

Computing Hilbert Class Polynomials

Juliana Belding¹, Reinier Bröker², Andreas Enge³, and Kristin Lauter²

¹ Dept. of Mathematics, University of Maryland, College Park, MD 20742, USA
jbelding@math.umd.edu

² Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA
{reinierb,klauter}@microsoft.com

³ INRIA Futurs & Laboratoire d'Informatique (CNRS/UMR 7161)
École polytechnique, 91128 Palaiseau cedex, France
enge@lix.polytechnique.fr

Abstract. We present and analyze two algorithms for computing the Hilbert class polynomial H_D . The first is a p -adic lifting algorithm for inert primes p in the order of discriminant $D < 0$. The second is an improved Chinese remainder algorithm which uses the class group action on CM-curves over finite fields. Our run time analysis gives tighter bounds for the complexity of all known algorithms for computing H_D , and we show that all methods have comparable run times.

1 Introduction

For an imaginary quadratic order $\mathcal{O} = \mathcal{O}_D$ of discriminant $D < 0$, the j -invariant of the complex elliptic curve \mathbf{C}/\mathcal{O} is an algebraic integer. Its minimal polynomial $H_D \in \mathbf{Z}[X]$ is called the *Hilbert class polynomial*. It defines the ring class field $K_{\mathcal{O}}$ corresponding to \mathcal{O} , and within the context of explicit class field theory, it is natural to ask for an algorithm to *explicitly compute* H_D .

Algorithms to compute H_D are also interesting for elliptic curve primality proving [2] and for cryptographic purposes [6]; for instance, pairing-based cryptosystems using ordinary curves rely on complex multiplication techniques to generate the curves. The classical approach to compute H_D is to approximate the values $j(\tau_{\mathfrak{a}}) \in \mathbf{C}$ of the complex analytic j -function at points $\tau_{\mathfrak{a}}$ in the upper half plane corresponding to the ideal classes \mathfrak{a} for the order \mathcal{O} . The polynomial H_D may be recovered by rounding the coefficients of $\prod_{\mathfrak{a} \in \text{Cl}(\mathcal{O})} (X - j(\tau_{\mathfrak{a}})) \in \mathbf{C}[X]$ to the nearest integer. It is shown in [9] that an optimized version of that algorithm has a complexity that is essentially linear in the output size.

Alternatively one can compute H_D using a p -adic lifting algorithm [7,3]. Here, the prime p splits completely in $K_{\mathcal{O}}$ and is therefore relatively large: it satisfies the lower bound $p \geq |D|/4$. In this paper we give a p -adic algorithm for *inert* primes p . Such primes are typically much smaller than totally split primes, and under GRH there exists an inert prime of size only $O((\log |D|)^2)$. The complex multiplication theory underlying all methods is more intricate for inert primes p , as the roots of $H_D \in \mathbf{F}_{p^2}[X]$ are now j -invariants of *supersingular* elliptic curves. In Section 2 we explain how to define the canonical lift of a supersingular elliptic curve, and in Section 4 we describe a method to explicitly compute this lift.

In another direction, it was suggested in [1] to compute H_D modulo several totally split primes p and then combine the information modulo p using the Chinese remainder theorem to compute $H_D \in \mathbf{Z}[X]$. The first version of this algorithm was quite impractical, and in Section 3 we improve this ‘multi-prime approach’ in two different ways. We show how to incorporate inert primes, and we improve the original approach for totally split primes using the class group action on CM-curves. We analyze the run time of the new algorithm in Section 5 in terms of the logarithmic height of H_D , its degree, the largest prime needed to generate the class group of \mathcal{O} and the discriminant D . Our tight bounds on the first two quantities from Lemmata 1 and 2 apply to all methods to compute H_D . For the multi-prime approach, we derive the following result.

Theorem 1. *The algorithm presented in Section 3 computes, for a discriminant $D < 0$, the Hilbert class polynomial H_D . If GRH holds true, the algorithm has an expected run time $O(|D|(\log |D|)^{7+o(1)})$. Under heuristic assumptions, the complexity becomes $O(|D|(\log |D|)^{3+o(1)})$.*

We conclude by giving examples of the presented algorithms in Section 6.

2 Complex Multiplication in Characteristic p

Throughout this section, $D < -4$ is any discriminant, and we write \mathcal{O} for the imaginary quadratic order of discriminant D . Let $E/K_{\mathcal{O}}$ be an elliptic curve with endomorphism ring isomorphic to \mathcal{O} . As \mathcal{O} has rank 2 as a \mathbf{Z} -algebra, there are two isomorphisms $\varphi : \text{End}(E) \xrightarrow{\sim} \mathcal{O}$. We always assume we have chosen the *normalized* isomorphism, i.e., for all $y \in \mathcal{O}$ we have $\varphi(y)^*\omega = y\omega$ for all invariant differentials ω . For ease of notation, we write E for such a ‘normalized elliptic curve,’ the isomorphism φ being understood.

For a field F , let $\text{Ell}_D(F)$ be the set of isomorphism classes of elliptic curves over F with endomorphism ring \mathcal{O} . The ideal group of \mathcal{O} acts on $\text{Ell}_D(K_{\mathcal{O}})$ via

$$j(E) \mapsto j(E)^{\mathfrak{a}} = j(E/E[\mathfrak{a}]),$$

where $E[\mathfrak{a}]$ is the group of \mathfrak{a} -torsion points, i.e., the points that are annihilated by all $\alpha \in \mathfrak{a} \subset \mathcal{O} = \text{End}(E)$. As principal ideals act trivially, this action factors through the class group $\text{Cl}(\mathcal{O})$. The $\text{Cl}(\mathcal{O})$ -action is transitive and free, and $\text{Ell}_D(K_{\mathcal{O}})$ is a principal homogeneous $\text{Cl}(\mathcal{O})$ -space.

Let p be a prime that splits completely in the ring class field $K_{\mathcal{O}}$. We can embed $K_{\mathcal{O}}$ in the p -adic field \mathbf{Q}_p , and the reduction map $\mathbf{Z}_p \rightarrow \mathbf{F}_p$ induces a bijection $\text{Ell}_D(\mathbf{Q}_p) \rightarrow \text{Ell}_D(\mathbf{F}_p)$. The $\text{Cl}(\mathcal{O})$ -action respects reduction modulo p , and the set $\text{Ell}_D(\mathbf{F}_p)$ is a $\text{Cl}(\mathcal{O})$ -torsor, just like in characteristic zero. This observation is of key importance for the improved ‘multi-prime’ approach explained in Section 3.

We now consider a prime p that is *inert* in \mathcal{O} , fixed for the remainder of this section. As the principal prime $(p) \subset \mathcal{O}$ splits completely in $K_{\mathcal{O}}$, all primes of $K_{\mathcal{O}}$ lying over p have residue class degree 2. We view $K_{\mathcal{O}}$ as a subfield of the

unramified degree 2 extension L of \mathbf{Q}_p . It is a classical result, see [8] or [15, Th. 13.12], that for $[E] \in \text{Ell}_D(L)$, the reduction E_p is *supersingular*. It can be defined over the finite field \mathbf{F}_{p^2} , and its endomorphism ring is a maximal order in the unique quaternion algebra $\mathcal{A}_{p,\infty}$ which is ramified at p and ∞ . The reduction map $\mathbf{Z}_L \rightarrow \mathbf{F}_{p^2}$ also induces an embedding $f : \mathcal{O} \hookrightarrow \text{End}(E_p)$. This embedding is not surjective, as it is in the totally split case, since $\text{End}(E_p)$ has rank 4 as a \mathbf{Z} -algebra, and \mathcal{O} has rank 2.

We let $\text{Emb}_D(\mathbf{F}_{p^2})$ be the set of isomorphism classes of pairs (E_p, f) with E_p/\mathbf{F}_{p^2} a supersingular elliptic curve and $f : \mathcal{O} \hookrightarrow \text{End}(E_p)$ an embedding. Here, (E_p, f) and (E'_p, f') are isomorphic if there exists an isomorphism $h : E_p \xrightarrow{\sim} E'_p$ of elliptic curves with $h^{-1}f'(\alpha)h = f(\alpha)$ for all $\alpha \in \mathcal{O}$. As an analogue of picking the normalized isomorphism $\mathcal{O} \xrightarrow{\sim} \text{End}(E)$ in characteristic zero, we now identify (E_p, f) and (E'_p, f') if f equals the complex conjugate of f' .

Theorem 2. *Let $D < -4$ be a discriminant. If p is inert in $\mathcal{O} = \mathcal{O}_D$, the reduction map $\pi : \text{Ell}_D(L) \rightarrow \text{Emb}_D(\mathbf{F}_{p^2})$ is a bijection. Here, L is the unramified extension of \mathbf{Q}_p of degree 2.*

Proof. By the Deuring lifting theorem, see [8] or [15, Th. 13.14], we can lift an element of $\text{Emb}_D(\mathbf{F}_{p^2})$ to an element of $\text{Ell}_D(L)$. Hence, the map is surjective.

Suppose that we have $\pi(E) = \pi(E')$. As E and E' both have endomorphism ring \mathcal{O} , they are isogenous. We let $\varphi_a : E \rightarrow E^a = E'$ be an isogeny. Writing $\mathcal{O} = \mathbf{Z}[\tau]$, we get

$$f' = f^a : \tau \mapsto \overline{\varphi}_a f(\tau) \widehat{\varphi}_a \otimes (\deg \overline{\varphi}_a)^{-1} \in \text{End}(E_p) \otimes \mathbf{Q}.$$

The map $\overline{\varphi}_a$ commutes with $f(\tau)$ and is thus contained in $S = f(\text{End}(E)) \otimes \mathbf{Q}$.

Write $\mathcal{O}' = S \cap \text{End}(E_p)$, and let m be the index $[\mathcal{O}' : f(\text{End}(E))]$. For any $\delta \in \mathcal{O}'$, there exists $\gamma \in \text{End}(E)$ with $m\delta = f(\gamma)$. As $f(\gamma)$ annihilates the m -torsion $E_p[m]$, γ annihilates $E[m]$, thus it is a multiple of m inside $\text{End}(E)$. We derive that δ is contained in $f(\text{End}(E))$, and $\mathcal{O}' = f(\text{End}(E))$. Hence, φ_a is an endomorphism of E , and E and E^a are isomorphic. \square

The *canonical lift* \widetilde{E} of a pair $(E_p, f) \in \text{Emb}_D(\mathbf{F}_{p^2})$ is defined as the inverse $\pi^{-1}(E_p, f) \in \text{Ell}_D(L)$. This generalizes the notion of a canonical lift for ordinary elliptic curves, and the main step of the p -adic algorithm described in Section 4 is to compute \widetilde{E} : its j -invariant is a zero of the Hilbert class polynomial $H_D \in L[X]$.

The reduction map $\text{Ell}_D(L) \rightarrow \text{Emb}_D(\mathbf{F}_{p^2})$ induces a transitive and free action of the class group on the set $\text{Emb}_D(\mathbf{F}_{p^2})$. For an \mathcal{O} -ideal \mathfrak{a} , let $\varphi_a : E \rightarrow E^a$ be the isogeny of CM-curves with kernel $E[\mathfrak{a}]$. Writing $\mathcal{O} = \mathbf{Z}[\tau]$, let $\beta \in \text{End}(E)$ be the image of τ under the normalized isomorphism $\mathcal{O} \xrightarrow{\sim} \text{End}(E)$. The normalized isomorphism for E^a is now given by

$$\tau \mapsto \varphi_a \beta \widehat{\varphi}_a \otimes (\deg \varphi_a)^{-1}.$$

We have $E_p^a = (E^a)_p$ and f^a is the composition $\mathcal{O} \xrightarrow{\sim} \text{End}(E^a) \hookrightarrow \text{End}(E_p^a)$. Note that principal ideals indeed act trivially: φ_a is an endomorphism in this case and, as $\text{End}(E)$ is commutative, we have $f = f^a$.

To explicitly compute this action, we fix one supersingular curve E_p/\mathbf{F}_{p^2} and an isomorphism $i_{E_p} : \mathcal{A}_{p,\infty} \xrightarrow{\sim} \text{End}(E_p) \otimes \mathbf{Q}$ and view the embedding f as an injective map $f : \mathcal{O} \hookrightarrow \mathcal{A}_{p,\infty}$. Let $R = i_{E_p}^{-1}(\text{End}(E_p))$ be the maximal order of $\mathcal{A}_{p,\infty}$ corresponding to E_p . For \mathfrak{a} an ideal of \mathcal{O} , we compute the curve $E_p^\mathfrak{a} = \overline{\varphi}_\mathfrak{a}(E_p)$ and choose an auxiliary isogeny $\varphi_\mathfrak{b} : E_p \rightarrow E_p^\mathfrak{a}$. This induces an isomorphism $g_\mathfrak{b} : \mathcal{A}_{p,\infty} \xrightarrow{\sim} \text{End}(E_p^\mathfrak{a}) \otimes \mathbf{Q}$ given by

$$\alpha \mapsto \varphi_\mathfrak{b} i_{E_p}(\alpha) \widehat{\varphi}_\mathfrak{b} \otimes (\deg \varphi_\mathfrak{b})^{-1}.$$

The left R -ideals $Rf(\mathfrak{a})$ and \mathfrak{b} are left-isomorphic by [22, Th. 3.11] and thus we can find $x \in \mathcal{A}_{p,\infty}$ with $Rf(\mathfrak{a}) = \mathfrak{b}x$. As $y = f(\tau)$ is an element of $Rf(\mathfrak{a})$, we get the embedding $\tau \mapsto xyx^{-1}$ into the right order $R_\mathfrak{b}$ of \mathfrak{b} . By construction, the induced embedding $f^\mathfrak{a} : \mathcal{O} \hookrightarrow \text{End}(E_p^\mathfrak{a})$ is precisely

$$f^\mathfrak{a}(\tau) = g_\mathfrak{b}(xyx^{-1}) \in \text{End}(E_p^\mathfrak{a}),$$

and this is independent of the choice of \mathfrak{b} . For example, if $E_p^\mathfrak{a} = E_p$, then choosing $\varphi_\mathfrak{b}$ as the identity, we find x with $Rf(\mathfrak{a}) = Rx$ to get the embedding $f^\mathfrak{a} : \tau \mapsto i_{E_p}(xyx^{-1}) \in \text{End}(E_p)$.

3 The Multi-prime Approach

This section is devoted to a precise description of the new algorithm for computing the Hilbert class polynomial $H_D \in \mathbf{Z}[X]$ via the Chinese remainder theorem.

Algorithm 1

INPUT: an imaginary quadratic discriminant D

OUTPUT: the Hilbert class polynomial $H_D \in \mathbf{Z}[X]$

0. Let $(A_i, B_i, C_i)_{i=1}^{h(D)}$ be the set of primitive reduced binary quadratic forms of discriminant $B_i^2 - 4A_iC_i = D$ representing the class group $\text{Cl}(\mathcal{O})$. Compute

$$n = \left\lceil \log_2 \left(2.48 h(D) + \pi \sqrt{|D|} \sum_{i=1}^{h(D)} \frac{1}{A_i} \right) \right\rceil + 1, \tag{1}$$

which by [9] is an upper bound on the number of bits in the largest coefficient of H_D .

1. Choose a set \mathcal{P} of primes p such that $N = \prod_{p \in \mathcal{P}} p \geq 2^n$ and each p is either inert in \mathcal{O} or totally split in $K_\mathcal{O}$.
2. For all $p \in \mathcal{P}$, depending on whether p is split or inert in \mathcal{O} , compute $H_D \bmod p$ using either Algorithm 2 or 3.
3. Compute $H_D \bmod N$ by the Chinese remainder theorem, and return its representative in $\mathbf{Z}[X]$ with coefficients in $(-\frac{N}{2}, \frac{N}{2})$.

The choice of \mathcal{P} in Step 1 leaves some room for different flavors of the algorithm. Since Step 2 is exponential in $\log p$, the primes should be chosen as small as possible. The simplest case is to only use split primes, to be analyzed in Section 5. As the run time of Step 2 is worse for inert primes than for split primes, we view the use of inert primes as a practical improvement.

3.1 Split Primes

A prime p splits completely in $K_{\mathcal{O}}$ if and only if the equation $4p = u^2 - v^2D$ has a solution in integers u, v . For any prime p , we can efficiently test if such a solution exists using an algorithm due to Cornacchia. In practice, we generate primes satisfying this relation by varying u and v and testing if $(u^2 - v^2D)/4$ is prime.

Algorithm 2

INPUT: an imaginary quadratic discriminant D and a prime p that splits completely in $K_{\mathcal{O}}$

OUTPUT: $H_D \bmod p$

1. Find a curve E over \mathbf{F}_p with endomorphism ring \mathcal{O} . Set $j = j(E)$.
2. Compute the Galois conjugates $j^{\mathfrak{a}}$ for $\mathfrak{a} \in \text{Cl}(\mathcal{O})$.
3. Return $H_D \bmod p = \prod_{\mathfrak{a} \in \text{Cl}(\mathcal{O})} (X - j^{\mathfrak{a}})$.

Note: The main difference between this algorithm and the one proposed in [1] is that the latter determines *all* curves with endomorphism ring \mathcal{O} via exhaustive search, while we search for one and obtain the others via the action of $\text{Cl}(\mathcal{O})$ on the set $\text{Ell}_D(\mathbf{F}_p)$.

Step 1 can be implemented by picking j -invariants at random until one with the desired endomorphism ring is found. With $4p = u^2 - v^2D$, a necessary condition is that the curve E or its quadratic twist E' has $p + 1 - u$ points. In the case that D is fundamental and $v = 1$, this condition is also sufficient. To test if one of our curves E has the right cardinality, we pick a random point $P \in E(\mathbf{F}_p)$ and check if $(p + 1 - u)P = 0$ or $(p + 1 + u)P = 0$ holds. If neither of them does, E does not have endomorphism ring \mathcal{O} . If E survives this test, we select a few random points on both E and E' and compute the orders of these points *assuming* they divide $p + 1 \pm u$. If the curve E indeed has $p + 1 \pm u$ points, we quickly find points $P \in E(\mathbf{F}_p)$, $P' \in E'(\mathbf{F}_p)$ of maximal order, since we have $E(\mathbf{F}_p) \cong \mathbf{Z}/n_1\mathbf{Z} \times \mathbf{Z}/n_2\mathbf{Z}$ with $n_1 \mid n_2$ and a fraction $\varphi(n_2)/n_2$ of the points have maximal order. For P and P' of maximal order and $p > 457$, either the order of P or the order of P' is at least $4\sqrt{p}$, by [19, Theorem 3.1], due to J.-F. Mestre. As the Hasse interval has length $4\sqrt{p}$, this then *proves* that E has $p + 1 \pm u$ points.

Let $\Delta = \frac{D}{f^2}$ be the fundamental discriminant associated to D . For $f \neq 1$ or $v \neq 1$ (which happens necessarily for $D \equiv 1 \pmod{8}$), the curves with $p + 1 \pm u$ points admit any order $\mathcal{O}_{g^2\Delta}$ such that $g \mid fv$ as their endomorphism rings. In this case, one possible strategy is to use Kohel’s algorithm described in [13, Th. 24] to compute g , until a curve with $g = f$ is found. This variant is easiest to analyze and enough to prove Theorem 1.

In practice, one would rather keep a curve that satisfies $f \mid g$, since by the class number formula $g = vf$ with overwhelming probability. As v and thus $\frac{fv}{g}$ is small, it is then possible to use another algorithm due to Kohel and analyzed in detail by Fouquet–Morain [13,11] to quickly apply an isogeny of degree $\frac{fv}{g}$ leading to a curve with endomorphism ring \mathcal{O} .

Concerning Step 2, let $\text{Cl}(\mathcal{O}) = \bigoplus \langle \mathfrak{l}_i \rangle$ be a decomposition of the class group into a direct product of cyclic groups generated by invertible degree 1 prime ideals \mathfrak{l}_i of order h_i and norm ℓ_i not dividing pv . The $j^{\mathfrak{a}}$ may then be obtained successively by computing the Galois action of the \mathfrak{l}_i on j -invariants of curves with endomorphism ring \mathcal{O} over \mathbf{F}_p , otherwise said, by computing ℓ_i -isogenous curves: $h_1 - 1$ successive applications of \mathfrak{l}_1 yield $j^{\mathfrak{l}_1}, \dots, j^{\mathfrak{l}_1^{h_1-1}}$; to each of them, \mathfrak{l}_2 is applied $h_2 - 1$ times, and so forth.

To explicitly compute the action of $\mathfrak{l} = \mathfrak{l}_i$, we let $\Phi_{\ell}(X, Y) \in \mathbf{Z}[X]$ be the classical modular polynomial. It is a model for the modular curve $Y_0(\ell)$ parametrizing elliptic curves together with an ℓ -isogeny, and it satisfies $\Phi_{\ell}(j(z), j(\ell z)) = 0$ for the modular function $j(z)$. If $j_0 \in \mathbf{F}_p$ is the j -invariant of some curve with endomorphism ring \mathcal{O} , then all the roots in \mathbf{F}_p of $\Phi_{\ell}(X, j_0)$ are j -invariants of curves with endomorphism ring \mathcal{O} by [13, Prop. 23]. If \mathfrak{l} is unramified, there are two roots, $j_0^{\mathfrak{l}}$ and $j_0^{\mathfrak{l}^{-1}}$. For ramified \mathfrak{l} , we find only one root $j_0^{\mathfrak{l}} = j_0^{\mathfrak{l}^{-1}}$. So Step 2 is reduced to determining roots of univariate polynomials over \mathbf{F}_p .

3.2 Inert Primes

Algorithm 3

INPUT: an imaginary quadratic discriminant D and a prime p that is inert in \mathcal{O}

OUTPUT: $H_D \bmod p$

1. Compute the list of supersingular j -invariants over \mathbf{F}_{p^2} together with their endomorphism rings inside the quaternion algebra $\mathcal{A}_{p,\infty}$.
2. Compute an optimal embedding $f : \mathcal{O} \hookrightarrow \mathcal{A}_{p,\infty}$ and let R be a maximal order that contains $f(\mathcal{O})$.
3. Select a curve E/\mathbf{F}_{p^2} in the list with $\text{End}(E) \cong R$, and let j be its j -invariant.
4. Compute the Galois conjugates $j^{\mathfrak{a}}$ for $\mathfrak{a} \in \text{Cl}(\mathcal{O})$.
5. Return $H_D \bmod p = \prod_{\mathfrak{a} \in \text{Cl}(\mathcal{O})} (X - j^{\mathfrak{a}})$.

As the number of supersingular j -invariants grows roughly like $(p - 1)/12$, this algorithm is only feasible for *small* primes. For the explicit computation, we use an algorithm due to Cerviño [4] to compile our list. The list gives a bijection between the set of $\text{Gal}(\mathbf{F}_{p^2}/\mathbf{F}_p)$ -conjugacy classes of supersingular j -invariants and the set of maximal orders in $\mathcal{A}_{p,\infty}$.

In Step 2 we compute an element $y \in \mathcal{A}_{p,\infty}$ satisfying the same minimal polynomial as a generator τ of \mathcal{O} . For non-fundamental discriminants we need to ensure that the embedding is optimal, i.e., does not extend to an embedding of the maximal overorder of \mathcal{O} into $\mathcal{A}_{p,\infty}$. Using standard algorithms for quaternion algebras, Step 2 poses no practical problems. To compute the action of an ideal \mathfrak{a} in Step 4, we note that the right order R' of the left R -ideal $Rf(\mathfrak{a})$ is isomorphic to the endomorphism ring $\text{End}(E')$ of a curve E' with $j(E') = j^{\mathfrak{a}}$ by [22, Prop. 3.9]. The order R' is isomorphic to a unique order in the list, and we get a conjugacy class of supersingular j -invariants. Since roots of $H_D \bmod p$ which are not in \mathbf{F}_p come in conjugate pairs, this allows us to compute all the Galois conjugates $j^{\mathfrak{a}}$.

4 Computing the Canonical Lift of a Supersingular Curve

In this section we explain how to compute the Hilbert class polynomial H_D of a discriminant $D < -4$ using a p -adic lifting technique for an inert prime $p \equiv 1 \pmod{12}$. Our approach is based on the outline described in [7]. The condition $p \equiv 1 \pmod{12}$ ensures that the j -values $0, 1728 \in \mathbf{F}_p$ are not roots of $H_D \in \mathbf{F}_p[X]$. The case where one of these two values is a root of $H_D \in \mathbf{F}_p[X]$ is more technical due to the extra automorphisms of the curve, and will be explained in detail in the first author’s PhD thesis.

Under GRH, we can take p to be *small*. Indeed, our condition amounts to prescribing a Frobenius symbol in the degree 8 extension $\mathbf{Q}(\zeta_{12}, \sqrt{D})/\mathbf{Q}$, and by effective Chebotarev [14] we may take p to be of size $O((\log |D|)^2)$.

The first step of the algorithm is the same as for Algorithm 3 in Section 3: we compute a pair $(j(E_p), f_0) \in \text{Emb}_D(\mathbf{F}_{p^2})$. The main step of the algorithm is to compute to sufficient p -adic precision the canonical lift \tilde{E}_p of this pair, defined in Section 2 as the inverse under the bijection π of Theorem 2.

For an arbitrary element $\eta \in \text{Emb}_D(\mathbf{F}_{p^2})$, let

$$X_D(\eta) = \{(j(E), f) \mid j(E) \in \mathbf{C}_p, (j(E) \pmod p, f) = \eta\}$$

be a ‘disc’ of pairs lying over η . Here, \mathbf{C}_p is the completion of an algebraic closure of \mathbf{Q}_p . The disc $X_D(\eta)$ contains the points of $\text{Ell}_D(L)$ that reduce modulo p to the j -invariant corresponding to η .

These discs are similar to the discs used for the *split* case in [7,3]. The main difference is that now we need to keep track of the embedding as well. We can adapt the key idea of [7] to construct a p -adic analytic map from the set of discs to itself that has the CM-points as fixed points in the following way. Let \mathfrak{a} be an \mathcal{O} -ideal of norm N that is coprime to p . We define a map

$$\rho_{\mathfrak{a}} : \bigcup_{\eta} X_D(\eta) \rightarrow \bigcup_{\eta} X_D(\eta)$$

as follows. For $(j(E), f) \in X_D(\eta)$, the ideal $f(\mathfrak{a}) \subset \text{End}(E_p)$ defines a subgroup $E_p[f(\mathfrak{a})] \subset E_p[N]$ which lifts canonically to a subgroup $E[\mathfrak{a}] \subset E[N]$. We define $\rho_{\mathfrak{a}}((j(E), f)) = (j(E/E[\mathfrak{a}]), f^{\mathfrak{a}})$, where $f^{\mathfrak{a}}$ is as in Section 2. If the map f is clear, we also denote by $\rho_{\mathfrak{a}}$ the induced map on the j -invariants.

For principal ideals $\mathfrak{a} = (\alpha)$, the map $\rho_{\mathfrak{a}} = \rho_{\alpha}$ stabilizes every disc. Furthermore, as $\tilde{E}_p[(\alpha)]$ determines an endomorphism of \tilde{E}_p , the map ρ_{α} *fixes* the canonical lift $j(\tilde{E}_p)$. As $j(E_p)$ does not equal $0, 1728 \in \mathbf{F}_p$, the map ρ_{α} is *p -adic analytic* by [3, Theorem 4.2].

Writing $\alpha = a + b\tau$, the derivative of ρ_{α} in a CM-point $j(\tilde{E})$ equals $\alpha/\bar{\alpha} \in \mathbf{Z}_L$ by [3, Lemma 4.3]. For $p \nmid a, b$ this is a p -adic unit and we can use a modified version of Newton’s method to converge to $j(\tilde{E})$ starting from a random lift $(j_1, f_0) \in X_D(\eta)$ of the chosen point $\eta = (j(E_p), f_0) \in \text{Emb}_D(\mathbf{F}_{p^2})$. Indeed, the sequence

$$j_{k+1} = j_k - \frac{\rho_{\alpha}((j_k, f_0)) - j_k}{\alpha/\bar{\alpha} - 1} \tag{2}$$

converges quadratically to $j(\tilde{E})$. The run time of the resulting algorithm to compute $j(\tilde{E}) \in L$ up to the necessary precision depends heavily on the choice of α . We find a suitable α by sieving in the set $\{a + b\tau \mid a, b \in \mathbf{Z}, \gcd(a, b) = 1, a, b \neq 0 \pmod p\}$. We refer to the example in Section 6.3 for the explicit computation of the map ρ_α .

Once the canonical lift has been computed, the computation of the Galois conjugates is easier. To compute the Galois conjugate $j(\tilde{E}_p)^l$ of an ideal l of prime norm $\ell \neq p$, we first compute the value $j(E_p)^l \in \mathbf{F}_{p^2}$ as in Algorithm 3 in Section 3. We then compute all roots of the ℓ -th modular polynomial $\Phi_\ell(j(\tilde{E}_p), X) \in L[X]$ that reduce to $j(E_p)^l$. If there is only one such root, we are done: this is the Galois conjugate we are after. In general, if $m \geq 1$ is the p -adic precision required to distinguish the roots, we compute the value $\rho_l((j(\tilde{E}_p), f_0))$ to $m + 1$ p -adic digits precision to decide which root of the modular polynomial is the Galois conjugate. After computing all conjugates, we expand the product $\prod_{a \in \text{Cl}(\mathcal{O})} (X - j(\tilde{E}_p)^a) \in \mathbf{Z}_L[X]$ and recognize the coefficients as integers.

5 Complexity Analysis

This section is devoted to the run time analysis of Algorithm 1 and the proof of Theorem 1. To allow for an easier comparison with other methods to compute H_D , the analysis is carried out with respect to all relevant variables: the discriminant D , the class number $h(D)$, the logarithmic height n of the class polynomial and the largest prime generator $\ell(D)$ of the class group, before deriving a coarser bound depending only on D .

5.1 Some Number Theoretic Bounds

For the sake of brevity, we write llog for $\log \log$ and lllog for $\log \log \log$.

The bound given in Algorithm 1 on n , the bit size of the largest coefficient of the class polynomial, depends essentially on two quantities: the class number $h(D)$ of \mathcal{O} and the sum $\sum_{[A,B,C]} \frac{1}{A}$, taken over a system of primitive reduced quadratic forms representing the class group $\text{Cl}(\mathcal{O})$.

Lemma 1. *We have $h(D) = O(|D|^{1/2} \log |D|)$. If GRH holds true, we have $h(D) = O(|D|^{1/2} \text{llog} |D|)$.*

Proof. By the analytic class number formula, we have to bound the value of the Dirichlet L -series $L(s, \chi_D)$ associated to D at $s = 1$. The unconditional bound follows directly from [20], the conditional bound follows from [16]. \square

Lemma 2. *We have $\sum_{[A,B,C]} \frac{1}{A} = O((\log |D|)^2)$. If GRH holds true, we have $\sum_{[A,B,C]} \frac{1}{A} = O(\log |D| \text{llog} |D|)$.*

Proof. The bound $\sum_{[A,B,C]} \frac{1}{A} = O((\log |D|)^2)$ is proved in [18] with precise constants in [9]; the argument below will give a different proof of this fact.

By counting the solutions of $B^2 \equiv D \pmod{4A}$ for varying A and using the Chinese remainder theorem, we obtain

$$\sum_{[A,B,C]} \frac{1}{A} \leq \sum_{A \leq \sqrt{|D|}} \frac{\prod_{p|A} \left(1 + \frac{D}{p}\right)}{A}.$$

The Euler product expansion bounds this by $\prod_{p \leq \sqrt{|D|}} \left(1 + \frac{1}{p}\right) \left(1 + \frac{D}{p}\right)$. By Mertens theorem, this is at most $c \log |D| \prod_{p \leq \sqrt{|D|}} \frac{1}{1 - \frac{D}{p}}$ for some constant $c > 0$. This last product is essentially the value of the Dirichlet L -series $L(1, \chi_D)$ and the same remarks as in Lemma 1 apply. \square

Lemma 3. *If GRH holds true, the primes needed for Algorithm 1 are bounded by $O(h(D) \max(h(D)(\log |D|)^4, n))$.*

Proof. Let $k(D)$ be the required number of splitting primes. We have $k(D) \in O\left(\frac{n}{\log |D|}\right)$, since each prime has at least $\log_2 |D|$ bits.

Let $\pi_1(x, K_{\mathcal{O}}/\mathbf{Q})$ be the number of primes up to $x \in \mathbf{R}_{>0}$ that split completely in $K_{\mathcal{O}}/\mathbf{Q}$. By [14, Th. 1.1] there is an effectively computable constant $c \in \mathbf{R}_{>0}$, independent of D , such that

$$\left| \pi_1(x, K_{\mathcal{O}}/\mathbf{Q}) - \frac{\text{Li}(x)}{2h(D)} \right| \leq c \left(\frac{x^{1/2} \log(|D|^{h(D)} x^{2h(D)})}{2h(D)} + \log(|D|^{h(D)}) \right), \quad (3)$$

where we have used the bound $\text{disc}(K_{\mathcal{O}}/\mathbf{Q}) \leq |D|^{h(D)}$ proven in [3, Lemma 3.1]. It suffices to find an $x \in \mathbf{R}_{>0}$ for which $k(D) - \text{Li}(x)/(2h(D))$ is larger than the right hand side of (3). Using the estimate $\text{Li}(x) \sim x/\log x$, we see that the choice $x = O(\max(h(D)^2 \log^4 |D|, h(D)n))$ works. \square

5.2 Complexity of Algorithm 2

Let us fix some notation and briefly recall the complexities of the asymptotically fastest algorithms for basic arithmetic. Let $M(\log p) \in O(\log p \lceil \log p \rceil \log p)$ be the time for a multiplication in \mathbf{F}_p and $M_X(\ell, \log p) \in O(\ell \log \ell M(\log p))$ the time for multiplying two polynomials over \mathbf{F}_p of degree ℓ .

As the final complexity will be exponential in $\log p$, we need not worry about the detailed complexity of polynomial or subexponential steps. Writing $4p = u^2 - v^2 D$ takes polynomial time by the Cornacchia and Tonelli–Shanks algorithms [5, Sec 1.5]. By Lemma 3, we may assume that v is polynomial in $\log |D|$.

Concerning Step 2, we expect to check $O(p/h(D))$ curves until finding one with endomorphism ring \mathcal{O} . To test if a curve has the desired cardinality, we need to compute the orders of $O(\lceil \log p \rceil)$ points, and each order computation takes time $O((\log p)^2 M(\log p))$. Among the curves with the right cardinality, a fraction of $\frac{h(D)}{H(v^2 D)}$, where $H(v^2 D)$ is the Kronecker class number, has the

desired endomorphism ring. So we expect to apply Kohel’s algorithm with run time $O(p^{1/3+o(1)})$ an expected $\frac{H(v^2D)}{h(D)} \in O(v \log v)$ times. As $p^{1/3}$ is dominated by $p/h(D)$ of order about $p^{1/2}$, Step 2 takes time altogether

$$O\left(\frac{p}{h(D)}(\log p)^2 M(\log p) \log p\right). \tag{4}$$

Heuristically, we only check if some random points are annihilated by $p + 1 \pm u$ and do not compute their actual orders. The $(\log p)^2$ in (4) then becomes $\log p$.

In Step 3, the decomposition of the class group into a product of cyclic groups takes subexponential time. Furthermore, since all involved primes ℓ_i are of size $O((\log |D|)^2)$ under GRH, the time needed to compute the modular polynomials is negligible. Step 3 is thus dominated by $O(h(D))$ evaluations of reduced modular polynomials and by the computation of their roots.

Once $\Phi_\ell \bmod p$ is computed, it can be evaluated in time $O(\ell^2 M(\log p))$. Finding its roots is dominated by the computation of X^p modulo the specialized polynomial of degree $\ell + 1$, which takes time $O(\log p M_X(\ell, \log p))$. Letting $\ell(D)$ denote the largest prime needed to generate the class group, Step 3 takes time

$$O(h(D)\ell(D) M(\log p)(\ell(D) + \log |D| \log p)). \tag{5}$$

Under GRH, $\ell(D) \in O((\log |D|)^2)$, and heuristically, $\ell(D) \in O((\log |D|)^{1+\varepsilon})$.

By organizing the multiplications of polynomials in a tree of height $O(\log h)$, Step 4 takes $O(\log h(D) M_X(h(D), \log p))$, which is dominated by Step 3. We conclude that the total complexity of Algorithm 2 is dominated by Steps 2 and 3 and given by the sum of (4) and (5).

5.3 Proof of Theorem 1

We assume that $\mathcal{P} = \{p_1, p_2, \dots\}$ is chosen as the set of the smallest primes p that split into principal ideals of \mathcal{O} . Notice that $\log p, \log h(D) \in O(\log |D|)$, so that we may express all logarithmic quantities with respect to D .

The dominant part of the algorithm are the $O(n/\log |D|)$ invocations of Algorithm 2 in Step 2. Specializing (4) and (5), using the bound on the largest prime of Lemma 3 and assuming that $\ell(D) \in \Omega(\log |D| \log |D|)$, this takes time

$$O\left(n M(\log |D|) \left(h(D) \frac{\ell(D)^2}{\log |D|} + \log |D| \log |D| \max(h(D)(\log |D|)^4, n)\right)\right). \tag{6}$$

Finally, the fast Chinese remainder algorithm takes $O(M(\log N) \log N)$ by [12, Th. 10.25], so that Step 3 can be carried out in $O(h(D) M(n) \log |D|)$, which is also dominated by Step 2. Plugging the bounds of Lemmata 1 and 2 into (6) proves the rigorous part of Theorem 1.

For the heuristic result, we note that Lemma 3 overestimates the size of the primes, since it gives a very high bound already for the *first* split prime. Heuristically, one would rather expect that all primes are of size $O(nh)$. Combined with the heuristic improvements to (4) and (5), we find the run time

$$O\left(n M(\log |D|) \left(n + h(D) \frac{\ell(D)^2}{\log |D|}\right)\right). \quad \square$$

5.4 Comparison

The bounds under GRH of Lemmata 1 and 2 also yield a tighter analysis for other algorithms computing H_D . By [9, Th. 1], the run time of the complex analytic algorithm turns out to be $O(|D|(\log |D|)^3(\log \log |D|)^3)$, which is essentially the same as the heuristic bound of Theorem 1.

The run time of the p -adic algorithm becomes $O(|D|(\log |D|)^{6+o(1)})$. A heuristic run time analysis of this algorithm has not been undertaken, but it seems likely that again $O(|D|(\log |D|)^{3+o(1)})$ would be reached.

6 Examples and Practical Considerations

6.1 Inert Primes

For very small primes there is a unique supersingular j -invariant in characteristic p . For example, for $D \equiv 5 \pmod 8$, the prime $p = 2$ is inert in \mathcal{O}_D and we immediately have $H_D \pmod 2 = X^{h(D)}$.

More work needs to be done if there is more than one supersingular j -invariant in \mathbf{F}_{p^2} , as illustrated by computing $H_{-71} \pmod{53}$. The ideal $\mathfrak{a} = (2, 3 + \tau)$ generates the order 7 class group of $\mathcal{O} = \mathbf{Z}[\tau]$. The quaternion algebra $\mathcal{A}_{p,\infty}$ has a basis $\{1, i, j, k\}$ with $i^2 = -2, j^2 = -35, ij = k$, and the maximal order R with basis $\{1, i, 1/4(2 - i - k), -1/2(1 + i + j)\}$ is isomorphic to the endomorphism ring of the curve with j -invariant 50. We compute the embedding $f : \tau \mapsto y = 1/2 - 3/2i + 1/2j \in R$, where y satisfies $y^2 - y + 18 = 0$. Calculating the right orders of the left R -ideals $Rf(\mathfrak{a}^i)$ for $i = 1, \dots, 7$, we get a sequence of orders corresponding to the j -invariants $28 \pm 9\sqrt{2}, 46, 0, 46, 28 \pm 9\sqrt{2}, 50, 50$ and compute $H_{-71} \pmod{53} = X(X - 46)^2(X - 50)^2(X^2 + 50X + 39)$.

6.2 Totally Split Primes

For $D = -71$, the smallest totally split prime is $p = 107 = \frac{12^2+4\cdot 71}{4}$. Any curve over \mathbf{F}_p with endomorphism ring \mathcal{O} is isomorphic to a curve with $m = p + 1 \pm 12 = 96$ or 120 points. By trying randomly chosen j -invariants, we find that $E : Y^2 = X^3 + X + 35$ has 96 points. We either have $\text{End}(E) = \mathcal{O}_D$ or $\text{End}(E) = \mathcal{O}_{4D}$. In this simple case there is no need to apply Kohel’s algorithm. Indeed, $\text{End}(E)$ equals \mathcal{O}_D if and only if the complete 2-torsion is \mathbf{F}_p -rational. The curve E has only the point $P = (18, 0)$ as rational 2-torsion point, and therefore has endomorphism ring \mathcal{O}_{4D} . The 2-isogenous curve $E' = E/\langle P \rangle$ given by $Y^2 = X^3 + 58X + 59$ of j -invariant 19 has endomorphism ring \mathcal{O}_D .

The smallest odd prime generating the class group is $\ell = 3$. The third modular polynomial $\Phi_\ell(X, Y)$ has the two roots 46, 63 when evaluated in $X = j(E') = 19 \in \mathbf{F}_p$. Both values are roots of $H_D \pmod p$. We successively find the other Galois conjugates 64, 77, 30, 57 using the modular polynomial Φ_ℓ and expand

$$H_{-71} \pmod{107} = X^7 + 72X^6 + 93X^5 + 73X^4 + 46X^3 + 29X^2 + 30X + 19.$$

6.3 Inert Lifting

We illustrate the algorithm of Section 4 by computing H_D for $D = -56$.

The prime $p = 37$ is inert in $\mathcal{O} = \mathcal{O}_D$. The supersingular j -invariants in characteristic p are $8, 3 \pm 14\sqrt{-2}$. We fix a curve $E = E_p$ with j -invariant 8. We take the basis $\{1, i, j, k\}$ with $i^2 = -2, j^2 = j - 5, ij = k$ of the quaternion algebra $\mathcal{A}_{p,\infty}$. This basis is also a \mathbf{Z} -basis for a maximal order $R \subset \mathcal{A}_{p,\infty}$ that is isomorphic to the endomorphism ring $\text{End}(E_p)$.

Writing $\mathcal{O}_D = \mathbf{Z}[\tau]$, we compute an element $y = [0, 1, 1, -1] \in R$ satisfying $y^2 + 56 = 0$. This determines the embedding $f = f_0$ and we need to lift the pair (E, f) to its canonical lift. As element α for the ‘Newton map’ ρ_α , we use a generator of \mathfrak{a}^4 where $\mathfrak{a} = (3, 1 + \tau)$ is a prime lying over 3.

To find the kernel $E[f(\mathfrak{a})]$ we check which 3-torsion points $P \in E[3]$ are killed by $f(1 + \tau) \in \text{End}(E)$. We find $P = 18 \pm 9\sqrt{-2}$, and use Vélú’s formulas to find $E^\mathfrak{a} \cong E$ of j -invariant 8. As E and $E^\mathfrak{a}$ are isomorphic, it is easy to compute $f^\mathfrak{a}$. We compute a left-generator $x = [1, 1, 0, 0] \in R$ of the left R -ideal $Rf(\mathfrak{a})$ to find $f^\mathfrak{a}(\tau) = xy/x = [-1, 0, 1, 1] \in R$.

Next, we compute the \mathfrak{a} -action on the pair $(E^\mathfrak{a}, f^\mathfrak{a}) = (E, f^\mathfrak{a})$. We find that $P = 19 \pm 12\sqrt{a}$ is annihilated by $f^\mathfrak{a}(1 + \tau) \in \text{End}(E)$. The curve $E^{\mathfrak{a}^2}$ of j -invariant $3 - 14\sqrt{-2}$ is not isomorphic to E . We pick a 2-isogeny $\varphi_\mathfrak{b} : E^\mathfrak{a} \rightarrow E^{\mathfrak{a}^2}$ with kernel $\langle 19 + 23\sqrt{-2} \rangle$. The ideal \mathfrak{b} has basis $\{2, i + j, 2j, k\}$ and is left-isomorphic to $Rf^\mathfrak{a}(\mathfrak{a})$ via left-multiplication by $x' = [-1, 1/2, 1/2, -1/2] \in R$. We get $f^{\mathfrak{a}^2}(\tau) = x'y/x' = [0, 1, 1, -1] \in R_\mathfrak{b}$ and we use the map $g_\mathfrak{b}$ from Section 2 to view this as an embedding into $\text{End}(E^{\mathfrak{a}^2})$.

The action of \mathfrak{a}^3 and \mathfrak{a}^4 is computed in the same way. We find a cycle of 3-isogenies

$$(E, f) \rightarrow (E^\mathfrak{a} = E, f^\mathfrak{a}) \rightarrow (E^{\mathfrak{a}^2}, f^{\mathfrak{a}^2}) \rightarrow (E^{\mathfrak{a}^3}, f^{\mathfrak{a}^3}) \rightarrow (E^{\mathfrak{a}^4}, f^{\mathfrak{a}^4}) = (E, f)$$

where each element of the cycle corresponds uniquely to a root of H_D . We have now also computed $H_D \bmod p = (X - 8)^2(X^2 - 6X - 6)$.

As a lift of E we choose the curve defined by $Y^2 = X^3 + 210X + 420$ over the unramified extension L of degree 2 of \mathbf{Q}_p . We lift the cycle of isogenies over \mathbf{F}_{p^2} to L in 2 p -adic digits precision using Hensel’s lemma, and update according to the Newton formula (2) to find $j(\tilde{E}) = -66 + 148\sqrt{-2} + O(p^2)$. Next we work with 4 p -adic digits precision, lift the cycle of isogenies and update the j -invariant as before. In this example, it suffices to work with 16 p -adic digits precision to recover $H_D \in \mathbf{Z}[X]$.

Since we used a generator of an ideal generating the class group, we get the Galois conjugates of $j(\tilde{E})$ as a byproduct of our computation. In the end we expand the polynomial $H_{-56} = \prod_{\mathfrak{a} \in \text{Cl}(\mathcal{O})} (X - j(\tilde{E})^\mathfrak{a}) \in \mathbf{Z}[X]$ which has coefficients with up to 23 decimal digits.

6.4 Chinese Remainder Theorem

As remarked in Section 5.4, the heuristic run time of Theorem 1 is comparable to the expected run times of both the complex analytic and the p -adic approaches from [9] and [7,3]. To see if the CRT-approach is comparable *in practice* as well, we computed an example with a reasonably sized discriminant $D = -108708$, the first discriminant with class number 100.

The *a posteriori* height of H_D is 5874 bits, and we fix a target precision of $n = 5943$. The smallest totally split prime is 27241. If only such primes are used, the largest one is 956929 for a total of 324 primes. Note that these primes are indeed of size roughly $|D|$, in agreement with Lemma 3. We have partially implemented the search for a suitable curve: for each $4p = u^2 - v^2D$ we look for the first j -invariant such that for a random point P on an associated curve, $(p+1)P$ and uP have the same X -coordinate. This allows us to treat the curve and its quadratic twist simultaneously. The largest occurring value of v is 5. Altogether, 487237 curves need to be checked for the target cardinality.

On an Athlon-64 2.2 GHz computer, this step takes roughly 18.5 seconds. As comparison, the third authors' complex analytic implementation takes 0.3 seconds on the same machine. To speed up the multi-prime approach, we incorporated some inert primes. Out of the 168 primes less than 1000, there are 85 primes that are inert in \mathcal{O} . For many of them, the computation of $H_D \bmod p$ is trivial. Together, these primes contribute 707 bits and we only need 288 totally split primes, the largest one being 802597. The required 381073 curve cardinalities are tested in 14.2 seconds.

One needs to be careful when drawing conclusions from only few examples, but the difference between 14.2 and 0.3 seconds suggests that the implicit constants in the O -symbol are worse for the CRT-approach.

6.5 Class Invariants

For many applications, we are mostly interested in a generating polynomial for the ring class field $K_{\mathcal{O}}$. As the Hilbert class polynomial has very large coefficients, it is then better to use 'smaller functions' than the j -function to save a constant factor in the size of the polynomials. We refer to [17,21] for the theory of such *class invariants*.

There are theoretical obstructions to incorporating class invariants into Algorithm 1. Indeed, if a modular function f has the property that there are class invariants $f(\tau_1)$ and $f(\tau_2)$ with different minimal polynomials, we cannot use the CRT-approach. This phenomenon occurs for instance for the double eta quotients described in [10]. For the discriminant D in Section 6.4, we can use the double eta quotient of level $3 \cdot 109$ to improve the 0.3 seconds of the complex analytic approach. For CRT, we need to consider less favourable class invariants.

Acknowledgement

We thank Dan Bernstein, François Morain and Larry Washington for helpful discussions.

References

1. Agashe, A., Lauter, K., Venkatesan, R.: Constructing elliptic curves with a known number of points over a prime field. In: van der Poorten, A.J., Stein, A. (eds.) High Primes and Misdemeanours: Lectures in Honour of the 60th Birthday of H C Williams. Fields Inst. Commun., vol. 41, pp. 1–17 (2004)

2. Atkin, A.O.L., Morain, F.: Elliptic curves and primality proving. *Math. Comp.* 61(203), 29–68 (1993)
3. Bröker, R.: A p -adic algorithm to compute the Hilbert class polynomial. *Math. Comp.* (to appear)
4. Cerviño, J.M.: Supersingular elliptic curves and maximal quaternionic orders. In: *Math. Institut G-A-Univ. Göttingen*, pp. 53–60 (2004)
5. Cohen, H.: A Course in Computational Algebraic Number Theory. In: *Graduate Texts in Mathematics*, vol. 138, Springer, Heidelberg (1993)
6. Cohen, H., Frey, G., Avanzi, R., Doche, C., Lange, T., Nguyen, K., Vercauteren, F.: Handbook of Elliptic and Hyperelliptic Curve Cryptography. In: *Discrete Mathematics and its Applications*, Chapman & Hall/CRC (2005)
7. Couveignes, J.-M., Henocq, T.: Action of modular correspondences around CM points. In: Fieker, C., Kohel, D.R. (eds.) *ANTS 2002. LNCS*, vol. 2369, pp. 234–243. Springer, Heidelberg (2002)
8. Deuring, M.: Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abh. Math. Sem. Univ. Hamburg* 14, 197–272 (1941)
9. Enge, A.: The complexity of class polynomial computation via floating point approximations. HAL-INRIA 1040 = arXiv:cs/0601104, INRIA (2006), <http://hal.inria.fr/inria-00001040>
10. Enge, A., Schertz, R.: Constructing elliptic curves over finite fields using double eta-quotients. *J. Théor. Nombres Bordeaux* 16, 555–568 (2004)
11. Fouquet, M., Morain, F.: Isogeny volcanoes and the SEA algorithm. In: Fieker, C., Kohel, D.R. (eds.) *ANTS 2002. LNCS*, vol. 2369, pp. 276–291. Springer, Heidelberg (2002)
12. von zur Gathen, J., Gerhard, J.: *Modern Computer Algebra*. Cambridge University Press, Cambridge (1999)
13. Kohel, D.: *Endomorphism Rings of Elliptic Curves over Finite Fields*. PhD thesis, University of California at Berkeley (1996)
14. Lagarias, J.C., Odlyzko, A.M.: Effective versions of the Chebotarev density theorem. In: Fröhlich, A. (ed.) *Algebraic Number Fields (L-functions and Galois properties)*, pp. 409–464. Academic Press, London (1977)
15. Lang, S.: *Elliptic Functions*. In: *GTM 112*, 2nd edn., Springer, New York (1987)
16. Littlewood, J.E.: On the class-number of the corpus $P(\sqrt{-k})$. *Proc. London Math. Soc.* 27, 358–372 (1928)
17. Schertz, R.: Weber’s class invariants revisited. *J. Théor. Nombres Bordeaux* 14(1), 325–343 (2002)
18. Schoof, R.: The exponents of the groups of points on the reductions of an elliptic curve. In: van der Geer, G., Oort, F., Steenbrink, J. (eds.) *Arithmetic Algebraic Geometry*, pp. 325–335. Birkhäuser, Basel (1991)
19. Schoof, R.: Counting points on elliptic curves over finite fields. *J. Théor. Nombres Bordeaux* 7, 219–254 (1995)
20. Schur, I.: Einige Bemerkungen zu der vorstehenden Arbeit des Herrn G. Pólya: Über die Verteilung der quadratischen Reste und Nichtreste. *Nachr. Kön. Ges. Wiss. Göttingen, Math.-Phys. Kl.*, pp. 30–36 (1918)
21. Stevenhagen, P.: Hilbert’s 12th problem, complex multiplication and Shimura reciprocity. In: Miyake, K. (ed.) *Class Field Theory—its Centenary and Prospect*, pp. 161–176. Amer. Math. Soc. (2001)
22. Waterhouse, W.C.: Abelian varieties over finite fields. *Ann. Sci. École Norm. Sup.* (4) 2, 521–560 (1969)