, j(E)) \in \mathbf{Q}_p[X]$ and let I be an invertible $\mathcal{O}$ -ideal of norm $l \nmid Np$. Let $\Phi = \Phi_{f,f_l}$ be the polynomial defined above. From the moduli interpretation of $X(f; l)_a$, it is clear that one of the roots of $\Phi(x, X) \in \overline{\mathbf{Q}}_p[X]$ equals x^I . To see what the other roots are, we look at the diagram



again. Above $x \in \mathbf{A}^1(\bar{\mathbf{Q}}_p)$ there are $\deg(f)$ distinct points $(E_i, P_i, Q_i) \in Y(f)_a(\bar{\mathbf{Q}}_p)$. Above $(E_i, P_i, Q_i) \in Y(f)_a(\bar{\mathbf{Q}}_p)$, there are $l+1$ points $(E_i, P_i, Q_i, G_j) \in Y(f; l)_a(\bar{\mathbf{Q}}_p)$. Here, G_j ranges over the $l+1$ subgroups of order l of $E_i[l]$. The points (E_i, P_i, Q_i, G_j) all map to $x \in \mathbf{A}^1(\bar{\mathbf{Q}}_p)$ under F . The images under $F' : X(f; l)_a \rightarrow \mathbf{A}^1$ are exactly the roots of $\Phi(x, X)$.

REMARK. The curve $X(f; l)$ is a quotient of $X(N)$. Since $X(N)$ has good reduction outside N , the curve $X(f; l)$ has good reduction outside lN by [33, Proposition 4.2]. Hence, the description of the roots of $\Phi(x, X)$ remains valid over $\mathbf{F}_p$.

We want to decide which root of $\Phi(x, X)$ is actually x^I . The first observation is that it suffices to look at the roots in $\mathbf{Q}_p$. Indeed, if x is a class invariant then we automatically have $x \in \mathbf{Q}_p$. If x is not a class invariant, then x^I need not lie in $\mathbf{Q}_p$. But if it does not, we have automatically proven that x is not a class invariant.

Usually, x^I is the only root of $\Phi(x, X)$ that is also a root of $\Psi_f(X, j(E^I))$. Hence, we test for all roots $\alpha \in \mathbf{Q}_p$ of $\Phi(x, X)$ whether $\Psi_f(\alpha, j(E^I)) = 0$ holds. If x is a class invariant, we find at least one such α . If we find exactly one root with this property, we have computed x^I .

The problem is that in the general context we are working in, we cannot prove much on the number of roots of $\Phi(x, X)$ with specific properties. A moduli interpretation of f would seem the least to require here. In practice there is never a problem however. We choose our function f to be ‘small’, and we expect that the degree $\deg(f)$ will be small in general. We show that for $\deg(f) = 1$ and a class invariant x , there is exactly one root $\alpha \in \mathbf{Q}_p$ of both $\Phi(x, X) \in \mathbf{Q}_p[X]$ and $\Psi_f(X, j(E^I)) \in \mathbf{Q}_p[X]$.

Let $(E, P, Q) \in Y(f)_a(\bar{\mathbf{Q}}_p)$ be the unique point of $Y(f)_a$ with $f(E, P, Q) = x \in \mathbf{Q}_p$. Of the $l+1$ points $(E, P, Q, G_i) \in Y(f; l)_a(\bar{\mathbf{Q}}_p)$ lying over $(E, P, Q) \in Y(f)_a(\bar{\mathbf{Q}}_p)$, only for the 2 points $G_i = E[I]$ and $G_i = E[\bar{I}]$, the value $j(t(E, P, Q, G_i))$ is contained in $\mathbf{Q}_p$, cf. lemma 5.11. If both $F'(E, P, Q, E[I])$ and $F'(E, P, Q, E[\bar{I}])$ are roots of $\Phi(x, X) \in \mathbf{Q}_p[X]$, then we must have $[I] = [\bar{I}] \in \text{Pic}(\mathcal{O})$. Since x is a class invariant, we then have $F'(E, P, Q, E[I]) = F'(E, P, Q, E[\bar{I}])$.

6.9 Further improvements

It is of great help that for many class invariants the coefficients of the modular polynomial $\Phi_l^f = \Phi_{f, f_l}$ are a lot smaller than those of the classical modular polynomial for j . As an example, we consider the classical Weber function f . For small primes l the coefficients of the polynomial are really small, like

$$\begin{aligned}\Phi_5^f(X, Y) &= (X^5 - Y)(X - Y^5) + 5XY \\ \Phi_7^f(X, Y) &= (X^7 - Y)(X - Y^7) + 7(XY - X^4Y^4).\end{aligned}$$

For $l = 13$ it takes at least two of these pages to write down the classical polynomial Φ_l , but we have

$$\begin{aligned}\Phi_{13}^f(X, Y) &= (X^{13} - Y)(X - Y^{13}) + 5 \cdot 13XY \\ &\quad + 13(X^2Y^{12} + X^{12}Y^2 + 4X^{10}Y^4 + 4X^{10}Y^4 + 6X^6Y^8 + 6X^8Y^6).\end{aligned}$$

This can be used to give a significant practical speed up of the algorithm. We now give the final algorithm to compute on input $\Delta < -4$ a generating polynomial for the ring class field $H_{\mathcal{O}}$ corresponding to the order of discriminant Δ . Assume that we are given a modular function f that yields class invariants when evaluated at appropriate generators $\tau \in \mathbf{H}$ of the $\mathbf{Z}$ -algebra $\mathcal{O} = \mathbf{Z}[\tau]$. We need the Fourier expansion of f and the minimal polynomial $\Psi_f(f, X) \in \mathbf{Z}[j, X]$. Assume furthermore that the minimal polynomial P_{Δ}^f of a class invariant $f(\tau)$ has integer coefficients.

Step 1. Find a prime $p \nmid N$ and an elliptic curve $E/\mathbf{F}_p$ with $\text{End}(E) \cong \mathcal{O}$. This is done using the endomorphism ring algorithm from section 5.2. We compute the zeroes $x_1, \dots, x_k \in \mathbf{F}_p$ of $\Psi_f(X, j(E)) \in \mathbf{F}_p[X]$.

Step 2. We have to decide which of these zeroes is the reduction of a class invariant. Compute smooth primitive generators $y_1, \dots, y_t$ of $(\mathcal{O}/N\mathcal{O})^*$. We will show how to compute $x_1^{y_1}$.

Write $y_1 = \alpha_1 \cdot \dots \cdot \alpha_s$, with $N(\alpha_i) = l_i \in \mathbf{Z}$ prime. Compute the cycle

$$j(E) \xrightarrow{\bar{\rho}_{\alpha_1}} j(E^{(\alpha_1)}) \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_t}} j(E^{(y_1)}) = j(E)$$

of j -invariants over $\mathbf{F}_p$. This is done exactly as in chapter 5, employing the modular polynomials for j . Compute all roots $\eta_i \in \mathbf{F}_p$ of $\Phi_{l_1}^f(x_1, X) \in \mathbf{F}_p[X]$ that also satisfy $\Psi_f(\eta_i, j(E^{(\alpha_1)})) = 0$. Here, $\Phi_{l_1}^f$ is the modular polynomial for f . In practice, we find either zero or one such root η_i , cf. section 6.8. If we find zero roots, we know

that x_1 is not the reduction of a class invariant. We then repeat this computation with x_2 , etc.

The only root of both $\Phi_{l_1}^f(x_1, X)$ and $\Psi_f(X, j(E^{(\alpha_1)}))$ has to be $x_1^{(\alpha_1)}$. Continuing like this, we compute a series

$$x_1 \xrightarrow{\bar{\rho}_{\alpha_1}} x_1^{(\alpha_1)} \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_t}} x_1^{(y_1)}.$$

If we have $x_1^{y_1} = x_1$, we compute $x_1^{y_2}$, etc. If x_1 is invariant under all generators $y_1, \dots, y_t$ of $(\mathcal{O}/N\mathcal{O})^*$, it is the reduction of a class invariant. Otherwise, we repeat this computation with x_2 , etc., until we find a reduction of a class invariant.

Step 3. Say that $x \in \mathbf{F}_p$ is the reduction of a class invariant. We choose a smooth $\mathcal{O}$ -ideal $(\alpha) = \alpha_1 \dots \alpha_u$ for the map ρ_α from chapter 5. Here, we require that the norm of α is coprime to the level N .

We compute a cycle

$$j(E) \xrightarrow{\bar{\rho}_{\alpha_1}} j(E^{(\alpha_1)}) \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_u}} j(E^\alpha) = j(E)$$

and using this cycle the corresponding cycle

$$x \xrightarrow{\bar{\rho}_{\alpha_1}} x^{(\alpha_1)} \xrightarrow{\bar{\rho}_{\alpha_2}} \dots \xrightarrow{\bar{\rho}_{\alpha_u}} x^\alpha = x$$

for x , just like we did in step 2.

Step 4. Lift $E/\mathbf{F}_p$ to $E_1/\mathbf{Q}_p$ by lifting the coefficients of the Weierstraß equation for E arbitrarily. We use two p -adic digits accuracy in this step.

Step 5. Lift $x_1 \in \mathbf{F}_p$ to $x_1 \in \mathbf{Z}_p/(p^2)$ as a root of $\Psi(X, j(E_1)) \in \mathbf{Z}_p[X]$. It is easy to compute $x_1^{\alpha_1}$. Indeed, we know the reduction $\overline{x_1^{\alpha_1}} = x^{\alpha_1}$ modulo p , hence we know which root of $\Phi_l^f(x_1, X) \in (\mathbf{Z}_p/(p^2))[X]$ to pick. In this way we compute $x_1^\alpha \in \mathbf{Z}_p/(p^2)$.

Next we compute $\rho_\alpha(j(E_1))$ as the unique root of $\Psi(x^{(\alpha)}, X) \in \mathbf{Z}_p[X]$ that reduces to $j(E)$ modulo p .

Step 6. Update $\rho_\alpha(j(E_1))$ according to the same formula as in chapter 5:

$$j(E_{k+1}) = j(E_k) - \frac{\rho_\alpha(j(E_k)) - j(E_k)}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1}.$$

Step 7. Repeat step 5 with $j(E_1)$ replaced by $j(E_2)$. We now work with four p -adic digits precision. We obtain $j(E_3)$. Continue this iteration process until we have computed the canonical lift $j(\tilde{E})$ with high enough accuracy.

Step 8. Compute the ‘canonical lift’ $\tilde{x} \in \mathbf{Z}_p$ of $x \in \mathbf{F}_p$ as a root of $\Psi_f(X, j(\tilde{E}))$. Again, since we know $\tilde{x} = x \in \mathbf{F}_p$, we know which root to pick.

Step 9. It remains to compute the conjugates of $\tilde{x}$ under $\text{Pic}(\mathcal{O})$. This is done exactly as before. For an invertible $\mathcal{O}$ -ideal I of norm l coprime to N , we compute $j(E^I) \in \mathbf{F}_p$ using the methods of chapter 5. Knowing $j(E^I)$, we compute a root $\beta \in \mathbf{F}_p$ of $\Phi_l^f(x, X) \in \mathbf{F}_p[X]$ that also satisfies $\Psi_f(\beta, j(E^I)) = 0$. Just as in step 2, in practice we only find *one* such β , and we then have $x^I = \beta$.

We know the reduction $\tilde{x}^I = \beta$ of $\tilde{x}^I$, and consequently, we know which root of $\Phi_l^f(\tilde{x}, X) \in \mathbf{Z}_p[X]$ is $\tilde{x}^I$.

Step 10. Expand the polynomial

$$P_D^f = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - \tilde{x}^I) \in \mathbf{Z}[X],$$

just like we did in chapter 5.

REMARK. We only need the (large) modular polynomials for j when we are working over $\mathbf{F}_p$. In the lifting process we only need to know the smaller modular polynomials for f . In practice this is a significant speed up.

REMARK. The algorithm makes clear that we constantly have to work with two functions: f and j . Just knowing an f -value $x \in \mathbf{A}^1(\mathbf{Q}_p)$ is not enough, we also need to know a j -value. This should come as no surprise, since the function field of the curve $X(f)$ is generated by *two* functions: f and j .

7.1 A cryptographic curve

As an example of the techniques described in this thesis, we construct a curve that can readily be used for elliptic curve cryptography. The discrete log problem in $E(\mathbf{F}_q)$ is considered to be hard if the order $N = \#E(\mathbf{F}_q)$ is a prime number of roughly 200 bits. We will construct a curve of prime order

$$N = 123456789012345678901234567890123456789012345678901234568197 \approx 10^{59}.$$

As in chapter 4, we look for a solution $x, y \in \mathbf{Z}_{\geq 1}$ to $x^2 - Dy^2 = 4N$ for varying D with $q = N + 1 - x$ prime. The ‘smallest’ D that admits a solution is $D = -2419$. We have

$$x^2 + 2419y^2 = 4N$$

and

$$q = N + 1 + x = 123456789012345678901234567890654833374525085966737125236501$$

is prime for

$$x = 531376585512740287835890668303$$

$$y = 9349802208089011828618119329.$$

The fact that the first 30 digits of N and q are the same is no coincidence: Hasse’s theorem tells us that the group order N differs at most $2\sqrt{q}$ from the size q of the finite field.

The class polynomial P_D has degree 8 and $P_D \in \mathbf{F}_q[X]$ splits completely. Any of its zeroes is a j -invariant of a curve with N points. In the next subsection we explain the non-archimedean algorithm to compute $P_D \in \mathbf{Z}[X]$. With

$$a = 78876029697996107120563826094864556580999965110862558799913 \in \mathbf{F}_q,$$

the curve E defined by $Y^2 = X^3 + aX - a$ has CM by $\mathcal{O}_D$. The point $P = (1, 1) \in E(\mathbf{F}_q)$ does not satisfy $N \cdot P = O_E$, so the quadratic twist E' of E defined by

$$Y^2 = X^3 + 4aX - 8a$$

has exactly N points over $\mathbf{F}_q$. This can easily be checked by computing NP' for a random point $P' \in E'(\mathbf{F}_q)$.

► Computing the class polynomial

Our goal is to compute the class polynomial $P_D \in \mathbf{Z}[X]$ for $D = -2419$. We will do this p -adically, so we first find an elliptic curve over some finite field $\mathbf{F}_p$ that has CM by $\mathcal{O}_D$. As in section 5.1, we look for the smallest integer $t > 0$ for which $(t^2 - D)/4$ is prime. This yields $t = 3$ and $p = 607$. We fix p for the rest of section 7.1. Since p splits completely in the Hilbert class field of $\mathbf{Q}(\sqrt{D})$, there exists an elliptic curve over $\mathbf{F}_p$ with endomorphism ring $\mathcal{O}_D$, i.e., a curve with trace of Frobenius $t = \pm 3$. We apply the naïve algorithm from chapter 2 and look for a curve with $p + 1 \pm t$ points over $\mathbf{F}_p$. We find that the curve $E/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + X + 56$$

of j -invariant $j(E) = 137 \in \mathbf{F}_p$ has trace of Frobenius 3 and consequently endomorphism ring $\mathcal{O}_D$.

We need to compute the canonical lift $j(\tilde{E}) \in \mathbf{Q}_p$ of $j(E) \in \mathbf{F}_p$ up to k p -adic digits precision with

$$k = \frac{\pi\sqrt{|D|}}{\log p} \sum_{[a,b,c] \in \mathcal{F}_D^+} \frac{1}{a}. \quad (7.1)$$

The class group of the order $\mathcal{O}_D$ is cyclic of order 8 and representing the elements as binary quadratic forms as in section 3.3, we find $k \approx 43$. We will compute $j(\tilde{E}) \in \mathbf{Q}_p$ up to 45 p -adic digits precision.

Next we determine which element $\alpha \in \mathcal{O}_D \setminus \mathbf{Z}$ we will use for the map $\rho_\alpha : X_D(\mathbf{C}_p) \rightarrow X_D(\mathbf{C}_p)$. The ideal $(\alpha) = (3 + \pi_p)$ of norm $625 = 5^4$ factors as

$$(\alpha) = \mathfrak{p}_5^4 = (5, \pi_p + 3)^4.$$

Here, $\pi_p = \frac{3+\sqrt{D}}{2}$ is a prime element of norm p . We compute the action of the prime ideal $\mathfrak{p}_5$ on $j(E)$. The eigenvalue for the action of Frobenius on the 5-torsion $E[5]$ is

$-3 \in \mathbf{F}_5$. If we evaluate the modular polynomial $\Phi_5(X, Y)$ in $X = j(E) = 137 \in \mathbf{F}_p$, we get a polynomial which has 2 roots over $\mathbf{F}_p$, namely 214 and 309. From this we deduce that $\mathfrak{p}_5$ sends $j(E)$ to one of these 2 roots; we do not know which one yet.

We just *guess* that the correct j -invariant is $214 \in \mathbf{F}_p$. Using the techniques from section 5.6, we compute the eigenspace C of the 5-torsion corresponding to this isogeny. We get the x -coordinates of the points on E in C as zeroes of

$$\bar{f}_C = X^2 + 502X + 90 \in \mathbf{F}_p[X].$$

Since we know that the eigenvalue for the action of $\mathfrak{p}_5$ is $-3 \in \mathbf{F}_p$, we can now just check whether

$$(X^p, Y^p) = -3 \cdot (X, Y)$$

holds for points in C , i.e., we compute both (X^p, Y^p) and $-3 \cdot (X, Y)$ in the ring

$$\mathbf{F}_p[X, Y]/(\bar{f}_C, Y^2 - X^3 - X - 56).$$

Here, the $\cdot$ means adding *on the curve*. In this example, it turns out that (X^p, Y^p) and $-3 \cdot (X, Y)$ are not the same. It follows that the j -invariant of the $\mathfrak{p}_5$ -isogenous curve is the other value $309 \in \mathbf{F}_p$.

The action of $\mathfrak{p}_5$ on the j -invariant 309 is now easier to compute: the modular polynomial $\Phi_5(X, 309) \in \mathbf{F}_p[X]$ again has two roots, but one of these roots is $j(E)$. This root corresponds to the action of $\bar{\mathfrak{p}}_5$, so we pick the other root. If we compute the entire cycle corresponding to $\mathfrak{p}_5$, we get:

$$137 \xrightarrow{\mathfrak{p}_5} 309 \xrightarrow{\mathfrak{p}_5} 532 \xrightarrow{\mathfrak{p}_5} 214 \xrightarrow{\mathfrak{p}_5} 137.$$

It should come as no surprise that we get a cycle of length 4, since we know that $\mathfrak{p}_5^4 = (\alpha)$ acts trivially.

We now lift $E/\mathbf{F}_p$ to $E_1/\mathbf{Q}_p$ by lifting the coefficients of the Weierstraß equation arbitrarily. The polynomial $\Phi_5(X, j(E_1)) \in \mathbf{Z}_p[X]$ has exactly 2 roots, one of which reduces to 309 modulo p . Taking the lift $E_1/\mathbf{Q}_p$ defined by $Y^2 = X^3 + X + 56$, we find the roots $214 + 91p$ and $309 - 92p$. We have $j(E_1)^{\mathfrak{p}_5} = 309 - 92p$. Continuing like this, we get the ‘cycle’

$$137 - 41p \xrightarrow{\mathfrak{p}_5} 309 - 92p \xrightarrow{\mathfrak{p}_5} 532 - 133p \xrightarrow{\mathfrak{p}_5} 214 - 251p \xrightarrow{\mathfrak{p}_5} 137 - 28p$$

over $\mathbf{Q}_p$. We update $j(E_1)$ according to the ‘Newton formula’

$$j(E_{k+1}) = j(E_k) - \frac{\rho_\alpha(j(E_k)) - j(E_k)}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1} \quad (7.2)$$

and find that $j(E_2) = 137 - 15p$ is the two digit approximation of the j -invariant of the canonical lift.

Starting from $j(E_2)$, we now lift the cycle to four p -adic digits precision, compute $j(E_3)$ from this, and so on. We obtain

$$\begin{aligned}
 j(\tilde{E}) &= 137 + O(p) \\
 &= 137 - 15p + O(p^2) \\
 &= 137 - 15p - 290p^2 - 8p^3 + O(p^4) \\
 &= 137 - 15p - 290p^2 - 8p^3 - 107p^4 - 108p^5 - 192p^6 - 35p^7 + O(p^8) \\
 &= 137 - 15p - 290p^2 - 8p^3 - 107p^4 - 108p^5 - 192p^6 - 35p^7 - 172p^8 + 14p^9 \\
 &\quad - 160p^{10} + 98p^{11} - 195p^{12} + 303p^{13} - 212p^{14} - 283p^{15} + O(p^{16}).
 \end{aligned}$$

We continue this process until we have computed the canonical lift in 45 p -adic digits accuracy.

To compute the conjugates of $j(\tilde{E})$ under $\text{Gal}(H/K) \cong \text{Pic}(\mathcal{O}_D)$, we note that the class group $\text{Pic}(\mathcal{O}_D)$ is cyclic of order 8 and is generated by a prime $\mathfrak{p}_{11}$ of norm 11. To speed up the computation, we also use that the prime $\mathfrak{p}_5$ of norm 5 has order 4 in the class group. We compute the conjugates $\rho_{\mathfrak{p}_5}(j(\tilde{E}))$, $\rho_{\mathfrak{p}_5^2}(j(\tilde{E}))$ and $\rho_{\mathfrak{p}_5^3}(j(\tilde{E}))$ using the modular polynomial $\Phi_5(X, Y) \in \mathbf{Z}_p[X, Y]$. Although we know which root of $\Phi_5(X, j(\tilde{E})) \in \mathbf{Z}_p[X]$ is the j -invariant $\rho_{\mathfrak{p}_5}(j(\tilde{E}))$ – we know its reduction modulo p – this is not important at the moment, since we need to know all conjugates. The remaining four conjugates are computed using the modular polynomial $\Phi_{11}(X, Y) \in \mathbf{Z}_p[X, Y]$. Once we have computed all conjugates, we expand the degree 8 polynomial

$$P_{-2419} = \prod_{[I] \in \text{Pic}(\mathcal{O}_D)} (X - j(\tilde{E})^I) \in \mathbf{Z}[X].$$

The polynomial P_D has coefficients up to 119 decimal digits.

► Using a cube root of j

A generating polynomial for the Hilbert class field H of $K = \mathbf{Q}(\sqrt{D})$ with smaller coefficients than P_D can be achieved by working with a cube root of the j -function instead of with j itself. The j -function has a holomorphic cube root $\gamma_2 : \mathbf{H} \rightarrow \mathbf{C}$ with integral Fourier expansion, and it is known that γ_2 yields class invariants for $3 \nmid D$. Since we have $3 \nmid D = -2419$, there is cube root of $j(\tilde{E})$ that also lies in the Hilbert class field H . In this subsection we explain how to compute the minimal polynomial of this class invariant.

As we have $p \equiv 1 \pmod{3}$, there are three cube roots $\beta_1, \beta_2, \beta_3$ of $j(E) = 137 \in \mathbf{F}_p$. We need to decide which one is the reduction of a class invariant. As γ_2 is modular of level 3, we need to decide which root is invariant under the action of the Galois group $\text{Gal}(H_3/H) \cong (\mathcal{O}/3\mathcal{O})^*/\{\pm 1\}$ of the ray class field H_3 over H . Since 3 is inert in $\mathcal{O}$, we have $(\mathcal{O}/3\mathcal{O})^*/\{\pm 1\} \cong \mathbf{Z}/4\mathbf{Z}$.

The group $(\mathcal{O}/3\mathcal{O})^*/\{\pm 1\}$ is generated by $\alpha = \pi_p - 1$ of norm $605 = 5 \cdot 11^2$. We compute the action of the principal ideal (α) on $\beta_1, \beta_2, \beta_3$. Write $(\alpha) = \bar{\mathfrak{p}}_5 \cdot \mathfrak{p}_{11}^2$, with $\bar{\mathfrak{p}}_5 = (5, \pi_p - 1)$ and $\mathfrak{p}_{11} = (11, \pi_p - 1)$. In chapter 6 we explained how to compute $\beta_1^{\bar{\mathfrak{p}}_5}$ using an explicit description in terms of 3-torsion points. A faster way is to use the modular polynomials from sections 6.8 and 6.9, and this is the method that we will use in this section.

We need to know the modular polynomials of level 5 and 11 for γ_2 . We noted already in chapter 6 that, for $l \neq 3$, the modular polynomial $\Phi_l^{\gamma_2}(X, Y)$ of level l for γ_2 has integer coefficients and degree $l + 1$ in both X and Y . By comparing the Fourier expansions of $\gamma_2(q)$ and $\gamma_2(q^l)$, we can recursively find the coefficients of $\Phi_l^{\gamma_2}$. The following general lemma simplifies our computations.

LEMMA 7.1. *Let f be a modular function, and let l be a prime not dividing the level of f . Suppose that the modular polynomial Φ_l^f has integer coefficients. If f is invariant under the action of either $S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \in \text{SL}_2(\mathbf{Z})$ or $M = \begin{pmatrix} 0 & -l \\ 1 & 0 \end{pmatrix} \in \text{GL}_2(\mathbf{Q})$, then Φ_l^f is symmetric, i.e., we have*

$$\Phi_l^f(X, Y) = \Phi_l^f(Y, X).$$

PROOF. The proof is similar to the symmetry proof [37, Theorem 5.3] of the classical modular polynomial for the j -function. Assume first that f is invariant under S . If we replace z by $-1/(lz)$ in the equation $\Phi_l^f(f(z), f(lz)) = 0$, we obtain

$$\Phi_l^f(f(-1/(lz)), f(-1/z)) = 0.$$

Using the invariance of f under S , we derive

$$\Phi_l^f(f(lz), f(z)) = 0.$$

Since $\Phi_l^f(X, f)$ is irreducible in $\mathbf{C}[X, Y]$, we see that $\Phi_l^f(f, X)$ is a multiple of $\Phi_l^f(X, f)$. There exists a polynomial $g(X, Y)$ with

$$\Phi_l^f(f, X) = g(X, f)g(f, X)\Phi_l^f(f, X).$$

The Gauß lemma tells us that we have $g(X, Y) \in \mathbf{Z}[X, Y]$ and hence $g(X, Y) = \pm 1$. For $g(X, Y) = -1$, we get $\Phi_l^f(X, Y) = -\Phi_l^f(Y, X)$ and $\Phi_l^f(X, X) = 0$. Then $X - Y$ would be a factor of $\Phi_l^f(X, Y)$. This contradicts the irreducibility. Hence, we have $g(X, Y) = 1$ and Φ_l^f is symmetric.

If f is invariant under M , we replace z by $-1/z$ in the equation $\Phi_l^f(f(z), f(lz)) = 0$ to obtain

$$\Phi_l^f(f(-1/z), f(-l/z)) = 0.$$

Using the invariance of f under M , we derive

$$\Phi_l^f(f(lz), f(z)) = 0,$$

and the proof proceeds as before. □

Applying the lemma to $f = \gamma_2$, which is invariant under S , we know that $\Phi_l^{\gamma_2}$ starts with $X^{l+1} + Y^{l+1}$. For $l = 5$ we compute

$$\gamma_2(q^{1/3})^6 + \gamma_2(q^{5/3})^6 = q^{-10} + \dots$$

and hence we have

$$\Phi_5^{\gamma_2}(X, Y) = X^6 + Y^6 - X^5Y^5 + \text{lower order terms}.$$

We compute $\gamma_2(q^{1/3})^6 + \gamma_2(q^{5/3})^6 - \gamma_2(q^{1/3})^5\gamma_2(q^{5/3})^5$ and find

$$\Phi_5^{\gamma_2}(X, Y) = X^6 + Y^6 - X^5Y^5 + 1240(X^5Y^2 + Y^2X^5) + \text{lower order terms}.$$

Continuing like this, we find the modular polynomial $\Phi_5^{\gamma_2}$ of level 5 for γ_2 :

$$\begin{aligned} \Phi_5^{\gamma_2}(X, Y) = & X^6 + Y^6 - X^5Y^5 + 1240(X^5Y^2 + X^2Y^5) + 20620X^4Y^4 \\ & + 66211200(X^4Y + XY^4) - 125915650X^3Y^3 \\ & + 654403829760(X^3 + Y^3) + 229282790400X^2Y^2 \\ & - 82577379557376XY + 5209253090426880 \in \mathbf{Z}[X, Y]. \end{aligned}$$

The computation of $\Phi_{11}^{\gamma_2}$ proceeds similarly. The time needed for this computation is negligible.

For the cube root $\beta_1 = 208 \in \mathbf{F}_p$ of $j(E)$, the polynomial $\Phi_5^{\gamma_2}(X, \beta_1) \in \mathbf{F}_p[X]$ has 2 roots, namely 176 and 328. In order to decide which one is $\beta_1^{\bar{p}^5}$ and which one is $\beta_1^{p^5}$, we compute $j(E)^{p^5}$ as in the previous subsection. We find $j(E)^{p^5} = 309 \in \mathbf{F}_p$. We compute $176^3 = 309 \in \mathbf{F}_p$ and conclude that we have

$$\beta_1^{\bar{p}^5} = 328 \quad \text{and} \quad \beta_1^{p^5} = 176.$$

Next we compute the action of $\mathfrak{p}_{11}$ and find

$$b_1 = 208 \xrightarrow{\bar{\mathfrak{p}}_5} 328 \xrightarrow{\mathfrak{p}_{11}} 157 \xrightarrow{\mathfrak{p}_{11}} 208.$$

Hence, $b_1 = 208 \in \mathbf{F}_p$ is the reduction of a class invariant. For completeness sake, we note that we have

$$b_2 = 423 \xrightarrow{\bar{\mathfrak{p}}_5} 289 \xrightarrow{\mathfrak{p}_{11}} 258 \xrightarrow{\mathfrak{p}_{11}} 583,$$

and

$$b_3 = 583 \xrightarrow{\bar{\mathfrak{p}}_5} 597 \xrightarrow{\mathfrak{p}_{11}} 192 \xrightarrow{\mathfrak{p}_{11}} 423,$$

showing that b_2 and b_3 are not reductions of class invariants. In fact, we knew this already. If b_2 would be the reduction of a class invariant, then ζ_3 would be contained in the Hilbert class field H and the extension $H/\mathbf{Q}$ would be ramified at 3.

As smooth ideal (α) that we use for the map $\rho_\alpha : X_D(\mathbf{C}_p) \rightarrow X_D(\mathbf{C}_p)$ we pick $\alpha = 3 + \pi_p$, just like we did for the j -function. We compute the cycle

$$137 \xrightarrow{\mathfrak{p}_5} 309 \xrightarrow{\mathfrak{p}_5} 532 \xrightarrow{\mathfrak{p}_5} 214 \xrightarrow{\mathfrak{p}_5} 137$$

of j -invariants over $\mathbf{F}_p$ as before. Using this cycle, we compute the corresponding cycle for b_1 :

$$208 \xrightarrow{\mathfrak{p}_5} 176 \xrightarrow{\mathfrak{p}_5} 562 \xrightarrow{\mathfrak{p}_5} 328 \xrightarrow{\mathfrak{p}_5} 208.$$

Next we lift $E/\mathbf{F}_p$ arbitrarily to $E_1/\mathbf{Q}_p$ and compute a cube root of $j(E_1) \in \mathbf{Q}_p$ that reduces to b_1 modulo p . We denote this cube root again by b_1 . Taking the lift E_1 defined by $Y^2 = X^3 + X + 56$, we find $b_1 = 208 - 43p \in \mathbf{Z}_p$. The polynomial $\Phi_5^{72}(X, b_1) \in \mathbf{Z}_p[X]$ has 2 roots, one of which reduces to $176 \in \mathbf{F}_p$. We compute the ‘cycle’ corresponding to b_1 :

$$208 - 43p \xrightarrow{\mathfrak{p}_5} 176 - 74p \xrightarrow{\mathfrak{p}_5} 562 - 118p \xrightarrow{\mathfrak{p}_5} 328 + 104p \xrightarrow{\mathfrak{p}_5} 208 - 255p.$$

Note that we don’t use the modular polynomial for the j -function for this computation over $\mathbf{Q}_p$: the modular polynomials for γ_2 suffice. We lift $j(E)$ to $j_1 = \rho_\alpha(b_1)^3 = 137 + 121p \in \mathbf{Z}_p$. We update j_1 according to the ‘Newton formula’ (7.2) and find that $j(E_2) = 137 - 15p$ is the two digit approximation of the j -invariant of the canonical lift. Accordingly, we see that $b_2 = 208 + 140p$ is the two digit approximation of the canonical lift $\tilde{b}$.

We continue this process until we have computed $\tilde{b}$ up to $45/3 = 15$ p -adic digits. We find

$$\begin{aligned}\tilde{b}_1 = & 208 + 140p + 96p^2 - 249p^3 - 37p^4 - 34p^5 + 40p^6 - 222p^7 + 142p^8 \\ & - 139p^9 - 159p^{10} + 118p^{11} + 225p^{12} + 103p^{13} - 234p^{14} - 140p^{15} + O(p^{16}).\end{aligned}$$

Computing the conjugates for $\tilde{b}$ under $\text{Pic}(\mathcal{O}_D)$ proceeds just as in the computation of P_D for the j -function, using the modular polynomials $\Phi_5^{\gamma_2}$ and $\Phi_{11}^{\gamma_2}$. In the end we expand the polynomial of degree 8 to find the polynomial

$$\begin{aligned}P_{-2419}^{\gamma_2} &= \prod_{[I] \in \text{Pic}(\mathcal{O}_D)} (X - \tilde{b}^I) \\ &= X^8 + 23344847974866451112256X^7 \\ &\quad + 431537460087154644582865920X^6 \\ &\quad + 20716070070453749000805185224704X^5 \\ &\quad + 1917235980323082783654716721070080X^4 \\ &\quad + 300822183549446154017184276258226176X^3 \\ &\quad + 4961110370685787305744112066133753856X^2 \\ &\quad + 583359477884330290298868497942826713088X \\ &\quad - 6798285426905262621977757780174169964544 \in \mathbf{Z}[X].\end{aligned}$$

The polynomial $P_D^{\gamma_2} \in \mathbf{F}_q[X]$ splits completely. Let $x \in \mathbf{F}_q$ be a root. A curve $E/\mathbf{F}_q$ with j -invariant $j(E) = x^3$ has endomorphism ring $\mathcal{O}$. Either E or its quadratic twist has exactly N points over $\mathbf{F}_q$.

► Using the Weber function

Even better results can be obtained by using the Weber function $\mathfrak{f}(z) = \zeta_{48}^{-1} \frac{\eta(\frac{z+1}{2})}{\eta(z)}$ of level 48. The minimal polynomial

$$\Psi_{\mathfrak{f}} = (X^{24} - 16)^3 - jX^{24} \in \mathbf{Z}[j, X]$$

of $\mathfrak{f}$ over $\mathbf{Q}(j)$ has degree 72. As noted in chapter 6, the Weber function $\mathfrak{f}$ yields class invariants for discriminants $D \equiv 1 \pmod{8}$ with $3 \nmid D$. In the case $D \equiv 5 \pmod{8}$ and $3 \nmid D$, we get class invariants when we evaluate $\mathfrak{f}$ in an appropriate generator ω for the order $\mathcal{O}_{4D} = \mathbf{Z}[\omega]$ of conductor 2, cf. theorem 6.4.

We have $D = -2419 \equiv 5 \pmod{8}$. In fact, if we want to construct an elliptic curve with N points with N odd, then we will in practice always have $D \not\equiv 1 \pmod{8}$. Indeed, writing $4N = x^2 - Dy^2$ with $x, y \in \mathbf{Z}$ for a fundamental discriminant D with $N + 1 + x$ prime, we see that D cannot be congruent to 1 modulo 8.

We will compute $P_{4D}^{\mathfrak{f}} \in \mathbf{Z}[X]$. The polynomial $P_{4D}^{\mathfrak{f}}$ is a generating polynomial for the ray class field H_2 of conductor 2 of $K = \mathbf{Q}(\sqrt{D})$. First we find a small

prime l for which there exists an elliptic curve $E/\mathbf{F}_l$ with CM by $\mathcal{O}_{4D}$. The smallest $t > 0$ for which $(t^2 - 4D)/4$ is prime is $t = 4$, leading to the prime value $l = 2423$. Applying the naïve algorithm, we find that the elliptic curve $E/\mathbf{F}_l$ defined by

$$Y^2 = X^3 + X + 14$$

of j -invariant $1663 \in \mathbf{F}_l$ has trace of Frobenius 4. Since E has only one $\mathbf{F}_l$ -rational 2-torsion point, we have $\text{End}(E) \cong \mathcal{O}_{4D}$. In accordance with theorem 6.6, the polynomial

$$\Phi_f(X, j(E)) = (X^{24} - 16)^3 - j(E)X^{24} \in \mathbf{F}_l[X]$$

has 2 roots, namely $\pm 321 \in \mathbf{F}_l$. Both are reductions of class invariants. We will work with the root $b = 321$.

The Picard group $\text{Pic}(\mathcal{O}_{4D})$ is cyclic of order $3 \cdot 8 = 24$. We represent the elements as reduced binary quadratic forms and using the same notation as in formula (7.1), we compute $k \approx 107$. Hence, we need to know the canonical lift $\tilde{b} \in \mathbf{Q}_l$ of b in $[k/72] = 2$ l -adic digits accuracy. For the map $\rho_\alpha : X_{4D}(\mathbf{C}_l) \rightarrow X_{4D}(\mathbf{C}_l)$ we take $\alpha = \pi_l - 36$ of norm $3575 = 5^2 \cdot 11 \cdot 13$. The ideal (α) factors as

$$(\alpha) = \mathfrak{p}_5^2 \cdot \mathfrak{p}_{11} \cdot \mathfrak{p}_{13} = (5, \pi_l - 1)^2 \cdot (11, \pi_l - 3) \cdot (13, \pi_l - 10).$$

We start by computing the cycle of j -invariants over $\mathbf{F}_l$ for the map $\bar{\rho}_\alpha : \text{Ell}_{4D}(\mathbf{F}_l) \rightarrow \text{Ell}_{4D}(\mathbf{F}_l)$:

$$j(E) = 1663 \xrightarrow{\mathfrak{p}_5} 2355 \xrightarrow{\mathfrak{p}_5} 347 \xrightarrow{\mathfrak{p}_{11}} 1500 \xrightarrow{\mathfrak{p}_{13}} 1663.$$

To compute the corresponding cycle for the ‘Weber value’ b , we need to know the modular polynomials Φ_l^f of level $l = 5, 11, 13$ for f . The computation of these polynomials proceeds similarly to the computation of the modular polynomials for γ_2 . As f is invariant under $S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, the polynomials Φ_l^f are symmetric by lemma 7.1. We compute

$$\begin{aligned} \Phi_5^f(X, Y) &= X^6 + Y^6 + 4XY - X^5Y^5 \\ \Phi_{11}^f(X, Y) &= X^{12} + Y^{12} + 32XY - 88X^3Y^3 + 88X^5Y^5 - 44X^7Y^7 + 11X^9Y^9 \\ &\quad - X^{11}Y^{11} \\ \Phi_{13}^f(X, Y) &= X^{14} + Y^{14} + 64XY + 13(Y^2X^{12} + X^2Y^{12}) + 52(X^{10}Y^4 + Y^{10}X^4) \\ &\quad + 78(Y^6X^8 + Y^8X^6) - X^{13}Y^{13}. \end{aligned}$$

The polynomial $\Phi_5^f(X, b) \in \mathbf{F}_l[X]$ has 2 roots, namely 171 and 1600. For the root $r = 171$ we have $(r^{24} - 16)^3 - 2355r^{24} = 0$, so we have $b^{\mathfrak{p}_5} = r = 171 \in \mathbf{F}_l$. We compute the cycle for b :

$$b = 321 \xrightarrow{\mathfrak{p}_5} 171 \xrightarrow{\mathfrak{p}_5} 1665 \xrightarrow{\mathfrak{p}_{11}} 150 \xrightarrow{\mathfrak{p}_{13}} 321.$$

Next we choose an arbitrary lift of the curve $E/\mathbf{F}_l$ to $E_1/\mathbf{Q}_l$. We take the lift defined by $E_1 : Y^2 = X^3 + X + 14$ of j -invariant $j(E_1) = 1663 + 176l \in \mathbf{Q}_l$. This leads to the lift $b_1 = 321 + 618l$ of b . Just like we did in the previous subsection for γ_2 , we now compute the ‘cycle’ for b_1 :

$$321 + 618l \xrightarrow{\mathbf{p}_5} 171 - 1073l \xrightarrow{\mathbf{p}_5} 1665 + 696l \xrightarrow{\mathbf{p}_{11}} 150 - 223l \xrightarrow{\mathbf{p}_{13}} 321 + 440l.$$

The value $b_1^{(\alpha)} = 321 + 440l$ leads to $\rho_\alpha(j(E_1)) = 1663 + 629l$. We update this according to the Newton formula (7.2) and obtain $j(\tilde{E}) = 1663 - 1025l$. This is the canonical lift in 2 digits accuracy. We lift b to the root $\tilde{b} = 321 - 381l$ of $\Psi_f(X, j(\tilde{E})) \in \mathbf{Z}_l[X]$. The canonical lift $\tilde{b}$ is accurate in 2 l -adic digits.

We compute the conjugates of $\tilde{b}$ under $\text{Gal}(H_2/\mathbf{Q}(\sqrt{D})) \cong \text{Pic}(\mathcal{O}_{4D})$ by employing the modular polynomials Φ_5^f and Φ_{11}^f just like before. In the end we expand the polynomial

$$P_{4D}^f = \prod_{[I] \in \text{Pic}(\mathcal{O}_D)} (X - \tilde{b}^I) \in \mathbf{Z}[X]$$

and find

$$\begin{aligned} P_{-4 \cdot 2419}^f = & X^{24} + 624X^{23} - 756X^{22} - 820X^{21} + 7500X^{20} + 36424X^{19} \\ & + 91904X^{18} + 183248X^{17} + 286784X^{16} + 266736X^{15} + 111024X^{14} \\ & + 67328X^{13} + 133536X^{12} - 23552X^{11} - 150592X^{10} + 316736X^9 \\ & + 792896X^8 + 292224X^7 - 596736X^6 - 643840X^5 - 56832X^4 \\ & + 206080X^3 + 99584X^2 + 8192X + 256 \in \mathbf{Z}[X]. \end{aligned}$$

We return to the problem of constructing an elliptic curve with exactly

$$N = 123456789012345678901234567890123456789012345678901234568197$$

points. We have $x^2 - Dy^2 = 4N$, with $q = N + 1 + x$ prime. The following lemma describes the splitting behaviour of $P_{4D}^f \in \mathbf{F}_q[X]$.

LEMMA 7.2. *Let $p \in \mathbf{Z}$ be a prime that satisfies $t^2 - 4p = \Delta$ for $t \in \mathbf{Z}$ and a negative discriminant $\Delta \equiv 5 \pmod{8}$. Let D be the field discriminant of $\mathbf{Q}(\sqrt{\Delta})$. Then the polynomial $P_{4D}^f \in \mathbf{F}_p[X]$ splits into irreducible cubic factors.*

PROOF. The polynomial $P_{4D}^f \in \mathbf{Z}[X]$ is a generating polynomial for the ray class field H_2 of $K = \mathbf{Q}(\sqrt{D})$. The extension $H_2/\mathbf{Q}$ is Galois, and we need to show that we have $f_{\mathfrak{p}/p} = 3$. Here, $f_{\mathfrak{p}/p}$ denotes the residue class degree of a prime $\mathfrak{p}|p$ lying over p .

By assumption, the prime p splits completely in the Hilbert class field H of K . Since 2 is inert in K , the extension H_2/H has degree 3 and we have to show that p does not split completely in H_2 .

By class field theory, the prime $\pi_p = (t + \sqrt{\Delta})/2 \in \mathcal{O}_K$ of norm p splits completely in H_2/K if and only if we have $\pi_p \equiv 1 \pmod{(2)}$. The prime 2 is inert in K and t is odd, hence we have $\pi_p \not\equiv 1 \pmod{(2)}$. This shows that $P_{4D}^f \in \mathbf{F}_p[X]$ splits into irreducible cubic factors. $\square$

We see that $P_{4D}^f \in \mathbf{F}_q[X]$ splits into irreducible cubic factors. In order to construct a curve with N points over $\mathbf{F}_q$, we compute a root $\beta \in \mathbf{F}_{q^3}$ of $P_{4D}^f \in \mathbf{F}_{q^3}[X]$. An elliptic curve E with j -invariant $j = (\beta^{24} - 16)^3/\beta^{24} \in \mathbf{F}_{q^3}$ has endomorphism ring $\mathcal{O}_{4D}$. One of the three 2-isogenous curves of E has CM by $\mathcal{O}_D$ and is defined over $\mathbf{F}_q$. Hence, the classical modular polynomial $\Phi_2(X, j) \in \mathbf{F}_{q^3}[X]$ has a root j' in $\mathbf{F}_q$. An elliptic curve $E'/\mathbf{F}_q$ with j -invariant j' has CM by $\mathcal{O}_D$. Either E' or its quadratic twist has N points.

We illustrate lemma 7.2 by working over $\mathbf{F}_p$, with $p = 607$, rather than over $\mathbf{F}_q$. The polynomial $P_{4D}^f \in \mathbf{F}_p[X]$ factors as

$$\begin{aligned} P_{4D}^f &= (X^3 + 55X^2 + 343X + 2) \cdot (X^3 + 223X^2 + 337X + 2) \cdot \\ &\quad (X^3 + 319X^2 + 405X + 2) \cdot (X^3 + 402X^2 + 229X + 2) \cdot \\ &\quad (X^3 + 406X^2 + 168X + 2) \cdot (X^3 + 526X^2 + 177X + 2) \cdot \\ &\quad (X^3 + 544X^2 + 429X + 2) \cdot (X^3 + 577X^2 + 404X + 2) \in \mathbf{F}_p[X]. \end{aligned}$$

We represent the field $\mathbf{F}_{p^3}$ as $\mathbf{F}_p(\delta)$ with δ a zero of $f = X^3 + 55X^2 + 343X + 2 \in \mathbf{F}_p[X]$. An elliptic curve with j -invariant $j = (\delta^{24} - 16)^3/\delta^{24} \in \mathbf{F}_{p^3}$ has endomorphism ring $\mathcal{O}_{4D}$. The polynomial $\Phi_2(j, X) \in \mathbf{F}_{p^3}[X]$ has one root in $\mathbf{F}_p$, namely 298. An elliptic curve with j -invariant 298 $\in \mathbf{F}_p$ has endomorphism ring $\mathcal{O}_D$.

7.2 Large group orders

The cryptographic curve from the previous section is a rather small example, and our algorithm is capable of handling much larger inputs $N \in \mathbf{Z}_{\geq 1}$. Chapter 4 reports on the construction of a curve with exactly $10^{2004} + 4863 = \text{nextprime}(10^{2004})$ points. In this section we give more examples of constructing curves of prescribed order N .

The inputs N we choose in this section are ‘random’. To obtain suitable random integers, we encode all letters of the alphabet via the simple scheme $A \mapsto 01, B \mapsto 02, \dots, Z \mapsto 26$. The space is mapped to 00. For instance, if we take the title of this thesis

CONSTRUCTING ELLIPTIC CURVES OF PRESCRIBED ORDER

we get

$$N = 031514192018210320091407000512120916200903000321 \\ 182205190015060016180519031809020504001518040518.$$

We will construct a curve with exactly N points. The algorithm from chapter 4 requires the factorisation of N . As N is only of size 10^{95} , this is doable by the current factorisation algorithms. We find

$$N = 2 \cdot 7 \cdot 11 \cdot 2435345113 \cdot 15279732667 \cdot 496139791093 \cdot 7356751787309663 \\ \cdot 1506673061548025358525712547050019127952141103.$$

Next we look for a solution $x, y \in \mathbf{Z}_{\geq 1}$ to the equation $x^2 - Dy^2 = 4N$ for varying D with $p = N + 1 - x$ prime. The ‘smallest’ D that admits a solution is $D = -52477$. We have

$$x^2 + 52477y^2 = 4N$$

and

$$p = N + 1 - x = 31514192018210320091407000512120916200903000321 \\ 459768862031516742958137970371530378332687568921$$

is prime for

$$x = -277563672016456726777618938562509874331169528402 \\ y = 966453100641073489679726097861861130144188278.$$

Since $D \equiv 3 \pmod{4}$ is not a discriminant, we compute the class polynomial $P_{4D} \in \mathbf{Z}[X]$. The computation of P_{4D} proceeds just as in section 7.1. Alternatively, we can use the function γ_2 or one of the double η -quotients as explained in section 7.4.

The polynomial $P_{4D} \in \mathbf{F}_p[X]$ splits completely. Any of its zeroes yields a curve with N points. If we put

$$a = 31852030651423751802739993910297528236006693960 \\ 00863090565357151470459503211617791215206527455 \in \mathbf{F}_p,$$

then the curve defined by

$$Y^2 = X^3 + aX - a$$

has exactly N points over $\mathbf{F}_p$.

Another way of producing suitable large random numbers, is to write a date as an integer and work with $N = \text{nextprime}(10^{\text{date}})$.

For the date 6 December, we get $N = \text{nextprime}(10^{612}) = 10^{612} + 1411$. In this case, we get a solution to $x^2 - Dy^2 = 4N$ with $p = N + 1 - x$ prime for $D = -23837083$. For discriminants of this size, it is a bad idea to compute the class polynomial P_D corresponding to the j -function. Better results are obtained if we compute the polynomial P_{4D}^f corresponding to f . We do have to replace D by $4D$, but this is a relatively small price to pay.

The polynomial P_{4D}^f splits into irreducible cubic factors in $\mathbf{F}_p[X]$, cf. lemma 7.2. Using this factorisation, it is an easy matter to construct an elliptic curve over $\mathbf{F}_p$ with N points.

A larger example arises when we select 26 November. This time, we have $N = \text{nextprime}(10^{2611}) = 10^{2611} + 5641$ and we worked with $D = -80783323$. Although it is better to work with for instance the Weber function f in this case, it is still possible to compute the class polynomial P_D in a reasonable amount of time. The polynomial P_D has coefficients up to 32122 decimal digits, and we used 4400 digits in a 20195387-adic algorithm.

7.3 Simple η -quotients

We leave the problem of curve construction, and focus only on the problem of computing a generating polynomial for the Hilbert class field H of an imaginary quadratic number field $K = \mathbf{Q}(\sqrt{D})$. Displaying pages full of large numbers is not always pleasing to the human eye. Hence, we will work with relatively small values of D . Section 7.5 contains an example of the computation for a large discriminant.

For any choice of D , there are many ‘small’ functions that we can choose. We have seen already how to work with γ_2 and f in a p -adic setting. The Weber f -function can be ‘generalized’ in many ways. For a prime $l \in \mathbf{Z}$, define the function $g_l : \mathbf{H} \rightarrow \mathbf{C}$ by

$$g_l(z) = \frac{\eta(z/l)}{\eta(z)}.$$

Here, $\eta(z)$ denotes as usual the classical Dedekind eta function with Fourier expansion

$$\eta(z) = q^{1/24} \cdot \prod_{n=1}^{\infty} (1 - q^n), \quad q = e^{2\pi iz}.$$

We have $g_2(z) = f_1(z)$. Using the transformation rules [47, Theorem 1] for the η -function under an element $M \in \text{SL}_2(\mathbf{Z})$, one proves [25, Chapter 4] that g_l is a modular function of level $24l$ over $\mathbf{Q}(\zeta_{24l})$. In this section we consider the functions g_3 and g_{11} and explain how to work with these functions over $\mathbf{Q}_p$.

► The case $l = 3$

The function $g(z) = g_3(z) = \frac{\eta(z/3)}{\eta(z)}$ was already studied by Weber [65, §72]. Besides g , Weber considers the functions

$$\begin{aligned} g_{(1)}(z) &= \zeta_{24}^{-1} \frac{\eta(\frac{z+1}{3})}{\eta(z)} \\ g_{(2)}(z) &= \frac{\eta(\frac{z+2}{3})}{\eta(z)} \\ g_{(3)}(z) &= \sqrt{3} \frac{\eta(3z)}{\eta(z)}. \end{aligned}$$

The functions $g_{(i)}$ are conjugates of g over $\mathbf{Q}(j)$. Just like we have $\mathfrak{ff}_1 \mathfrak{f}_2 = \sqrt{2}$ for the classical Weber functions, we have $gg_{(1)}g_{(2)}g_{(3)} = \sqrt{3}$.

Class invariants arising from g are well studied [25, Theorem 4.1]. The techniques used are exactly the same as those used in the proof of theorem 6.4. We will illustrate our p -adic techniques by treating the same example $D = -479$ that Gee and Stevenhagen used to explain their complex analytic method for computing class invariants [27]. If we put $\omega = \frac{-1+\sqrt{D}}{2}$, then the value $\beta = \zeta_3^2 g_{(2)}^2(\omega)/\sqrt{-3}$ generates the Hilbert class field H of $K = \mathbf{Q}(\sqrt{D})$. Furthermore, the minimal polynomial of β over $\mathbf{Q}$ has integer coefficients.

For the p -adic computation, we compute the minimal polynomial

$$\Psi_h(j, X) = (X^6 + 27)(X^6 + 3)^3 - jX^6 \in \mathbf{Z}[j, X]$$

of $h = g^2$ over $\mathbf{Q}(j)$. This polynomial can in fact already be found in Weber [65, §72]. The functions h and $g_{(2)}^2$ are conjugates over $\mathbf{Q}(j)$. Since we know that $\beta = \zeta_3^2 g_{(2)}^2(\omega)/\sqrt{-3}$ is a class invariant, we know that there is a root $x \in \mathbf{C}$ of $\Psi_h(j(\omega), X) \in H[X]$ with the property that $x/\sqrt{-3}$ is a class invariant.

As in section 6.6, the logarithmic height of the coefficients of the minimal polynomial $P_D^h \in \mathbf{Z}[X]$ of the class invariant β is a factor

$$r(h) = \frac{\deg_X(\Psi_h(j, X))}{\deg_j(\Psi_h(j, X))} = \frac{24}{1} = 24$$

smaller than that of P^j .

The first thing we do is find a prime p and an elliptic curve $E/\mathbf{F}_p$ that has CM by $\mathcal{O} = \mathcal{O}_D$. Since we have $D \equiv 1 \pmod{8}$, the equation $t^2 - D = 4p$ has no solution with p prime. Hence, we look for the smallest integer $t > 0$ for which $(t^2 - 4D)/4$

is four times a prime p . We find $t = 8$ and $p = 487$. Applying the naïve algorithm from chapter 2, we find that the curve $E'/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + 253X - 253$$

of j -invariant $j(E) = 59 \in \mathbf{F}_p$ has trace of Frobenius 8. Since the polynomial $X^3 + 253X - 253 \in \mathbf{F}_p[X]$ has only one root, the curve E' has only one $\mathbf{F}_p$ -rational 2-torsion point and does not have CM by $\mathcal{O}$. The 2-isogenous curve $E/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + 219X + 306$$

has trace of Frobenius 8 and CM by $\mathcal{O}$.

The polynomial $\Psi_h(j(E), X) = X^{24} + 36X^{18} + 270X^{12} + 697X^6 + 729 \in \mathbf{F}_p[X]$ has six roots, namely $\pm 26, \pm 188, \pm 214 \in \mathbf{F}_p$. Since we have $\zeta_3 \in H$, for all six roots b the value $b/\sqrt{-3} \in \mathbf{F}_p$ is the reduction of a class invariant. Unfortunately, the minimal polynomial of the canonical lift $\tilde{b}/\sqrt{-3}$ has integer coefficients for only *two* choices of $\pm b$. We can just compute the minimal polynomial for all 3 choices of $\pm b$, and see whether the polynomial has integer coefficients. For the choice $b = 188 \in \mathbf{F}_p$ we get integer coefficients, and this is the root we will work with in this example.

Having fixed the root $b = 188 \in \mathbf{F}_p$, we show how to compute the canonical lift $\tilde{b} \in \mathbf{Q}_p$. We compute the required precision of $k = 47$ p -adic digits for the computation of P_D as in formula 7.1. The required precision for the computation of P_D^h is $\lceil k/24 \rceil = 2$ p -adic digits.

As smooth element $\alpha \in \mathcal{O} \setminus \mathbf{Z}$ for the map $\rho_\alpha : X_D(\mathbf{C}_p) \rightarrow X_D(\mathbf{C}_p)$ we take $\alpha = \pi_p - 2$ of norm $475 = 5^2 \cdot 19$. Writing $p = \pi_p \bar{\pi}_p \in \mathcal{O}$, we factor

$$(\alpha) = \mathfrak{p}_5^2 \cdot \mathfrak{p}_{19} = (5, \pi_p - 2)^2 \cdot (19, \pi_p - 2)$$

and compute the cycle of j -invariants for the map $\bar{\rho}_\alpha : \text{Ell}_D(\mathbf{F}_p) \rightarrow \text{Ell}_D(\mathbf{F}_p)$:

$$j(E) = 59 \xrightarrow{\mathfrak{p}_5} 481 \xrightarrow{\mathfrak{p}_5} 410 \xrightarrow{\mathfrak{p}_{19}} 59.$$

In order to compute the corresponding cycle for the ‘ h -values’ b_1 , we need the modular polynomials Φ_5^h and Φ_{19}^h . These polynomials are computed by relating the Fourier expansions of $h(z)$ and $h(lz)$, just as before. As an example, we give the polynomial Φ_5^h of level 5:

$$\Phi_5^h(X, Y) = X^6 + Y^6 - X^5Y^5 - 10X^4Y^4 - 45X^3Y^3 - 90X^2Y^2 - 81XY.$$

These polynomials are easily computed ‘on the fly’, but in fact the computation of the modular polynomials Φ_l^h must be seen as a *precomputation*, since this computation is independent of the discriminant.

The polynomial $\Phi_5^h(X, b) \in \mathbf{F}_p[X]$ has 2 roots, namely $17, 381 \in \mathbf{F}_p$. We have $\Psi_h(481, 381) = 0$, hence we have $b^{p^5} = 381$. Continuing like this, we compute the cycle

$$b = 188 \xrightarrow{p^5} 381 \xrightarrow{p^5} 472 \xrightarrow{p^{19}} 188$$

over $\mathbf{F}_p$ corresponding to b .

We lift the curve $E/\mathbf{F}_p$ to the curve $E_1/\mathbf{Q}_p$ defined by

$$Y^2 = X^2 + 219X + 306$$

of j -invariant $j_1 = j(E_1) = 59 - 145p \in \mathbf{Q}_p$. We lift b to $b_1 = 188 + 125 \in \mathbf{Q}_p$ as a root of $\Psi_h(j(E_1), X) \in \mathbf{Z}_p[X]$. Next we compute the ‘cycle’ belonging to b_1 . The polynomial $\Phi_5^h(b_1, X) \in \mathbf{Z}_p[X]$ has 2 roots, namely $17 - 114p$ and $381 + 197p$. Since we know the reduction of $b_1^{p^5}$ modulo p , we see that we have $b_1^{p^5} = 381 + 197p \in \mathbf{Q}_p$. Continuing like this, we compute

$$b_1 = 188 + 125p \xrightarrow{p^5} 381 + 197p \xrightarrow{p^5} 472 - 7p \xrightarrow{p^{19}} 188 + 56p = b_1^{(\alpha)}.$$

Knowing $b_1^{(\alpha)}$, we compute $j_1^{(\alpha)} = 59 - 235p$. We update $j(E_1)$ according to the Newton formula (7.2) and find that $j(E_2) = 137 + 31p$ is the two digit approximation of the j -invariant of the canonical lift. We compute $\tilde{b} = 188 + 195p + O(p^2)$.

The class group $\text{Pic}(\mathcal{O})$ of $\mathcal{O}$ is cyclic of order 16 and is generated by a prime of norm 5. We use the modular polynomial $\Phi_5^h \in \mathbf{Z}_p[X, Y]$ to compute the conjugates of $\tilde{b}$ under $\text{Gal}(H/K) \cong \text{Pic}(\mathcal{O})$.

The Artin symbol $[\mathfrak{p}_5, H/K]$ of $\mathfrak{p}_5$ sends $\sqrt{-3}$ to $-\sqrt{-3}$. Hence, the minimal polynomial P_d^h of $\tilde{b}$ is given by

$$P_D^h = (X - \tilde{b}/\sqrt{-3})(X + \tilde{b}^{p^5}/\sqrt{-3})(X - \tilde{b}^{p^5^2}/\sqrt{-3}) \cdots (X - \tilde{b}^{p^5^{14}}/\sqrt{-3})(X + \tilde{b}^{p^5^{15}}/\sqrt{-3}),$$

and we expand

$$\begin{aligned} P_{-479}^h = & X^{16} - 20X^{15} - 127X^{14} - 342X^{13} + 183X^{12} + 427X^{11} - 1088X^{10} \\ & - 794X^9 + 1333X^8 - 794X^7 - 1088X^6 + 427X^5 + 183X^4 - 342X^3 \\ & - 127X^2 - 20X + 1 \in \mathbf{Z}[X]. \end{aligned}$$

► The case $l = 11$

The situation becomes slightly more difficult if we consider the function $g(z) = g_{11}(z) = \frac{\eta(z/11)}{\eta(z)}$. Regarding class invariants, we have the following result [25, Theorem 4.2].

LEMMA 7.3. *Let K be an imaginary quadratic number field of discriminant D with $\gcd(D, 6) = 1$. Let $\omega = \frac{-1+\sqrt{D}}{2}$ be a generator for the maximal order $\mathcal{O} = \mathbf{Z}[\omega]$ of K . Write $f = f_{\mathbf{Q}}$ for the minimal polynomial of ω over $\mathbf{Q}$. If $k \in \mathbf{Z}$ satisfies*

$$f(-k) \equiv 0 \pmod{11} \quad \text{and} \quad k \equiv 0 \pmod{24},$$

then $\zeta_3^2 \left(\frac{\eta(\frac{\omega+k}{11})}{\eta(\omega)} \right)^2$ is a class invariant.

The existence of $k \in \mathbf{Z}$ with $f(-k) \equiv 0 \pmod{11}$ means exactly that 11 should not be inert in $\mathcal{O}$. The function $\zeta_3^2 \left(\frac{\eta(\frac{\omega+k}{11})}{\eta(\omega)} \right)^2$ is conjugate to $h = g^2$ over $\mathbf{Q}(j)$. We see that there is a root of $\Psi_h(j(\omega), X)$ that is a class invariant. Here, Ψ_h denotes the minimal polynomial of h over $\mathbf{Q}(j)$. Again, it is *not* guaranteed that the minimal polynomial of a class invariant $h(\omega)$ for $h = g^2$ has integer coefficients.

The function h is invariant under $\Gamma^0(11)$. The minimal polynomial

$$\begin{aligned} \Psi_h(j, X) = & X^{72} - 5940X^{66} + 14701434X^{60} - (139755j + 19264518900)X^{54} \\ & + (723797800j + 13849401061815)X^{48} \\ & + (67496j^2 - 1327909897380j - 4875351166521000)X^{42} \\ & + (2291468355j^2 + 1036871615940600j + 400050977713074380)X^{36} \\ & - (5346j^3 - 4231762569540j^2 + 310557763459301490j \\ & \quad - 122471154456433615800)X^{30} \\ & + (161201040j^3 + 755793774757450j^2 + 17309546645642506200j \\ & \quad + 6513391734069824031615)X^{24} \\ & + (132j^4 - 49836805205j^3 + 6941543075967060j^2 \\ & \quad - 64815179429761398660j + 104264884483130180036700)X^{18} \\ & + (468754j^4 + 51801406800j^3 + 214437541826475j^2 \\ & \quad + 77380735840203400j + 8041404949359194)X^{12} \\ & - (j^5 - 3732j^4 + 4586706j^3 - 2059075976j^2 + 253478654715j \\ & \quad - 2067305393340)X^6 \\ & + 1771561 \in \mathbf{Z}[j, X] \end{aligned}$$

of $h = g^2$ over $\mathbf{Q}(j)$ can be computed as in [44] by computing the conjugates of h over $\mathbf{Q}(j)$. The polynomial $\Psi_h(j, X)$ is a polynomial in X^6 . The highest power in j is j^5 . Since $X_0(11)$ has genus one, we knew beforehand that $\Psi_h(j, X)$ cannot be a linear polynomial in j .

As in section 6.6, the coefficients of the minimal polynomial $P^h \in \mathbf{Z}[X]$ of a

class invariant $g^2(\omega)$ will be a factor

$$r(g) = \frac{\deg_X(\Psi_h(j, X))}{\deg_j(\Psi_h(j, X))} = \frac{72}{5} \approx 14$$

smaller than the coefficients of P^j .

We also compute some modular polynomials. We have an upper bound

$$(l+1)\deg(h)$$

for the modular polynomial of degree l for h . Since the degree of j in the minimal polynomial Ψ_h is already 5, the degrees of the modular polynomials become rather large. To keep the coefficients small, we consider the modular polynomials Φ_l^g for g rather than for $h = g^2$. For $l = 5$, the polynomial Φ_5^g has degree $30 = 5 \cdot (5 + 1)$ in X and Y , as was to be expected. We compute

$$\begin{aligned} \Phi_5^g(X, Y) = & X^{30} - 5Y^5X^{29} + 10Y^{10}X^{28} + (-895Y^3 - 10Y^{15})X^{27} \\ & + (-3220Y^8 + 5Y^{20})X^{26} + (1030Y^{13} - 605Y - Y^{25})X^{25} \\ & + (-70Y^{18} + 76400Y^6)X^{24} + (-38825Y^{11} + 30Y^{23})X^{23} \\ & + (875Y^{16} - 389620Y^4)X^{22} + (603680Y^9 - 350Y^{21})X^{21} \\ & + (5Y^{26} - 26070Y^{14} + 146410Y^2)X^{20} \\ & + (-420Y^{19} - 4697825Y^7)X^{19} \\ & + (457275Y^{12} - 70Y^{24})X^{18} + (55775Y^{17} + 15080230Y^5)X^{17} \\ & + (-3154470Y^{10} + 875Y^{22})X^{16} \\ & + (-17715610Y^3 - 819470Y^{15} - 10Y^{27})X^{15} \\ & + (-26070Y^{20} + 12810875Y^8)X^{14} + (6748775Y^{13} + 1030Y^{25})X^{13} \\ & + (-124009270Y^6 + 457275Y^{18})X^{12} \\ & + (-6149220Y^{11} - 38825Y^{23})X^{11} \\ & + (-3154470Y^{16} + 1071794405Y^4 + 10Y^{28})X^{10} \\ & + (603680Y^{21} - 620046350Y^9)X^9 + (12810875Y^{14} - 3220Y^{26})X^8 \\ & + (-4697825Y^{19} + 6430766430Y^7)X^7 \\ & + (-124009270Y^{12} + 76400Y^{24})X^6 \\ & + (-5Y^{29} + 15080230Y^{17} - 25937424601Y^5)X^5 \\ & + (1071794405Y^{10} - 389620Y^{22})X^4 \\ & + (-895Y^{27} - 17715610Y^{15})X^3 + 146410Y^{20}X^2 \\ & - 605Y^{25}X + Y^{30} \in \mathbf{Z}[X, Y] \end{aligned}$$

by the same linear algebra techniques that we used for e.g. the classical Weber function f or γ_2 . Alternatively, we can compute ‘Schläfli modular polynomials’ [29] and convert those to modular polynomials. Since the Schläfli polynomials have smaller coefficients, the latter method is faster in this case.

The time needed to compute these modular polynomials is not completely negligible for this function. This should be considered as a *precomputation* however.

We will illustrate the use of h in a p -adic setting by putting $D = -359$ and compute the generating polynomial P_D^h for the Hilbert class field H of $K = \mathbf{Q}(\sqrt{D})$. As in lemma 7.3, we write $\mathcal{O}_D = \mathbf{Z}[\omega]$ with $\omega = \frac{-1+\sqrt{D}}{2}$. For $k = 192 = 8 \cdot 24$ and $k = 216 = 9 \cdot 24$, we have $f_{\mathcal{O}}^\omega(-k) \equiv 0 \pmod{11}$. From lemma 7.3 we derive that there are at least *two* roots of $\Psi_h(j(\omega), X)$ that are class invariants. The corresponding minimal polynomials are complex conjugates.

First we find a prime p and an elliptic curve $E/\mathbf{F}_p$ that has CM by $\mathcal{O} = \mathcal{O}_D$. We have $D \equiv 1 \pmod{8}$ and the smallest $t > 0$ for which there is a solution to $4p = t^2 - 4D$ with p prime is $t = 24$, leading to $p = 503$.

Applying the naïve algorithm, we find that the curve $E/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + 117X - 117$$

of j -invariant $j(E) = 15 \in \mathbf{F}_p$ has trace of Frobenius 24. Since E has its complete 2-torsion defined over $\mathbf{F}_p$, we have $\text{End}(E) = \mathcal{O}$.

We know that the polynomial $\Psi_h(X, j(E)) \in \mathbf{F}_p[X]$ has at least 2 roots. We will work with the polynomial $\Psi_g(X, j(E)) = \Psi_h(X^2, j(E)) \in \mathbf{F}_p[X]$ instead. The reason is that we have computed the modular polynomial for g rather than for $h = g^2$. The polynomial $\Psi_g(X, j(E))$ has four roots, namely $\pm 31, \pm 163 \in \mathbf{F}_p$. The *squares* $48, 9 \in \mathbf{F}_p$ of these roots are reductions of class invariants.

Taking $b = 31 \in \mathbf{F}_p$, we will compute the minimal polynomial $P_D^g \in \mathcal{O}[X]$ of the square of the ‘canonical lift’ $\tilde{b} \in \mathbf{Q}_p$. The class group $\text{Pic}(\mathcal{O})$ is cyclic of order 19 and is generated by a prime of norm 5. The required precision for the computation of $\tilde{b}$ is $\lceil k/14 \rceil = 4$ p -adic digits accuracy, where we compute $k \approx 46$ as in formula 7.1.

As smooth element $\alpha \in \mathcal{O} \setminus \mathbf{Z}$ we take the generator $\alpha = 216038\pi_p - 4114983$ of the principal ideal $\mathfrak{p}_5^{19} \subset \mathcal{O}$. Here, $\mathfrak{p}_5 = (5, \pi_p - 11)$ is an ideal of norm 5. Using the techniques from chapter 5 we compute the cycle of j -invariants over $\mathbf{F}_p$ belonging to the map $\bar{\rho}_\alpha : \text{Ell}_D(\mathbf{F}_p) \rightarrow \text{Ell}_D(\mathbf{F}_p)$.

$$j(E) = 15 \xrightarrow{\mathfrak{p}_5} 62 \xrightarrow{\mathfrak{p}_5} 79 \xrightarrow{\mathfrak{p}_5} \dots \xrightarrow{\mathfrak{p}_5} -215 \xrightarrow{\mathfrak{p}_5} 15$$

The polynomial $\Phi_5^g(X, b) \in \mathbf{F}_p[X]$ has 2 roots, namely $96, -124 \in \mathbf{F}_p$. The root 96 satisfies $\Psi_g(-215, 96) = 0$ and the root -124 satisfies $\Psi_g(62, -124) = 0$. We

conclude that we have $b^{p^5} = -124 \in \mathbf{F}_p$. For the next step in the cycle, we compute the roots of $\Phi_5^g(X, -124) \in \mathbf{F}_p[X]$. This polynomial has as *six* roots $2, 9, 31, -205, -112, -85 \in \mathbf{F}_p$. Only the root $x = -205$ satisfies $\Psi_g(79, x) = 0$, so we have $-205 = b^{p^5}$. Continuing like this, we compute the cycle for b :

$$31 \xrightarrow{p^5} -124 \xrightarrow{p^5} -205 \xrightarrow{p^5} -205 \dots \xrightarrow{p^5} -96 \xrightarrow{p^5} -31.$$

We do *not* have $b^{(\alpha)} = b$. This is no surprise, since we know that the *square* $b^2 = 48 \in \mathbf{F}_p$ is the reduction of a class invariant.

We lift $E/\mathbf{F}_p$ to $E_1/\mathbf{Q}_p$ defined by

$$Y^2 = X^3 + 117X - 117$$

of j -invariant $j_1 = j(E_1) = 15 + 241p \in \mathbf{Q}_p$. We lift b to $b_1 = 31 - 203p \in \mathbf{Q}_p$ as a zero of $\Psi_g(j(E_1), X) \in \mathbf{Z}_p[X]$. The polynomial $\Phi_5^g(X, b_1) \in \mathbf{Z}_p[X]$ has one root that reduces to $-124 \in \mathbf{F}_p$. We compute $b_1^{p^5} = -124 + 34p \in \mathbf{Q}_p$, and compute the entire ‘cycle’ for b_1 :

$$31 - 203p \xrightarrow{p^5} -124 + 34p \xrightarrow{p^5} \dots \xrightarrow{p^5} -31 - 2p.$$

Using the value $b_1^{(\alpha)} = -31 - 2p$, we compute $\rho_\alpha(j_1) = 15 + 14p$ as a root of $\Psi_g(X, b_1) \in \mathbf{Z}_p[X]$. Next we update j_1 according to the Newton formula (7.2) and obtain $j_2 = 15 + 403p$. Accordingly, we update b_1 to $b_2 = -31 + 17p$. This is the canonical lift $\tilde{b}$ of b in two p -adic digits accuracy.

Similarly, we compute $\tilde{b} = 31 + 17p - 195p^2 + 3p^3 + O(p^4) \in \mathbf{Q}_p$. To compute the conjugates of $\tilde{b}$ under $\text{Gal}(H/K) \cong \text{Pic}(\mathcal{O}) \cong \langle p_5 \rangle$ we employ the modular polynomial Φ_5^g once more. We compute the 19 conjugates of $\tilde{b}$ and expand the product

$$P_D^h = \prod_{i=1}^{19} \left(X - (\tilde{b}^{p^5})^2 \right) \in \mathbf{Z}_p[X].$$

Note that we *square* every conjugate. We cannot recognize the polynomial P_D^h as an element of $\mathcal{O}[X]$ yet.

To remedy this, we repeat the entire computation where we take as initial value the other root $b_1 = 163 \in \mathbf{F}_p$ of $\Psi_g(X, j(E)) \in \mathbf{F}_p[X]$. This yields another polynomial $(P_D^h)' \in \mathbf{Z}_p[X]$. The polynomials $P_D^h \in H[X]$ and $(P_D^h)' \in H[X]$ are complex conjugates, so their sum

$$\begin{aligned}
P_D^h + (P_D^h)' &= 2X^{19} + 108X^{18} - 2632X^{17} - 26939X^{16} + 41715X^{15} \\
&\quad + 774621X^{14} - 1706313X^{13} - 8416511X^{12} + 14265150X^{11} \\
&\quad + 241642630X^{10} - 760800808X^9 - 638692312X^8 \\
&\quad + 3092914049X^7 + 8791773739X^6 - 37066201629X^5 \\
&\quad - 28472524275X^4 + 149884243485X^3 + 183678558472X^2 \\
&\quad - 65029061950X - 15584009580
\end{aligned}$$

has integer coefficients. We compute a square root $\sqrt{D} = 12 - 21p - 186p^2 + 94p^3 + O(p^4) \in \mathbf{Q}_p$ of D . It is now an easy matter to write the coefficients of $P_D^g \in \mathbf{Z}_p[X]$ in the form $a + b\omega$ with $\omega = \frac{-1+\sqrt{D}}{2}$ and $a, b \in \mathbf{Z}$. We have

$$\begin{aligned}
P_{-359}^g &= X^{19} + (59 + 10\omega)X^{18} + (-1248 + 136\omega)X^{17} \\
&\quad + (-13882 + -825\omega)X^{16} + (17220 + -7275\omega)X^{15} \\
&\quad + (389937 + 5253\omega)X^{14} + (-762853 + 180607\omega)X^{13} \\
&\quad + (-4418582 + -420653\omega)X^{12} + (6112297 + -2040556\omega)X^{11} \\
&\quad + (124345093 + 7047556\omega)X^{10} + (-366954339 + 26892130\omega)X^9 \\
&\quad + (-355263701 + -71835090\omega)X^8 + (1454764983 + -183384083\omega)X^7 \\
&\quad + (4776144800 + 760515861\omega)X^6 + (-18321473928 + 423253773\omega)X^5 \\
&\quad + (-16122069915 + -3771615555\omega)X^4 \\
&\quad + (72171452159 + -5541339167\omega)X^3 \\
&\quad + (95835878728 + 7993198984\omega)X^2 \\
&\quad + (-28800117323 + 7428827304\omega)X \\
&\quad - 7756846897 + 70315786\omega \in \mathcal{O}[X].
\end{aligned}$$

7.4 Double η -quotients

Another generalisation of the Weber function is obtained by considering double η -products, as studied for instance by Enge and Schertz [23] and Cohen [11, Section 6.3]. For primes $p, q \in \mathbf{Z}$, put $s = 24/\gcd(24, (p-1)(q-1))$ and $N = pq$. Define the function $g_{p,q} : \mathbf{H} \rightarrow \mathbf{C}$ by

$$g_{p,q} = \left(\frac{\eta(z/p)\eta(z/q)}{\eta(z)\eta(z/pq)} \right)^s.$$

Using the transformation behaviour [47, Theorem 1] of the η -function under an element $M \in \mathrm{SL}_2(\mathbf{Z})$, we prove that that $g_{p,q}$ is invariant under the congruence

subgroup $\Gamma^0(N) \subset \mathrm{SL}_2(\mathbf{Z})$. One proves [23, Theorem 7] that $g_{p,q}$ is an element of the modular function field F_N of level N over $\mathbf{Q}(\zeta_N)$. Furthermore, the function $g_{p,q}$ is invariant under the Atkin-Lehner involution $z \mapsto -N/z$ associated to the matrix $\begin{pmatrix} 0 & -N \\ 1 & 0 \end{pmatrix} \in \mathrm{GL}_2(\mathbf{Q})$.

THEOREM 7.4. *Let $K = \mathbf{Q}(\sqrt{D})$ be imaginary quadratic of discriminant $D = \mathrm{disc}(K) < 0$. Take two primes p, q satisfying the condition:*

- ◇ *for $p \neq q$, the primes p, q are either split or ramified;*
- ◇ *for $p = q$, the prime p is split.*

Then $g_{p,q}(\omega)$ lies in the Hilbert class field of K when evaluated at an appropriate generator $\omega \in \mathbf{H}$ of the $\mathbf{Z}$ -algebra $\mathcal{O}_D = \mathbf{Z}[\omega]$. Furthermore, for $p, q \neq 2$, the minimal polynomial of $g_{p,q}(\omega)$ over $\mathbf{Q}$ has integer coefficients.

PROOF. This is part of [22, Theorem 3]. □

For any imaginary quadratic field K there is an infinite number of primes p, q such that $g_{p,q}$ yields class invariants for K . The minimal polynomial of a class invariant $g_{p,q}(\omega)$ has smaller coefficients than the minimal polynomial of $j(\omega)$. The difference in size depends on p and q , see [21] for a comparison.

Let $\Psi_{p,q}(j, X) \in \mathbf{Q}[j, X]$ be the minimal polynomial of $g_{p,q}$ over $\mathbf{Q}(j)$. As in section 6.6, the logarithmic height of the coefficients of the minimal polynomial $P^{g_{p,q}} \in \mathbf{Z}[X]$ of a class invariant $g(\omega)$ is a factor

$$r(g_{p,q}) = \frac{\deg_X(\Psi_{p,q}(j, X))}{\deg_j(\Psi_{p,q}(j, X))}$$

smaller than that of P^j . The polynomials $\Psi_{p,q}(j, X)$ for the function $g_{p,q}$ are studied in [23]. Enge and Schertz explicitly give the conjugates, i.e., the other zeroes of $\Psi_{p,q}(j, X) \in \mathbf{C}(j)[X]$ of $g_{p,q}$. Using this, they prove that we have $\Psi_{p,q}(j, X) \in \mathbf{Z}[j, X]$ and that $\Psi_{p,q}$ is an affine model for the modular curve $X_0(pq)$. The degree $\deg_X(\Psi_{p,q}(j, X))$ therefore equals the index $[\mathrm{SL}_2(\mathbf{Z}) : \Gamma_0(pq)]$. The degree $\deg_j(\Psi_{p,q}(j, X))$ is given by $s \frac{(p-1)(q-1)}{12}$.

We will illustrate the use of double η -quotients in a non-archimedean setting by selecting $p = 5$ and $q = 7$ and work with the function

$$g = g_{5,7} = \frac{\eta(z/5)\eta(z/7)}{\eta(z)\eta(z/35)}.$$

For this function, we have

$$r(g) = \frac{(5+1) \cdot (7+1)}{2} = 24,$$

and the size of the coefficients of P^g is a factor 24 smaller than the size of the coefficients of P^j .

We compute the minimal polynomial $\Psi_g(j, X) \in \mathbf{Z}[X, j]$ of g over $\mathbf{Q}(j)$. In accordance with the example in [23], we find

$$\begin{aligned} \Psi_{5,7}(j, X) = & X^{48} + (-j + 708)X^{47} + (35j + 171402)X^{46} \\ & + (-525j + 15185504)X^{45} + (4340j + 248865015)X^{44} \\ & + (-20825j + 1763984952)X^{43} + (52507j + 6992359702)X^{42} \\ & + (-22260j + 19325688804)X^{41} + (-243035j + 42055238451)X^{40} \\ & + (596085j + 70108209360)X^{39} + (-272090j + 108345969504)X^{38} \\ & + (-671132j + 121198179480)X^{37} + (969290j + 155029457048)X^{36} \\ & + (-1612065j + 97918126080)X^{35} + (2493785j + 141722714700)X^{34} \\ & + (647290j - 1509796288)X^{33} + (-3217739j + 108236157813)X^{32} \\ & + (3033590j - 93954247716)X^{31} + (-5781615j + 91135898154)X^{30} \\ & + (1744085j - 108382009680)X^{29} + (1645840j + 66862445601)X^{28} \\ & + (-2260650j - 66642524048)X^{27} + (6807810j + 38019611082)X^{26} \\ & + (-2737140j - 28638526644)X^{25} + (2182740j + 17438539150)X^{24} \\ & + (-125335j - 8820058716)X^{23} + (-1729889j + 5404139562)X^{22} \\ & + (1024275j - 1967888032)X^{21} + (-1121960j + 1183191681)X^{20} \\ & + (395675j - 370697040)X^{19} + (-54915j + 103145994)X^{18} \\ & + (15582j - 42145404)X^{17} + (34755j - 15703947)X^{16} \\ & + (-6475j - 3186512)X^{15} + (1120j - 4585140)X^{14} \\ & + (-176j + 1313040)X^{13} + (j^2 - 1486j - 38632)X^{12} \\ & + (-7j + 399000)X^{11} + (-19j + 211104)X^{10} \\ & + (-9j + 6771)X^8 + (8j - 6084)X^7 + (7j - 5258)X^6 \\ & + (j - 792)X^5 - 105X^4 + 16X^3 + 42X^2 + 12X + 1 \in \mathbf{Z}[j, X]. \end{aligned}$$

The degree in X of $\Psi_g(j, X)$ is indeed $48 = (5 + 1) \cdot (7 + 1)$, and j^2 appears exactly once.

Besides this polynomial, we also need to know modular polynomials $\Phi_l(X, Y) \in \mathbf{Z}[X, Y]$ for various primes l . In order to compute these, we note that we have an upper bound $(l + 1)\deg(g)$ for the degrees in X and Y of $\Phi_l^g(X, Y)$, where $\deg(g)$ is the degree of $g : X_0(35) \rightarrow \mathbf{P}^1$. We know that g generates the function field $\mathbf{C}(X_0(N))$ over $\mathbf{C}(j)$, and hence we have $\deg(g) = \deg_j(\Psi(j, X)) = 2$.

If we compute the modular polynomials, it turns out that the degree is in fact $l + 1$ and not $2(l + 1)$. Since g is invariant under the Atkin-Lehner involution, the polynomials Φ_l^g are symmetric by lemma 7.1. For $l = 2, 3$ we find the following polynomials.

$$\begin{aligned}\Phi_2^g &= X^3 + Y^3 - X^2Y^2 + 2(XY^2 + X^2Y) + XY \\ \Phi_3^g &= X^4 + Y^4 - X^3Y^3 + 3(X^2Y^3 + X^3Y^2) + 3(Y^3X + X^3Y) \\ &\quad + 6(X^2Y^2) - 3(Y^2X + X^2Y) - XY\end{aligned}$$

Since 5 and 7 divide the level 35 of g , we cannot use Φ_5^g and Φ_7^g . We computed all modular polynomials for primes up to 23. The time needed for this computation is a few minutes. Again, this computation should be considered as a precomputation.

Next we will show how to use g to compute a generating polynomial for the Hilbert class field of $K = \mathbf{Q}(\sqrt{D})$ with $D = -1571$. Both 5 and 7 split completely in $\mathcal{O} = \mathcal{O}_D$. Note that since 2 is inert in $\mathcal{O}$, we cannot use the Weber function $\mathfrak{f}$ in this case.

We start by looking for an elliptic curve E defined over some finite field $\mathbf{F}_p$ that has CM by $\mathcal{O}$. The smallest $t > 0$ for which there is a solution to $t^2 - 4p = D$ with p prime is $t = 15$, leading to $p = 449$. We fix p for the remainder of this section.

Applying the naïve algorithm, we find that the curve $E/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + X + 16$$

of j -invariant 383 has trace of Frobenius 15 and consequently CM by $\mathcal{O}_D$. The polynomial $\Psi_g(j(E), X) \in \mathbf{F}_p[X]$ has 4 roots in $\mathbf{F}_p$, namely $b_1 = 62$, $b_2 = 130$, $b_3 = 239$ and $b_4 = 358$. By examining theorem 3 of [22] more closely, we see in fact that all four of these roots are reductions of class invariants. To illustrate our p -adic techniques, we will reprove that $b_1, \dots, b_4$ are reductions of class invariants.

A root $b_i \in \mathbf{F}_p$ is the reduction of a class invariant if it is invariant under the action of the group $(\mathcal{O}/35\mathcal{O})^*/\{\pm 1\} \cong \mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/12\mathbf{Z} \times \mathbf{Z}/12\mathbf{Z}$. We take $\{\pi_p, 2\pi_p - 11, 2\pi_p - 19, -\pi_p - 28\}$ as a generating set for $(\mathcal{O}/35\mathcal{O})^*$. We choose this particular set of generators, because the elements have smooth norm (except π_p).

Since b_i is an element of $\mathbf{F}_p$, it is invariant under the action of π_p . Put $\alpha = 2\pi_p - 11$, the ‘next’ element in our generating set. Then α has order 12 in $(\mathcal{O}/35\mathcal{O})^*$, and the ideal (α) of norm $1587 = 3 \cdot 23^2$ factors as

$$(\alpha) = \mathfrak{p}_3 \cdot \mathfrak{p}_{23}^2 = (3, \pi_p - 1) \cdot (23, \pi_p - 17).$$

We compute the cycle of j -invariants over $\mathbf{F}_p$ for the map $\bar{\rho}_\alpha : \text{Ell}_D(\mathbf{F}_p) \rightarrow \text{Ell}_D(\mathbf{F}_p)$:

$$j(E) = 383 \xrightarrow{\mathfrak{p}_3} 13 \xrightarrow{\mathfrak{p}_{23}} 24 \xrightarrow{\mathfrak{p}_{23}} 383.$$

The modular polynomial $\Phi_3^g(b_1, X) \in \mathbf{F}_p[X]$ has 2 roots, namely $64, 95 \in \mathbf{F}_p$. We check that 64 satisfies $\Psi_{5,7}(13, 64) = 0$, where 13 is the j -invariant of E^{p^5} . Furthermore, the other root 95 does not satisfy $\Psi_{5,7}(13, 95) = 0$. We conclude that we have $b_1^{p^3} = 64 \in \mathbf{F}_p$. Continuing like this, we compute

$$b_1 = 62 \xrightarrow{p_3} 64 \xrightarrow{p_{23}} 34 \xrightarrow{p_{23}} 62.$$

The computation for b_2, b_3, b_4 proceeds similarly and they are also invariant under the action of (α) . The generator $2\pi_p - 19$ also has norm $3 \cdot 23^2$ and we compute

$$b_1 = 62 \xrightarrow{\bar{p}_3} 95 \xrightarrow{\bar{p}_{23}} 63 \xrightarrow{\bar{p}_{23}} 62.$$

The other roots b_2, b_3, b_4 are also invariant under the action of $(2\pi_p - 19)$. Finally, a similar computation shows that all 4 roots are invariant under the action of the last generator $-\pi_p - 28$. This proves that $b_1, \dots, b_4$ are reductions of class invariants.

We will work with $b = b_1 = 62 \in \mathbf{F}_p$. To estimate the required accuracy for the canonical lift $\tilde{b} \in \mathbf{Q}_p$ of b , we compute

$$k = \frac{\pi \sqrt{|D|}}{\log p} \sum_{[a,b,c] \in \mathcal{F}_D^+} \frac{1}{a} \approx 62,$$

as in formula 7.1. Hence, we have to compute $\tilde{b}$ up to $[62/24] = 3$ p -adic digits accuracy. As element α for the map $\rho_\alpha : X_D(\mathbf{C}_p) \rightarrow X_D(\mathbf{C}_p)$ we again take $\alpha = 2\pi_p - 11$ of norm $3 \cdot 23^2$. We lift $E/\mathbf{F}_p$ to the curve $E_1/\mathbf{Q}_p$ defined by $Y^2 = X^3 + X + 16$ of j -invariant $j(E_1) = 383 + 224p \in \mathbf{Q}_p$. This leads to the lift $b_1 = 62 + 45p \in \mathbf{Q}_p$.

We compute the ‘cycle’ for $b_1 \in \mathbf{Q}_p$ corresponding to the map ρ_α :

$$b_1 = 62 + 45p \xrightarrow{p_3} 64 + 175p \xrightarrow{p_{23}} 34 + 6p \xrightarrow{p_{23}} 62 - 198p = b_1^{(\alpha)}$$

The degree two polynomial $\Psi_{5,7}(X, b_1^{(\alpha)}) \in \mathbf{Z}_p[X]$ has roots $131 - 94p + O(p^2)$ and $383 - 119p + O(p^2)$. We conclude that we have $\rho_\alpha(j(E_1)) = 383 - 119p \in \mathbf{Q}_p$. We update this j -value according to the ‘Newton formula’ (7.2)

$$j(E_{k+1}) = j(E_k) - \frac{\rho_\alpha(j(E_k)) - j(E_k)}{(\alpha/\bar{\alpha}) - 1} \quad \text{for } k \in \mathbf{Z}_{\geq 1}.$$

and obtain $j(E_2) = 383 - 98p \in \mathbf{Q}_p$. This is the j -invariant of the canonical lift in two p -adic digits accuracy. We compute $\tilde{b} = 62 - 64p + O(p^2) \in \mathbf{Q}_p$.

Similarly, we compute $j(E_3) = 383 - 98p + 127p^2$ and $\tilde{b} = 62 - 64p + 66p^2$. To compute the conjugates of $\tilde{b}$ under $\text{Pic}(\mathcal{O}) \cong \mathbf{Z}/17\mathbf{Z} \cong \langle \mathfrak{p}_3 \rangle$ we use the modular polynomial Φ_3^g once more. In the end we expand the polynomial

$$P_D^g = \prod_{[I] \in \text{Pic}(\mathcal{O}_D)} (X - \tilde{b}^I) \in \mathbf{Z}[X]$$

and find

$$\begin{aligned} P_{-1571}^g = & X^{17} + 21X^{16} + 918X^{15} - 11046X^{14} + 49849X^{13} - 115187X^{12} \\ & + 112918X^{11} + 168294X^{10} - 275500X^9 + 361744X^8 - 403346X^7 \\ & + 181066X^6 - 10143X^5 - 3403X^4 - 4290X^3 + 1422X^2 \\ & - 71X + 1 \in \mathbf{Z}[X]. \end{aligned}$$

7.5 A large discriminant

The examples in the previous sections are relatively small. This is mostly for esthetic reasons, and our p -adic algorithm easily handles discriminants of size $\approx -10^{10}$. To illustrate this, we take $D = -92806391$, the same discriminant as in [7, Section 6].

We have $D \equiv 1 \pmod{8}$ and $3 \nmid D$, so we can use the classical Weber function $\mathfrak{f}$. Recall that the minimal polynomial $\Psi_{\mathfrak{f}}$ of $\mathfrak{f}$ over $\mathbf{Q}(j)$ is given by

$$\Psi_{\mathfrak{f}}(j, X) = (X^{24} - 16)^3 - jX^{24} \in \mathbf{Z}[j, X].$$

First we compute a prime p and an elliptic curve $E/\mathbf{F}_p$ with $\text{End}(E) = \mathcal{O} = \mathcal{O}_D$. The smallest positive integer $t \in \mathbf{Z}_{\geq 1}$ with $p = (t^2 - 4D)/4$ prime is $t = 132$, leading to $p = 92810747$. Applying the naïve algorithm, we find that the curve $E/\mathbf{F}_p$ defined by

$$Y^2 = X^3 + 1086X - 1086$$

of j -invariant $j(E) = 37202456 \in \mathbf{F}_p$ has trace of Frobenius 132. As E has all of its 2-torsion points defined over $\mathbf{F}_p$, its endomorphism ring is $\mathcal{O}_D$, not $\mathcal{O}_{4D}$.

The polynomial $\Psi_{\mathfrak{f}}(j(E), X) \in \mathbf{F}_p[X]$ has 2 roots, cf. theorem 6.6. Both roots ± 21677132 are reductions of class invariants. We will work with the root $b = 21677132 \in \mathbf{F}_p$.

For the smooth $\mathcal{O}$ -ideal (α) inducing $\rho_{\alpha} : X_D(\mathbf{C}_p) \rightarrow X_D(\mathbf{C}_p)$ we take $(\alpha) = (-420 + \pi_p)$ which factors as

$$(11, 8 + 2\pi_p) \cdot (17, 4 + 2\pi_p)^2 \cdot (23, 16 + 2\pi_p)^2 \cdot (31, 13 + 2\pi_p) \cdot (41, 30 + 2\pi_p).$$

Just as before, we compute the cycle in $\mathbf{F}_p$ for the j -invariants:

$$37202456 \xrightarrow{p_{11}} 4967239 \xrightarrow{p_{17}} \dots \xrightarrow{p_{31}} 21402782 \xrightarrow{p_{41}} 37202456.$$

Using this cycle, we can compute the cycle for b . The polynomial $\Phi_{11}^f(b, X) \in \mathbf{F}_p[X]$ has 2 roots: 32604444 and 60476019. Only the root 60476019 is also a root of $\Psi_f(j(E)^{p_{11}}, X)$, so this is the root we are after. We compute:

$$21677132 \xrightarrow{p_{11}} 60476019 \xrightarrow{p_{17}} \dots \xrightarrow{p_{31}} 53004472 \xrightarrow{p_{41}} 21677132.$$

We lift $E/\mathbf{F}_p$ to $E_1/\mathbf{Q}_p$ by lifting the coefficients of its Weierstraß equation, and we lift b to a root b_1 of $\Psi_f(j(E_1), X) \in \mathbf{Z}_p[X]$. Now we lift the cycle that we had for $b \in \mathbf{F}_p$ to a ‘cycle’ for $b_1 \in \mathbf{Z}_p$ by employing the small modular polynomials for f once more. Since we know that $b_1^{(\alpha)}$ is a root of $\Psi_f(j(E_1)^{(\alpha)}, X)$, we can compute

$$j(E_1^{(\alpha)}) = \frac{((\beta_1^{(\alpha)})^{24} - 16)^3}{(\beta_1^{(\alpha)})^{24}}$$

and use this value to update $j(E_1)$ according to the Newton formula (7.2).

Knowing $j(\tilde{E}) \bmod p^2$, we can lift $b \in \mathbf{F}_p$ to a root of $\Psi_f(j(\tilde{E}), X)$ in $\mathbf{Z}_p$ that is accurate to two p -adic digits. We continue this process of doubling the precision until we have $\tilde{b}$ with sufficient accuracy. The first four cycles yield:

$$\begin{aligned} \tilde{b} &= 21677132 + O(p) \\ &= 21677132 + 28966941p + O(p^2) \\ &= 21677132 + 28966941p + 7010373p^2 + 31182954p^3 + O(p^4) \\ &= 21677132 + 28966941p + 7010373p^2 + 31182954p^3 - 33808617p^4 \\ &\quad + 27519307p^5 - 31601027p^6 - 36195013p^7 + O(p^8) \\ &= 21677132 + 28966941p + 7010373p^2 + 31182954p^3 - 33808617p^4 \\ &\quad + 27519307p^5 - 31601027p^6 - 36195013p^7 - 8331811p^8 \\ &\quad - 33957007p^9 - 18191700p^{10} + 5895954p^{11} - 42670221p^{12} \\ &\quad + 23637278p^{13} - 40784695p^{14} + 7754196p^{15} + O(p^{16}) \in \mathbf{Z}_p. \end{aligned}$$

We expect to need $\lceil 313618/72 \rceil = 4356$ decimals digits of accuracy, so we compute $\tilde{b}$ up to 550 p -adic digits. The class group $\text{Pic}(\mathcal{O})$, which is cyclic of order 15610, is generated by a prime of norm 11. We can thus compute all the conjugates of $\tilde{b}$ under $\text{Gal}(H/K)$ to 550 p -adic digits using the modular polynomial Φ_{11}^f . In the end, we expand the polynomial of degree 15610 to find the class polynomial

$$P_D^f = \prod_{[I] \in \text{Pic}(\mathcal{O})} (X - \tilde{\beta}^I) \in \mathbf{Z}[X].$$

The time needed for this computation was roughly 15 minutes on our standard, 32-bit 2.8 GHz, PC. We implemented the algorithm in the programming language C, and employed several ‘computer science tricks’ – like Horner’s rule for the evaluation of a polynomial – to speed up the program. It takes a computer algebra package like Magma many hours to compute P_D^f .

References

- [1] A. Agashe, K. Lauter, R. Venkatesan, *Constructing elliptic curves with a known number of points over a prime field*, High Primes and Misdemeanours: lectures in honour of the 60th birthday of Hugh Cowie Williams, Fields Institute Communications Series, vol. 41, 2004, pp. 1–17.
- [2] M. Agrawal, N. Kayal, N. Saxena, *Primes is in P*, Ann. of Math. (2) **160** (2004), 781–793.
- [3] H. Baier, *Efficient computation of singular moduli with application in cryptography*, Fundamentals of computation theory, Springer Lecture Notes in Computer Science, vol. 2138, 2001, pp. 71–82.
- [4] R. C. Baker, G. Harman, J. Pintz, *The difference between consecutive primes II*, Proc. London Math. Soc. (3) **83** (2001), 532–562.
- [5] D. Bernstein, *Proving primality in essentially quartic random time*, to appear in Math. Comp.
- [6] R. Bröker, P. Stevenhagen, *Constructing elliptic curves in almost polynomial time*, Preprint, 2005.
- [7] R. Bröker, P. Stevenhagen, *Elliptic curves with a given number of points*, Algorithmic Number Theory Symposium VI, Springer Lecture Notes in Computer Science, vol. 3076, 2004, pp. 117–131.
- [8] J. Buhler, S. Wagon, *Basic algorithms in number theory*, Surveys in Algorithmic Number Theory, Cambridge University Press, 2006.
- [9] J. W. S. Cassels, *Lectures on elliptic curves*, Cambridge University Press, 1991.

- [10] H. Cohen, *A course in computational algebraic number theory*, Springer Graduate Texts in Mathematics, vol. 138, 1993.
- [11] H. Cohen, *Advanced topics in computational number theory*, Springer Graduate Texts in Mathematics, vol. 193, 2000.
- [12] J.-M. Couveignes, T. Henocq, *Action of modular correspondences around CM-points*, Algorithmic Number Theory Symposium V, Springer Lecture Notes in Computer Science, vol. 2369, 2002, pp. 234–243.
- [13] D. A. Cox, *Primes of the form $x^2 + ny^2$* , John Wiley & Sons (1989).
- [14] H. Cramér, *Some theorems concerning prime numbers*, Arkiv för Mat. Astronom. och Fysik **15** (1921), 1–33.
- [15] H. Darmon, F. Diamond, R. Taylor, *Fermat’s last theorem*, Elliptic curves, modular forms & Fermat’s last theorem, International Press Cambridge(1993), pp. 2–140.
- [16] M. Deuring, *Die Typen der Multiplikatorenringe elliptischer Funktionenkörper*, Abh. Math. Sem. Univ. Hamburg **14** (1941), 197–272.
- [17] F. Diamond, J. Im, *Modular forms and modular curves*, Seminar on Fermat’s last theorem, CMS conference proceedings, vol. 17, 1995.
- [18] S. J. Edixhoven, *Stable models of modular curves and applications*, PhD-thesis, Universiteit Utrecht, 1989.
- [19] N. D. Elkies, *Elliptic and modular curves over finite fields and related computational issues*, AMS/IP Stud. Adv. Math. **7** (1998), 21–76.
- [20] A. Enge, *The complexity of class polynomial computation via floating point approximations*, Preprint, 2006.
- [21] A. Enge, F. Morain, *Comparing invariants for class fields of imaginary quadratic fields*, Algorithmic Number Theory Symposium V, Springer Lecture Notes in Computer Science, vol. 2369, 2002, pp. 252–266.

- [22] A. Enge, R. Schertz, *Constructing elliptic curves over finite fields using double eta-quotients*, J. Théor. Nombres Bordeaux **16** (2004), 555–568.
- [23] A. Enge, R. Schertz, *Modular curves of composite level*, Acta Arith. **118** (2005), 129–141.
- [24] J. von zur Gathen, J. Gerhard, *Modern computer algebra*, Cambridge University Press, 1999.
- [25] A. Gee, *Class invariants by Shimura reciprocity*, PhD-thesis, Universiteit van Amsterdam, 2001.
- [26] A. Gee, *Class invariants by Shimura’s reciprocity law*, J. Théor. Nombres Bordeaux **11** (1999), 45–72.
- [27] A. Gee, P. Stevenhagen, *Generating class fields using Shimura reciprocity*, Algorithmic Number Theory, Springer Lecture Notes in Computer Science, vol. 1423, 1998, pp. 441–453.
- [28] G. H. Hardy, E. M. Wright, *An introduction to the theory of numbers*, Oxford University Press, 1938.
- [29] W. B. Hart, *Schläfli modular equations for generalized Weber functions*, Preprint, 2005.
- [30] D. R. Heath-Brown, *Gaps between primes and the pair correlation of zeros of the zeta-function*, Acta Arith. **41** (1982), 85–99.
- [31] G. Hoheisel, *Primzahlprobleme in der Analysis*, Sitz. Preuss. Akad. Wiss. **33** (1930), 3–11.
- [32] A. Ivić, *The Riemann zeta-function*, John Wiley & Sons, 1985.
- [33] A. J. de Jong, *Families of curves and alterations*, Ann. Inst. Fourier (Grenoble) **47** (1997), 599–621.

- [34] D. Kohel, *Endomorphism rings of elliptic curves over finite fields*, PhD-thesis, University of California at Berkeley, 1996.
- [35] J. C. Lagarias, A. M. Odlyzko, *Effective versions of the Chebotarev density theorem*, Algebraic Number Fields, ed. A. Fröhlich, Academic Press, 1977, 409–465.
- [36] S. Lang, *Elliptic curves: diophantine analysis*, Springer, Grundlehren der mathematischen Wissenschaften, vol. 231, 1978.
- [37] S. Lang, *Elliptic functions*, Springer Graduate Texts in Mathematics, vol. 112, 1987.
- [38] A. K. Lenstra, H. W. Lenstra, Jr., *Algorithms in number theory*, Handbook of theoretical computer science, volume A: algorithms and complexity, 1990, 673–715.
- [39] H. W. Lenstra, Jr., *Factoring integers with elliptic curves*, Ann. of Math. (2) **126** (1987), 649–673.
- [40] H. W. Lenstra, Jr., C. Pomerance, *A rigorous time bound for factoring integers*, J. Amer. Math. Soc. **5** (1992), 483–516.
- [41] H. W. Lenstra, Jr., C. Pomerance, *Primality testing with Gaussian periods*, Preprint, 2005.
- [42] K. Mahler, *On a class of non-linear functional equations connected with modular functions*, J. Amer. Math. Soc. **22** (1976), 65–118.
- [43] H. Matsumura, *Commutative ring theory*, Cambridge studies in advanced mathematics, vol. 8, 1986.
- [44] F. Morain, *Calcul du nombre de points sur une courbe elliptique dans un corps fini: aspects algorithmiques*, J. Théor. Nombres Bordeaux **7** (1995), 255–282.
- [45] F. Morain, *Implementing the asymptotically fast version of the elliptic curve primality proving algorithm*, preprint, arXiv:math.NT/0502097, 2005.

- [46] J. Neukirch, *Algebraic number theory*, Springer, Grundlehren der mathematischen Wissenschaften, vol. 322, 1999.
- [47] M. Newman, *Construction and applications of a class of modular functions (II)*, Proc. London Math. Soc. (3) **9** (1959), 373–387.
- [48] Prime gaps: internet page on the largest known prime gaps, see <http://hjem.get2net.dk/jka/math/primegaps/gaps20.htm>.
- [49] R. Schertz, *Weber's class invariants revisited*, J. Théor. Nombres Bordeaux **14** (2002), 325–343.
- [50] R. Schoof, *Counting points on elliptic curves over finite fields*, J. Théor. Nombres Bordeaux **7** (1993), 219–254.
- [51] R. Schoof, *Elliptic curves over finite fields and the computation of square roots mod p* , Math. Comp. **44** (1985), 483–494.
- [52] R. Schoof, *Nonsingular plane cubic curves over finite fields*, J. Combin. Theory Ser. A **46** (1987), 183–211.
- [53] R. Schoof, *The exponents of the groups of points of the reductions of an elliptic curve*, Arithmetic Algebraic Geometry, ed. G. van der Geer, Birkhäuser, New York, 1991.
- [54] A. Selberg, *On the normal density of primes in small intervals, and the difference between consecutive primes*, Archiv for Math. og Natur. **47** (1943), No 6, 87–105.
- [55] J.-P. Serre, *A course in arithmetic*, Springer, Graduate Texts in Mathematics, vol. 7, 1973.
- [56] G. Shimura, *Introduction to the arithmetic theory of automorphic forms*, Princeton University Press, 1971.
- [57] J. H. Silverman, *Advanced topics in the arithmetic of elliptic curves*, Springer Graduate Texts in Mathematics, vol. 151, 1994.

- [58] J. H. Silverman, *The arithmetic of elliptic curves*, Springer Graduate Texts in Mathematics, vol. 106, 1986.
- [59] P. Stevenhagen, *Hilbert's 12th problem, complex multiplication and Shimura reciprocity*, Class field theory – its centenary and prospect, ed. K. Miyake, Adv. studies in pure math., vol. 30, 2001, pp. 161–176.
- [60] H. P. F. Swinnerton-Dyer, *An application of computing to class field theory*, Algebraic Number Theory, ed. J. W. S. Cassels & A. Fröhlich, Academic Press, 1967.
- [61] T. Takagi, *Über eine Theorie des relativ Abel'schen Zahlkörpers*, J. College of Science, **41** no 9. (1920), 1–133, Imperial Univ. of Tokyo; in 'Collected Works', I. Shoten, 1973, pp. 73–167.
- [62] J. T. Tate, *Endomorphisms of abelian varieties over finite fields*, Invent. Math. **2** (1966), 134–144.
- [63] J. T. Tate, *The arithmetic of elliptic curves*, Invent. Math. **23** (1974), 179–206.
- [64] J. Vélú, *Isogénies entre courbes elliptiques*, C. R. Acad. Sci. Paris Sér. A–B **273** (1971), A238–A241.
- [65] H. Weber, *Lehrbuch der Algebra*, dritter Band, Friedrich Vieweg und Sohn, 1908.

Samenvatting

Dit proefschrift gaat over *algoritmen* in de getaltheorie. Het woord algoritme is een verbastering van de naam van de Perzische wiskundige Muhammad ibn Musa al-Khwarizmi (± 790 –850). Zijn boek *Kitab al jabr wa'l-muqabala* gaat over het oplossen van lineaire en kwadratische vergelijkingen, en is van grote invloed geweest op de wiskunde. Het woord algebra is afgeleid van de titel van zijn boek.


Een algoritme is volgens Van Dale een *rekenschema, voorschrift voor het uitvoeren van rekenkundige operaties en de volgorde daarvan*. Een bekende algoritme is de ‘lagere-schoolmethode’ voor het vermenigvuldigen van twee getallen. Voor het bepalen van het product 371×123 maken we een tabel

$$\begin{array}{r} 371 \\ 123 \times \\ 1113 \\ 7420 \\ \hline 37100 + \\ 45633 \end{array}$$

en berekenen $371 \times 123 = 45633$.


Naast de formulering van een algoritme, willen we ook graag weten hoe snel de algoritme is. Een gangbare maat voor de snelheid is het aantal ‘operaties’ dat een algoritme moet uitvoeren: het aantal optellingen en vermenigvuldigingen met cijfers tussen 0 en 9. In het voorbeeld moeten er $3 \times 2 + (1 + 2 \times 2) + 3 = 14$ optellingen en $3 \times 3 = 9$ vermenigvuldigingen gedaan worden voor de berekening van 371×123 . Als we weten hoeveel operaties een algoritme nodig heeft, dan kunnen we voorspellen hoe lang een computer erover zal doen om de algoritme uit te voeren. Onze vermenigvuldigalgoritme is erg snel, en een moderne computer kan het product van twee getallen van duizenden cijfers binnen een fractie van een seconde berekenen.

De eerste 4 hoofdstukken van dit proefschrift gaan over algoritmen om een *elliptische kromme* te construeren met een voorgeschreven *puntenaantal*. Het woord *orde* uit de titel van het proefschrift is de vakterm voor puntenaantal. Hier is de invoer van de algoritme telkens het puntenaantal, en de uitvoer is een elliptische kromme met dat puntenaantal. Een elliptische kromme wordt gegeven door een vergelijking van de vorm $Y^2 = X^3 + aX + b$, en ziet er typisch als volgt uit.



De oplossingen van de vergelijking $Y^2 = X^3 + aX + b$ heten de *punten* van de elliptische kromme. De elliptische kromme in het plaatje heeft oneindig veel punten. Krommen met oneindig veel punten zijn niet de krommen die ons interesseren in dit proefschrift. Wij willen graag krommen construeren met een *eindig* aantal punten. De truc om dit te bewerkstelligen is door niet te werken over de reële getallen, maar over een zogenaamd ‘eindig lichaam’. Rekenen in een eindig lichaam gaat precies hetzelfde als met gewone getallen, maar – de naam zegt het al – in een eindig lichaam hebben we maar eindig veel getallen. De gebruikelijke notatie voor een eindig lichaam is $\mathbf{F}_q$. Hier is de $\mathbf{F}$ een afkorting van het woord ‘field’, de Engelse vakterm voor lichaam. De index q geeft aan hoeveel getallen er in het eindige lichaam zitten. Het is niet ongebruikelijk dat vaktermen niet letterlijk vertaald worden. Het Nederlands volgt in dezen het Duits, waar een lichaam een ‘Körper’ genoemd wordt. In Vlaanderen heet een lichaam juist een ‘veld’, wat altijd tot enige spraakverwarring aanleiding geeft tussen Nederlandse en Vlaamse wiskundigen. . .

Twee gehele getallen zijn *gelijk* in het eindige lichaam $\mathbf{F}_7$ als hun verschil deelbaar is door 7. Zo zijn bijvoorbeeld 60 en 4 gelijk in $\mathbf{F}_7$, aangezien $60 - 4 = 56 = 8 \times 7$ deelbaar is door 7. We kunnen ook een plaatje tekenen van de elliptische kromme $Y^2 = X^3 - X$ over het eindige lichaam $\mathbf{F}_7$. De dikke punten in het plaatje zijn de punten van deze elliptische kromme. Het punt $(X, Y) = (4, 2)$ ligt bijvoorbeeld op de kromme, want $4^3 - 4 = 60$ en $2^2 = 4$ zijn gelijk in het eindige lichaam $\mathbf{F}_7$.



Het *puntenaantal* van deze elliptische kromme is *eindig*. De elliptische krommen met voorgeschreven puntenaantal die we in dit proefschrift construeren zijn altijd elliptische krommen over een eindig lichaam.

Elliptische krommen werden lang als een ‘wiskundig speeltje’ gezien waar je in het dagelijks leven niets mee kon. Dit veranderde in het midden van de jaren ’80 van de 20e eeuw, toen bleek dat elliptische krommen van grote waarde zijn in de *cryptografie*, de kunst van het versleutelen van informatie. Praktisch iedere mobiele telefoon bevat tegenwoordig een elliptische kromme, en ook bij het versturen van bijvoorbeeld creditcardgegevens over het internet worden ze steeds meer gebruikt. Het zal niet als een verrassing komen dat ook vele veiligheidsdiensten, zoals onze eigen AIVD, geïnteresseerd zijn geraakt in elliptische krommen. Bij iedere conferentie waar algoritmen voor elliptische krommen aan bod komen, zijn experts van veiligheidsdiensten aanwezig.

Een van de hoofdresultaten uit dit proefschrift is een algoritme die snel een elliptische kromme kan construeren welke gebruikt kan worden in de cryptografie. Als voorgeschreven puntenaantal kiezen we op dit moment een priemgetal van ongeveer 60 cijfers. Wie een ‘cryptografische kromme’ wil zien, bladere door paragraaf 7.1.

Wiskundigen willen graag een *rigoureuse* afschatting voor het aantal operaties dat een algoritme moet uitvoeren. Anders gezegd: als we beweren dat een algoritme op zijn hoogst x operaties nodig heeft, dan moeten we deze bewering ook kunnen *bewijzen*. Dit is iets anders dan bijvoorbeeld de algoritme 100 keer draaien en daaruit een schatting afleiden voor de snelheid. Enkel met een wiskundig bewijs hebben we absolute zekerheid over de snelheid. Ook als zou blijken dat Galileo het onverhoopt fout had en de zon toch om de aarde draait, dan nog is de algoritme zo snel als we beweren.

Soms is het echter nog niet mogelijk de snelheid van een algoritme rigoureus te analyseren. In sommige gevallen moeten we bijvoorbeeld een onbewezen vermoeden,

zoals de Riemannhypothese, aannemen om tot een sluitend bewijs te komen. Het kan echter zijn dat zelfs dit niet volstaat, en dan rest ons niets anders dan over te gaan op een *heuristiek*, jargon voor nattevingerwerk. We proberen dan op basis van wiskundige stellingen redelijke vermoedens te bedenken die ons in staat stellen een goede schatting te geven van de snelheid van een algoritme. Dit is zeker niet zo mooi als het geven van een echt wiskundig bewijs, maar als de heuristiek precies verklaart wat we in de praktijk zien gebeuren, dan is er eigenlijk niets aan de hand. In zekere zin lijkt deze aanpak een beetje op de natuurkunde: hier worden ook theorieën bedacht die de werkelijkheid moeten *verklaren*.



De algoritme voor het maken van cryptografische krommen wordt gegeven in hoofdstuk 4. Wie hier doorheen bladert, zal een tabel zien met ‘numerical support’. Dit geeft al aan dat de snelheidsanalyse die in hoofdstuk 4 gegeven wordt heuristisch van aard is. De aanname die we voor deze heuristiek nodig hebben, betreft een eigenschap van de verdeling van priemgetallen en heeft niets met elliptische krommen te maken. De uitkomsten van de experimenten worden prima verklaard door de aannames. Er bestaan overigens ook algoritmen om elliptische krommen van voorgeschreven orde te maken die een *rigoureuze* snelheidsanalyse hebben, zie bijvoorbeeld hoofdstuk 3, maar deze zijn erg langzaam. Langzaam betekent hier dat als we de algoritme uit hoofdstuk 3 zouden gebruiken voor het construeren van een cryptografische kromme, we langer op het antwoord zouden moeten wachten dan de tijd die sinds de oerknal verstreken is. Met de algoritme uit hoofdstuk 4 staat het antwoord er vrijwel direct.

Voor de algoritme in hoofdstuk 4 hebben we een subalgoritme nodig om een ‘Hilbertklassenpolynoom’ uit te rekenen. Bij ieder negatief geheel getal hoort een Hilbertklassenpolynoom, een polynoom met gehele coëfficiënten. Deze polynomen zijn vernoemd naar David Hilbert (1862–1943). Hilbert was een van de grondleggers van de *klassenlichamentheorie*, een theorie die tot vele doorbraken in de 20e eeuw heeft geleid. Algoritmen voor de berekening van een Hilbertklassenpolynoom vormen dan ook onderdeel van de computationele klassenlichamentheorie. Ze kunnen – zoals wij gedaan hebben – gebruikt worden om cryptografische elliptische krommen mee te maken, maar er zijn veel meer toepassingen.

Er is een klassieke algoritme om een Hilbertklassenpolynoom uit te rekenen, maar deze heeft een nadeel: we moeten heel erg uitkijken dat er geen *afrondfouten* optreden tijdens de berekening. We *benaderen* de coëfficiënten van het polynoom namelijk met reële getallen. Als we bijvoorbeeld 0,99832 als coëfficiënt berekenen, dan ronden we dit af tot 1. Als er veel afrondfouten optreden, dan kan het bijvoorbeeld zo zijn dat we 0,63781 als antwoord krijgen. Het is een te grote gok dit zomaar op 1 af te ronden. We kunnen afrondfouten vermijden door ‘*p*-adisch’ te rekenen. De *p*-adische getallen zijn uitgevonden door Kurt Hensel (1861–1941). Rekenen met *p*-adische getallen lijkt erg op rekenen met reële getallen, maar het grote voordeel is dat we in onze berekening *geen* afrondfouten kunnen krijgen. Het aanpassen van de klassieke algoritme tot een *p*-adische algoritme voor het berekenen van een Hilbertklassenpolynoom heeft nog aardig wat voeten in de aarde. Het idee achter de algoritme die we in hoofdstuk 5 uitleggen, komt van de Franse wiskundigen Couveignes en Henocq. In een artikel voor een getaltheorieconferentie in Sydney in 2002 gaven zij een manier om een Hilbertklassenpolynoom *p*-adisch te berekenen. Hun artikel was echter redelijk kort en bevatte weinig houvast voor mensen die de algoritme daadwerkelijk wilden uitvoeren op een computer. Er moesten nog aardig wat gaten ingevuld worden voordat er een echt bruikbare algoritme was. Het resultaat hiervan is hoofdstuk 5.

Een ander nadeel is dat de coëfficiënten van het Hilbertklassenpolynoom erg groot zijn. Als we voor onze algoritme als invoer -23 nemen, dan krijgen we

$$X^3 + 3491750X^2 - 5151296875X + 12771880859375$$

als polynoom. De coëfficiënten van dit polynoom staan in geen verhouding tot de bescheiden invoer -23 . Een natuurlijke vraag is of er geen ‘beter’ polynoom is dat we kunnen gebruiken. Deze vraag werd voor het eerst bestudeerd door Heinrich Weber (1842–1913). In het derde deel van zijn *Lehrbuch* geeft hij voorbeelden van betere polynomen met kleinere coëfficiënten. Voor -23 kunnen we bijvoorbeeld ook het

polynoom

$$X^3 - X - 1$$

gebruiken. Een systematische aanpak van de zoektocht naar betere polynomen wordt gegeven door Goro Shimura (1930–). Met zijn ‘wederkerigheidswet’ uit 1971 is het nu een redelijk mechanische berekening geworden om een ‘beter’ polynoom te vinden dan het grote Hilbertklassenpolynoom. Echter: zijn aanpak maakt geen gebruik van p -adische getallen, en dus hebben we weer het probleem van afrondfouten.

In hoofdstuk 6 laten we zien hoe we de technieken van Shimura kunnen uitbreiden tot een p -adische algoritme. Deze algoritme combineert beide verbeteringen van de klassieke algoritme voor het berekenen van een Hilbertklassenpolynoom: ook voor het berekenen van de betere polynomen hoeven we ons nu geen zorgen meer te maken over afrondfouten. Hoofdstuk 7 is volledig gewijd aan voorbeelden van deze nieuwe algoritme. Dit hoofdstuk neemt met zijn 28 pagina’s een flink deel in van dit proefschrift. Het is een zeer bewuste keuze geweest dit laatste hoofdstuk zo lang te maken: algoritmen zijn er immers voor om ook daadwerkelijk uitgevoerd te worden!

Curriculum Vitae

Reinier Bröker is op 12 juli 1979 geboren in Geldermalsen. Na de basisschool ging hij naar het Koningin Wilhelmina College te Culemborg, waar hij in 1997 het VWO-diploma behaalde. Aansluitend begon hij aan een studie wiskunde aan de Vrije Universiteit in Amsterdam. In de eerste paar jaar volgde hij naast de wiskundecolleges ook een aantal vakken bij de opleidingen natuurkunde en informatica. Hij vervulde diverse student-assistentschappen, zowel bij wiskunde als bij informatica. Onder begeleiding van prof. dr. E. J. Ditters en dr. M. L. J. van de Vel schreef hij zijn afstudeerscriptie *Formele groepen binnen de algebraïsche topologie*, waarmee hij in 2001 cum laude afstudeerde.

Na een kortstondige studie Nederlands, begon hij in april 2002 als Assistent in Opleiding aan het Mathematisch Instituut in Leiden. Het promotieonderzoek, waarvan dit proefschrift het resultaat is, werd begeleid door prof. dr. P. Stevenhagen. Naast zijn promotieonderzoek was Reinier lid van de voorlichtingscommissie, en was hij actief betrokken bij de jaarlijkse wetenschapsdag. Hij gaf voordrachten op congressen in Nederland, België, Frankrijk, Duitsland, Griekenland en de Verenigde Staten.

Per 1 september is hij als postdoc verbonden aan het Fields Institute in Toronto en de University of Calgary.