# Accelerating the CM method 

Andrew V. Sutherland


#### Abstract

Given a prime $q$ and a negative discriminant $D$, the CM method constructs an elliptic curve $E / \mathbf{F}_{q}$ by obtaining a root of the Hilbert class polynomial $H_{D}(X)$ modulo $q$. We consider an approach based on a decomposition of the ring class field defined by $H_{D}$, which we adapt to a CRT setting. This yields two algorithms, each of which obtains a root of $H_{D} \bmod q$ without necessarily computing any of its coefficients. Heuristically, our approach uses asymptotically less time and space than the standard CM method for almost all $D$. Under the GRH, and reasonable assumptions about the size of $\log q$ relative to $|D|$, we achieve a space complexity of $O((m+n) \log q)$ bits, where $m n=h(D)$, which may be as small as $O\left(|D|^{1 / 4} \log q\right)$. The practical efficiency of the algorithms is demonstrated using $|D|>10^{16}$ and $q \approx 2^{256}$, and also $|D|>10^{15}$ and $q \approx 2^{33220}$. These examples are both an order of magnitude larger than the best previous results obtained with the CM method.


## 1. Introduction

The CM method is a widely used technique for constructing elliptic curves over finite fields. To illustrate, let us construct an elliptic curve $E / \mathbf{F}_{q}$ with exactly $N$ points. We shall assume that $q$ is prime, and require $t=q+1-N$ to be nonzero and satisfy $|t|<2 \sqrt{q}$. We may write $4 q=t^{2}-v^{2} D$, for some nonzero integer $v$ and negative discriminant $D$, and then proceed as follows.
(i) Compute the Hilbert class polynomial $H_{D} \in \mathbf{Z}[X]$.
(ii) Find a root $x$ of $H_{D}(X)$ modulo $q$.

The root $x$ is the $j$-invariant of an elliptic curve $E$ with $\# E\left(\mathbf{F}_{q}\right)=N$; an explicit equation for $E$ can be obtained via [37]. The endomorphism $\operatorname{ring} \operatorname{End}(E)$ is isomorphic to the imaginary quadratic order $\mathcal{O}$ with discriminant $D$, and we say that $E$ has complex multiplication (CM) by $\mathcal{O}$.

In principle, the CM method can construct any ordinary elliptic curve $E / \mathbf{F}_{q}$. In practice, it is feasible only when $|D|$ is fairly small. The main difficulty lies in step 1. The Hilbert class polynomial is notoriously large, as may be seen in Table 1.

The value $h(D)$ is the class number of $D$, which is the degree of $H_{D}$. The size listed is an upper bound on the total size of $H_{D}$ derived from known bounds on its largest coefficient [42, Lemma 8], and is generally accurate to within ten percent. These discriminants were

Table 1. Size of $H_{D}(X)$.

| $\|D\|$ | $h(D)$ | Size | $\|D\|$ | $h(D)$ | Size |
| :--- | ---: | :---: | :--- | ---: | ---: |
| $10^{5}+4$ | 152 | 152 KB | $10^{11}+4$ | 145981 | 323 GB |
| $10^{6}+104$ | 472 | 1.67 MB | $10^{12}+135$ | 465872 | 3.53 TB |
| $10^{7}+47$ | 1512 | 22.3 MB | $10^{13}+15$ | 1463328 | 38.5 TB |
| $10^{8}+20$ | 5056 | 239 MB | $10^{14}+4$ | 4658184 | 384 TB |
| $10^{9}+15$ | 15216 | 2.73 GB | $10^{15}+15$ | 14635920 | 4.45 PB |
| $10^{10}+47$ | 48720 | 31.4 GB | $10^{16}+135$ | 46275182 | 47.2 PB |

Received 9 May 2011; revised 22 April 2012.
2010 Mathematics subject classification 11 Y 16 (primary), 11G15, 11G20, 14H52 (secondary).
chosen so that the ratio $h(D) / \sqrt{|D|}$ is within ten percent of its asymptotic average ( $0.461559 \ldots$. . , so they represent typical examples.

There are at least three different ways to compute $H_{D}$ : the complex analytic method $[\mathbf{3}, \mathbf{2 1}$, 26], a $p$-adic approach $[\mathbf{1 2}, \mathbf{1 8}]$, or by computing $H_{D}$ modulo many small primes and applying the Chinese remainder theorem (CRT) [1, 7, 15]. Under suitable heuristic assumptions all three approaches can achieve quasi-linear running times: $O\left(|D| \log ^{c}|D|\right)$ for some constant $c$. However, the $O\left(|D| \log ^{1+\epsilon}|D|\right)$ space needed to compute $H_{D}$ makes it difficult to apply these algorithms when $|D|$ is large. As noted in [21], space is typically the limiting factor.

Here we build on the CRT approach of [42], which gives a probabilistic (Las Vegas) algorithm to compute $H_{D}$, with an expected running time of $O\left(|D| \log ^{5+\epsilon}|D|\right)$ under the generalized Riemann hypothesis (GRH), and a heuristic running time of $O\left(|D| \log ^{3+\epsilon}|D|\right)$. Most critically for the CM method, it directly computes $H_{D} \bmod q$ without computing $H_{D}$ over $\mathbf{Z}$. This yields a space complexity of $O\left(|D|^{1 / 2+\epsilon} \log q\right)$, allowing it to handle much larger values of $|D|$.

As a practical optimization, the CM method may use alternative class polynomials that are smaller than $H_{D}$ by a large constant factor. The algorithm in [42] has recently been adapted to compute such class polynomials [24]. To simplify our presentation we focus on the Hilbert class polynomial $H_{D}$, but our results apply to all the class polynomials considered in [24], a feature we exploit in $\S 6$.

When $h(D)$ is composite, a root of $H_{D}$ can be obtained via a decomposition of the field extension defined by $H_{D}(X)$, as described in $[\mathbf{2 3}, \mathbf{2 8}]$. The algorithm in $[\mathbf{2 3}]$ computes integer polynomials that describe this decomposition, which can be used in place of $H_{D}$. This allows a single root-finding operation to be replaced by two root-finding operations of smaller degree, speeding up step (ii) of the CM method. We adapt this technique to a CRT setting, where we find it also accelerates step (i), which is the asymptotically dominant step as a function of $|D|^{\dagger}$.

Provided $h(D)$ is sufficiently composite, we may choose a decomposition that significantly reduces the size of the coefficients in the defining polynomials. In order to do so, we derive an explicit height bound that can be efficiently computed for each of the possible decompositions available. By choosing the optimal decomposition we gain nearly a $\log |D|$ factor in the running time, on average, based on the heuristic analysis in $\S 5.4$. This claim is supported by empirical data, and we give practical examples that achieve more than a tenfold speedup.

We are also able to improve the space complexity of the CM method.

Proposition. Assume the $G R H$, and fix real constants $\delta \geqslant 0$ and $\epsilon>0$. Let $\mathcal{O}$ be an imaginary quadratic order with discriminant $D$ and class number $h=m n$, with $m \leqslant$ $O\left(|D|^{1 / 2-\delta}\right)$. Let $q$ be a prime of the form $4 q=t^{2}-v^{2} D$, and assume $\log q=O\left(|D|^{\delta} \log |D|\right)$. An elliptic curve $E / \mathbf{F}_{q}$ with $\operatorname{End}(E) \cong \mathcal{O}$ can be constructed in $O\left(|D| \log ^{6+\epsilon}|D|\right)$ expected time using $O((m+n) \log q)$ space.

This is achieved by interleaving the computation of the defining polynomials modulo many small primes $p$ with root-finding operations modulo $q$. If we additionally require $m=\Omega\left(|D|^{1 / 2-\gamma}\right)$, for some $\gamma>\delta$, this yields an $O\left(|D|^{1 / 4+\gamma} \log q\right)$ space bound, improving the $O\left(|D|^{1 / 2+\epsilon} \log q\right)$ result in [42].

The organization of the paper is as follows. We begin with some necessary preparation in $\S 2$, and then present two algorithms to obtain a root of the Hilbert (or other) class polynomial modulo a prime $q$ in $\S \S 3$ and 4 . The optimization of the height bound is addressed in $\S 5$. Finally, we present performance data in $\S 6$, including the construction of an elliptic curve over a 256 bit prime field with $|D|>10^{16}$, and an elliptic curve over a 10000 -digit prime field with $|D|>10^{15}$, both of which are new records for the CM method.

[^0]
## 2. Background

### 2.1. Hilbert class polynomials

We first recall some facts from the theory of complex multiplication, referring to $[\mathbf{1 9}, \mathbf{3 4}, \mathbf{4 0}]$ for proofs and further background. Let $\mathcal{O}$ be an imaginary quadratic order, identified by its discriminant $D$. The $j$-invariant of the lattice $\mathcal{O}$ is an algebraic integer whose minimal polynomial is the Hilbert class polynomial $H_{D}$. If $\mathfrak{a}$ is an invertible $\mathcal{O}$-ideal (including $\mathfrak{a}=\mathcal{O}$ ), then the torus $\mathbf{C} / \mathfrak{a}$ corresponds to an elliptic curve $E / \mathbf{C}$ with CM by $\mathcal{O}$, and every such curve arises in this fashion. Equivalent ideals yield isomorphic elliptic curves, and this gives a bijection between the ideal class group $\operatorname{cl}(\mathcal{O})$ and the set

$$
\operatorname{Ell}_{\mathcal{O}}(\mathbf{C})=\{j(E / \mathbf{C}): \operatorname{End}(E) \cong \mathcal{O}\}
$$

the $j$-invariants of the elliptic curves defined over $\mathbf{C}$ with CM by $\mathcal{O}$. We then have

$$
\begin{equation*}
H_{D}(X)=\prod_{j_{i} \in \operatorname{Ell}_{\mathcal{O}}(\mathbf{C})}\left(X-j_{i}\right) \tag{1}
\end{equation*}
$$

The splitting field of $H_{D}$ over $K=\mathbf{Q}(\sqrt{D})$ is the ring class field $K_{\mathcal{O}}$. It is an abelian extension whose Galois group is isomorphic to $\operatorname{cl}(\mathcal{O})$, via the Artin map.

This isomorphism can be made explicit via isogenies. Let $E / \mathbf{C}$ be an elliptic curve with CM by $\mathcal{O}$ and let $\mathfrak{a}$ be an invertible $\mathcal{O}$-ideal. After fixing an isomorphism $\operatorname{End}(E) \cong \mathcal{O}$, there is a uniquely determined separable isogeny whose kernel is the group of points annihilated by every endomorphism in $\mathfrak{a} \subset \mathcal{O} \cong \operatorname{End}(E)$. The image of this isogeny also has CM by $\mathcal{O}$, and this defines an action of the ideal group of $\mathcal{O}$ on the set $E l_{\mathcal{O}}(\mathbf{C})$. Principal ideals act trivially, and the induced action of the class group is regular. Thus the set $E l_{\mathcal{O}}(\mathbf{C})$ is a principal homogeneous space, a torsor, for the group $\operatorname{cl}(\mathcal{O})$. For a $j$-invariant $j_{i}$ in $\operatorname{Ell}_{\mathcal{O}}(\mathbf{C})$ and an ideal class $[\mathfrak{a}]$ in $\operatorname{cl}(\mathcal{O})$, we write $[\mathfrak{a}] j_{i}$ to denote the image of $j_{i}$ under the action of $[\mathfrak{a}]$.

If $p$ is a prime that splits completely in $K_{\mathcal{O}}$, equivalently (for $p>3$ ), a prime that satisfies the norm equation

$$
\begin{equation*}
4 p=t^{2}-v^{2} D \tag{2}
\end{equation*}
$$

for some nonzero integers $t$ and $v$, then $H_{D}$ splits completely in $\mathbf{F}_{p}[X]$ and its roots form the set $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$. Conversely, every ordinary (not supersingular) elliptic curve $E / \mathbf{F}_{p}$ has CM by some imaginary quadratic order $\mathcal{O}$ in which the Frobenius endomorphism corresponds to an element of norm $p$ and trace $t$.

### 2.2. Computing $H_{D}$ with the CRT method

The above theory suggests the following algorithm to compute $H_{D}$ modulo a prime $p$ that splits completely in $K_{\mathcal{O}}$.
(i) Find an elliptic curve $E$ with $j$-invariant $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$.
(ii) Enumerate $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ as $\left\{[\mathfrak{a}] j_{1}:[\mathfrak{a}] \in \operatorname{cl}(\mathcal{O})\right\}=\left\{j_{1}, \ldots, j_{h}\right\}$.
(iii) Compute $H_{D}(X) \bmod p$ as $\left(X-j_{1}\right)\left(X-j_{2}\right) \ldots\left(X-j_{h}\right)$.

Step (i) essentially involves testing curves at random, so this algorithm is feasible only when $p$ is fairly small, but see $[42, \S 3]$ for various ways in which this step may be optimized. Step (ii) is addressed in $\S 2.6$. Step (iii) is just polynomial arithmetic, but this is usually the most expensive step $[42, \S 7.1]$.

Under the GRH we can we always work with primes $p$ that are roughly the same size as $|D|$, no matter how big $q$ is. We select a sufficiently large set of such small primes that split completely in $K_{\mathcal{O}}$, and compute $H_{D} \bmod p$ for each of them. Provided the product of these primes is larger than $2 B$, where $B$ bounds the coefficients of $H_{D}$, the polynomial $H_{D}$
is uniquely determined by the Chinese Remainder Theorem (CRT). Explicit values for $B$ are given in $[\mathbf{2 2}]$ and [42, Lemma 8].

However, to compute $H_{D} \bmod q$ in $O\left(|D|^{1 / 2+\epsilon} \log q\right)$ space, one cannot simply compute $H_{D}$ over $\mathbf{Z}$ and then reduce it modulo $q$, since writing down $H_{D}$ would already take too much space. Instead, one may use the explicit CRT $(\bmod q)[9,10]$, applying it in an online fashion to accumulate results modulo $q$ that are updated after each computation of $H_{D} \bmod p$, as described in $[42, \S 6]$ and summarized in §2.3. Here we also use the explicit CRT, but to obtain a root of $H_{D} \bmod q$ with a space complexity that may be as small as $O\left(|D|^{1 / 4+\epsilon} \log q\right)$, we must even avoid writing down $H_{D} \bmod q$. This requires some new techniques that are described in §4.3.2.

### 2.3. Explicit CRT

Let $p_{1}, \ldots, p_{n}$ be primes with product $M$, let $M_{i}=M / p_{i}$, and let $a_{i} M_{i} \equiv 1 \bmod p_{i}$. If $c \in \mathbf{Z}$ satisfies $c \equiv c_{i} \bmod p_{i}$, then $c \equiv \sum_{i} c_{i} a_{i} M_{i} \bmod M$. If $M>2|c|$, this congruence uniquely determines $c$. This is the usual CRT method.

Now suppose $M>4|c|$ and let $q$ be a prime (or any integer, in fact). Then we may apply the $\operatorname{explicit} C R T \bmod q[\mathbf{1 0}$, Theorem 3.1] to compute

$$
\begin{equation*}
c \equiv\left(\sum_{i} c_{i} a_{i} M_{i}-r M\right) \bmod q, \tag{3}
\end{equation*}
$$

where $r$ is the closest integer to $\sum_{i} c_{i} a_{i} / p_{i}$; when computing $r$, it suffices to approximate each $c_{i} a_{i} / p_{i}$ to within $1 /(4 n)$, by $[\mathbf{1 0}$, Theorem 2.2$]$.

When applying the explicit CRT to compute many values $c \bmod q$, say, the coefficients of a polynomial, one first computes the $a_{i}\left(\bmod p_{i}\right)$ and $M_{i}(\bmod q)$ using a product tree as described in $[42, \S 6.1]$; this is CRT precomputation step, and it does not depend on the coefficients $c$. Then, as the reduced coefficient values $c_{i}=c \bmod p_{i}$ are computed for a particular $p_{i}$, the sum $\sum c_{i} a_{i} M_{i} \bmod q$ and an approximation to $\sum c_{i} a_{i} / p_{i}$ are updated for each coefficient, after which the $c_{i}$ may be discarded; this is the CRT update step. Finally, when these sums have been updated for every prime $p_{i}$, one applies (3) to obtain $c \bmod q$ for each coefficient, using the sums $\sum c_{i} a_{i} M_{i}$ and the approximations $r \approx \sum c_{i} a_{i} / p_{i}$; this is the CRT postcomputation step. For further details, including explicit algorithms for each step, see [42, §6.2]

### 2.4. Assuming the GRH

The Chebotarev density theorem guarantees that a set of primes $S$ sufficient to compute $H_{D}$ with the CRT method exists. We even have effective bounds on their size [33], but these are too large for our purpose. To obtain better bounds, we assume the GRH (for the Dedekind zeta function of $K_{\mathcal{O}}$ ), which implies that the primes in $S$ are no more than a polylogarithmic factor larger than $|D|$, see [42, Lemma 3]. Having made this assumption, we are then in a position to apply other bounds that depend on some instance of the extended or generalized Riemann hypothesis, all of which we take to be included when we assume the GRH without qualification. In particular, we make frequent use of the bound $h(D)=O\left(|D|^{1 / 2} \log \log |D|\right)$, proven in [35], as well as the bound

$$
\begin{equation*}
\sum_{[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})} \frac{1}{N(\mathfrak{a})}=O(\log |D| \log \log |D|), \tag{4}
\end{equation*}
$$

where $\mathfrak{a}$ is the invertible ideal of least norm in [ $\mathfrak{a}$ ], from [7, Lemma 2]. The sum in (4) may also be written as $\sum 1 / A_{i}$, where $\left(A_{i}, B_{i}, C_{i}\right)$ ranges over the primitive reduced binary quadratic forms $A_{i} x^{2}+B_{i} x y+C_{i} y^{2}$ of discriminant $D=B_{i}^{2}-4 A_{i} C_{i}$.

### 2.5. Decomposing the class equation

We now describe an explicit method for 'decomposing' a polynomial via a decomposition of the field extension it defines. This is based on material in [23] and [28] that we adapt to our purpose here.

Let $K$ be a number field, and let $P \in \mathbf{Z}[X]$ be a monic polynomial, irreducible over $K$, with splitting field $M$. We have in mind $K=\mathbf{Q}(\sqrt{D}), P=H_{D}$, and $M=K_{\mathcal{O}}$. We shall assume the action of $\operatorname{Gal}(M / K)$ is regular, equivalently, that $[M: K]=\operatorname{deg} P$, which holds in the case of interest.

Given a normal tower of fields $K \subset L \subset M$, we may decompose the extension $M / K$ into extensions $M / L$ and $L / K$ via polynomials $U \in \mathbf{Z}(Y)[X]$ and $V \in \mathbf{Z}[X]$, where $P(x)=0$ if and only if $U(x, y)=0$ and $V(y)=0$ for some $y \in L$. The polynomial $V(Y)$ defines the extension $L / K$, and for any root $y$ of $V$, the polynomial $U(X, y)$ defines the extension $M / L$. In our application $L$ will be identified as the fixed field of a given normal subgroup $G$ of $\operatorname{Gal}(M / K)$.

Let us fix a root $x$ of $P$, and let $\gamma x$ denote the conjugate of $x$ under the action of $\gamma \in \operatorname{Gal}(M / K)$. Let $\beta_{1}, \ldots, \beta_{n}$ be the elements of $G$, and let $\alpha_{1} G, \ldots, \alpha_{m} G$ be the cosets of $G$ in $\operatorname{Gal}(M / K)$. We may factor $P$ in $L[X]$ as $P=P_{1}, \ldots, P_{m}$, where

$$
\begin{equation*}
P_{i}(X)=\prod_{k=1}^{n}\left(X-\alpha_{i} \beta_{k} x\right)=\sum_{k=0}^{n} \theta_{i k} X^{k} . \tag{5}
\end{equation*}
$$

Now let us pick a symmetric function $s$ in $\mathbf{Z}\left[T_{1}, \ldots, T_{n}\right]$ for which each of the $m$ values $y_{i}=s\left(\alpha_{i} \beta_{1} x, \ldots, \alpha_{i} \beta_{n} x\right)$ are distinct, equivalently, for which $L=K\left(y_{i}\right) ;$ Lemma 1 below shows that this is easily achieved. We then define the polynomial

$$
\begin{equation*}
V(Y)=\prod_{i=1}^{m}\left(Y-y_{i}\right) . \tag{6}
\end{equation*}
$$

The coefficients of $V$ lie in $\mathbf{Z}$, since each may be expressed as a symmetric integer polynomial in the roots of the monic polynomial $P \in \mathbf{Z}[X]$. For $0 \leqslant k \leqslant n$ let

$$
\begin{equation*}
W_{k}(Y)=\sum_{i=1}^{m} \theta_{i k} \frac{V(Y)}{\left(Y-y_{i}\right)} . \tag{7}
\end{equation*}
$$

As with $V$, we have $W_{k} \in \mathbf{Z}[Y]$. Note that $W_{k}(Y)$ is the unique polynomial of degree less than $m$ for which $W_{k}\left(y_{i}\right)=\theta_{i k} V^{\prime}\left(y_{i}\right)$, by the Lagrange interpolation formula. This definition of the $W_{k}$ is referred to as the Hecke representation in [23].

Finally, let

$$
\begin{equation*}
U(X, Y)=\frac{1}{V^{\prime}(Y)} \sum_{k=0}^{n} W_{k}(Y) X^{k} \tag{8}
\end{equation*}
$$

For each root $y_{i}$ of $V(Y)$ we then have $U\left(X, y_{i}\right)=P_{i}(X)$, with $V^{\prime}\left(y_{i}\right) \neq 0$, since $V$ has distinct roots. Each root of $U\left(X, y_{i}\right)$ is a root of $P(X)$, and every root of $P(X)$ may be obtained in this way. Notice that this construction does not require us to know the coefficients of $P$, but we must be able enumerate the $G$-orbits of its roots.

As noted in [28], for any given $K, P$, and $G$, there are only finitely many linear combinations of the elementary symmetric functions, up to scalar factor, that do not yield a symmetric function $s$ suitable for the construction above. More precisely, we have the following lemma.

Lemma 1. Let $M=K(x)$ be a finite Galois extension of a number field $K$, let $G$ be a normal subgroup of $\operatorname{Gal}(M / K)$ with order $n$ and fixed field $L$, and let $m=[L: K]$. Let $\boldsymbol{x}=\left(x_{1}, \ldots, x_{n}\right)$ denote the $G$-orbit of $x$, and let $e_{1}, \ldots, e_{n}$ denote the elementary symmetric functions on $n$ variables. There are at most $m^{n-1}(m-1)^{n-1}$ linear combinations of the form $s=e_{1}+c_{2} e_{2}+\ldots+c_{n} e_{n}$, with $c_{2}, \ldots, c_{n} \in K$, for which $L \neq K(s(\boldsymbol{x}))$.

Proof. Let $z_{k}=e_{k}(\boldsymbol{x})$ and $L_{k}=K\left(z_{1}, \ldots, z_{k}\right)$. We have $L_{1} \subset \ldots \subset L_{n}=L$, where the last equality follows from the fact that $x$ is a root of the monic polynomial $X^{n}+\sum_{k=1}^{n}(-1)^{k} e_{k} X^{n-k}$ in $L_{n}[X]$, since $[L(x): L]=n$. From the proof of the primitive element theorem in $[44, \S 6.10]$, we know that $K\left(z_{1}, z_{2}\right)=K\left(z_{1}+c_{2} z_{2}\right)$ for all $c_{2} \in K$ not of the form $\left(u_{i}-u_{1}\right) /\left(v_{j}-v_{1}\right)$, with $j>1$, where $z_{1}=u_{1}, u_{2}, \ldots, u_{r}$ are conjugates in $L / K$, and $z_{2}=v_{1}, v_{2}, \ldots, v_{s}$ are conjugates in $L / K$. Since $r, s \leqslant m$, there can be at most $m(m-1)$ such $c_{2}$. The same argument shows that $K\left(z_{1}+c_{2} z_{2}+\ldots+c_{k-1} z_{k-1}, z_{k}\right)=K\left(z_{1}+c_{2} z_{2}+\ldots+c_{k} z_{k}\right)$ for all but at most $m(m-1)$ values of $c_{k} \in K$, and the lemma follows.

In the construction above, if $L=K(s(\boldsymbol{x}))$ holds for any $G$-orbit $\boldsymbol{x}$, then it necessarily holds for every $G$-orbit contained in the same $\operatorname{Gal}(M / K)$-orbit, and in this case the $s$-images $y_{i}$ of these $G$-orbits are distinct, since they are $\operatorname{Gal}(L / K)$-conjugates. To find such an $s$ explicitly, we pick random integer coefficients $c_{2}, \ldots, c_{n}$ uniformly distributed over $\left[0,2 m^{2}-1\right]$, and let $s=e_{1}+c_{2} e_{2}+\ldots+c_{k} e_{k}$. By Lemma 1, we then have $L=K(s(\boldsymbol{x}))$ with probability at least $1-2^{1-n}$. In the event that $L \neq K(s(\boldsymbol{x}))$, we simply pick a new set of random coefficients and repeat until we succeed, yielding a Las Vegas algorithm. In the trivial case, $n=1$ and we succeed on the first try with $s=e_{1}$; for $n>1$ we expect to succeed with at most $1 /\left(1-2^{1-n}\right) \leqslant 2$ attempts, on average.

### 2.6. Orbit enumeration

Recall that $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ is a torsor for $\operatorname{cl}(\mathcal{O})$. Each subgroup $G$ of $\operatorname{cl}(\mathcal{O})$ partitions $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ into $G$-orbits that correspond to cosets of $G$. Given $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$, the set $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ may be enumerated via [42, Algorithm 1.3], but to correctly identify its $G$-orbits we require some further refinements.

As in $[42, \S 5]$, we represent $\operatorname{cl}(\mathcal{O})$ using a polycyclic presentation defined by a sequence of ideals $\mathfrak{l}_{1}, \ldots, \mathfrak{l}_{k}$ of prime norm whose classes generate $\operatorname{cl}(\mathcal{O})$. The relative order of $\left[\mathfrak{l}_{i}\right]$ is the least positive integer $r_{i}$ for which $\left[r_{i}^{r_{i}}\right] \in\left\langle\left[\mathfrak{l}_{1}\right], \ldots,\left[\mathfrak{l}_{i-1}\right]\right\rangle$. Every element $[\mathfrak{a}]$ of $\operatorname{cl}(\mathcal{O})$ may be uniquely written in the form

$$
[\mathfrak{a}]=\left[\mathfrak{l}_{1}^{e_{1}}\right] \ldots\left[\mathfrak{l}_{k}^{e_{k}}\right]
$$

with $0 \leqslant e_{i}<r_{i}$. Having fixed a polycyclic presentation, we may enumerate $\operatorname{cl}(\mathcal{O})$ by enumerating the corresponding exponent vectors $\left(e_{1}, \ldots, e_{k}\right)$. Moreover, for any subgroup $G$ of $\operatorname{cl}(\mathcal{O})$, we can easily identify the exponent vectors corresponding to each coset of $G$. This allows us to distinguish the $G$-orbits of $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$, provided that we can unambiguously compute the actions of $\left[\mathfrak{l}_{1}\right], \ldots,\left[\mathfrak{l}_{k}\right]$.

Let $j_{1}=j\left(E_{1}\right)$ be an element of $\in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ and let $\mathfrak{l}$ be an invertible $\mathcal{O}$-ideal of prime norm $\ell \neq p$. Then $j_{2}=[l] j_{1}$ is the $j$-invariant of an elliptic curve $E_{2}$ that is $\ell$-isogenous to $E_{1}$. We may obtain $j_{2} \in \mathbf{F}_{p}$ as a root of $\phi(X)=\Phi_{\ell}\left(j_{1}, X\right)$, where $\Phi_{N} \in \mathbf{Z}[X, Y]$ is the classical modular polynomial $[46, \S 69]$ that parameterizes cyclic isogenies of degree $N$. When [l] has order 2, there is just one distinct root of $\phi(X)$ that lies in $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$, but in general there will be two: $[l] j_{1}$ and $[\bar{l}] j_{1}$. Typically these are the only roots of $\phi(X)$ in $\mathbf{F}_{p}$, and when they are not, they may be distinguished as the roots that lie on the surface of an $\ell$-volcano, see [42, §4] or [32]. The only ambiguity lies in distinguishing the actions of $[\mathfrak{l}]$ and $[\bar{l}]$, which correspond to the two directions we may walk along the cycle of $\ell$-isogenies on the surface.

To enumerate $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ it is not necessary to distinguish these directions, as shown in [42, Proposition 5]. However, it is necessary, in general, if we wish to identify the $G$-orbits of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ in this enumeration. To see why, let $\operatorname{cl}(\mathcal{O})$ be a cyclic group of order 6 generated by $[\mathfrak{a}]$, and suppose our polycyclic presentation has $\left[\mathfrak{l}_{1}\right]=\left[\mathfrak{a}^{2}\right]$ with $r_{1}=3$, and $\left[\mathfrak{l}_{2}\right]=[\mathfrak{a}]$ with $r_{2}=2$. Below are four of the eight possible enumerations of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ that might be produced
by [42, Algorithm 1.3] in this scenario.





Each node corresponds to a $j$-invariant $\left[\mathfrak{a}^{e}\right] j_{1}$ and is labeled by the exponent $e$. The arrows indicate the action of the ideal that appears in the label, and bold arrows indicate where an arbitrary choice was made. The two lightly shaded nodes in each configuration correspond to the expected positions of the elements of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ corresponding to the subgroup $G=\left\langle\mathfrak{a}^{3}\right\rangle$ of order 2 . The two configurations on the right yield a correct enumeration of each $G$-orbit, but the two on the left do not.

REMARK 1. Of course we could have used a polycyclic presentation with $\left[\mathfrak{l}_{1}\right]=[\mathfrak{a}]$ and $r_{1}=6$ in this example, but unless $[\mathfrak{a}]$ happens to contain the invertible $\mathcal{O}$-ideal of least prime norm, this is suboptimal, since the cost of computing the action of $[\mathfrak{l}]$ increases quadratically with the norm of $\mathfrak{l}$. The optimal polycyclic presentation rarely has one generator, even when $\operatorname{cl}(\mathcal{O})$ is cyclic.

We now consider ways to enumerate $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ that allow us to identify its $G$-orbits. First, the GCD technique of $[\mathbf{2 4}, \S 2.3]$ may be used to consistently choose $\mathfrak{l}_{i}$ or $\overline{\mathfrak{l}}_{i}$ each time a new path of $\ell_{i}$-isogenies is begun (this actually speeds up the enumeration, so it should be done in any case). We are then left with just $k$ choices, one for each $\mathfrak{l}_{i}$. To consistently orient these choices we may either: (a) use auxiliary relations $\left[\mathfrak{a}_{i}\right]=\left[\mathfrak{l}_{1} \ldots \mathfrak{l}_{i}\right]$, where $\mathfrak{a}_{i}$ has (small) prime norm different from $\ell_{1}, \ldots, \ell_{i}$, as described in $[\mathbf{2 4}, \S 4.3]$; or (b) compute the action of Frobenius on the kernel of the two isogenies, as discussed in $[\mathbf{1 3}, \S 4]$ and the references therein. Option (a) is fast and easy to implement, but for the best space complexity we should use (b), since it does not require us to store the entire enumeration. A minor drawback to option (b) is that it requires $\ell_{i} \nmid v$, where $v$ is as in the norm equation (2).

However, in many cases neither (a) nor (b) is necessary. If $G$ is of the form

$$
\begin{equation*}
G=\left\langle\left[\mathfrak{l}_{1}\right], \ldots,\left[\mathfrak{l}_{d-1}\right],\left[\mathfrak{l}_{d}^{e}\right]\right\rangle \tag{9}
\end{equation*}
$$

for some $d \leqslant k$ and $e \mid r_{d}$, then it follows from the proof of [42, Proposition 5] that any enumeration of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ output by [42, Algorithm 1.3] also gives a correct enumeration of the $G$-orbits of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$. In the context of the CM method, we are free to choose $G$, and we can always choose one that satisfies (9). This may limit our choices for $G$, but in practice this restriction is usually not a burden.

## 3. A first algorithm

Our first algorithm is a direct implementation of the theory presented in § 2, which can significantly accelerate the CM method in the typical case where the class number $h(D)$ is composite. At this stage we shall not be concerned with improving space complexity; this will be the focus of our second algorithm. Both algorithms are probabilistic (of Las Vegas type).

As above, $\mathcal{O}$ is an imaginary quadratic order with discriminant $D$, and we shall assume $h=h(D)>1$. Let $\mathcal{P}_{D}$ be the set of primes that split completely in $K_{\mathcal{O}}$, equivalently, primes that satisfy (2). Given $D$ and a prime $q \in \mathcal{P}_{D}$ we wish to obtain a root of the Hilbert class polynomial $H_{D}$ over $\mathbf{F}_{q}$. Such a root is the $j$-invariant of an elliptic curve $E / \mathbf{F}_{q}$ with CM by $\mathcal{O}$, as required by the CM method.

The first step is to choose a subgroup $G$ of $\operatorname{cl}(O)$, which determines $n=|G|$ and $m=h / n$. The quantities $\theta_{i k}$ and $y_{i}, V(Y), W_{k}(Y)$, and $U(X, Y)$, are then defined as in $\S 2.5$, with $1 \leqslant i \leqslant m$ and $0 \leqslant k<n$ (we do not need $k=n$ ).

The $y_{i}$ (and therefore $V(Y)$ ) depend on the choice of a function $s \in \mathbf{Z}\left[X_{1}, \ldots, X_{n}\right]$ that is a randomly chosen linear combination of the elementary symmetric functions $e_{1}, \ldots, e_{n}$, as described at the end of $\S 5$. In the unlikely event that we pick a bad $s$, we may find that $V$ is a perfect power. In this case the algorithm simply repeats the computation with a new choice of $s$.

We also require a bound $B$ on the coefficients of $V$ and $W_{k}$. The computation of $B$ is addressed in $\S 5$. We now give the algorithm.

Algorithm 1. Given $D$ and $q \in \mathcal{P}_{D}$, find a root $x$ of $H_{D}$ in $\mathbf{F}_{q}$ as follows:

1. Select a subgroup $G$ of $\operatorname{cl}(O)$ with $2|G|^{2} \leqslant q$ and let $n=|G|$.
2. Generate random integers $c_{2}, \ldots, c_{n}$ uniformly distributed over $\left[0,2 m^{2}-1\right]$ and set $s=e_{1}+c_{2} e_{2}+\ldots+c_{n} e_{n}$.
3. Compute a bound $B$ as described in $\S 5$.
4. As in steps $1-3$ of $[42$, Alg. 2$]$, use $B$ to select $S \subset \mathcal{P}_{D}$, compute a polycyclic presentation $\Gamma$ for $\operatorname{cl}(O)$, and perform CRT precomputation (see $\S 2.3$ ).
5. For each $p \in S$ :
(a) Find $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ as in [42, Alg. 1].
(b) Enumerate the $G$-orbits $G_{i}$ of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ using $j_{1}$ and $\Gamma$.
(c) Compute the $\theta_{i k}$ and the $y_{i} \bmod p($ using $s)$, as described in $\S 2.5$
(d) Compute $V$ and the $W_{k} \bmod p$.
(e) Update CRT data for the coefficients of $V$ and the $W_{k}$.
6. Perform CRT postcomputation to obtain $V$ and the $W_{k} \bmod q$.
7. Working in $\mathbf{F}_{q}$, find a root $y$ of $V$ that is not a root of $V^{\prime}$. If no such root exists then return to step 2 .
8. Compute $U_{y}(X)=U(X, y) \bmod q$ and output a root $x$ of $U_{y}$ in $\mathbf{F}_{q}$.

It follows from $\S 2.5$ that the output value $x$ is a root of $H_{D}$ in $\mathbf{F}_{q}$, thus the algorithm is correct. Lemma 1 guarantees that it always terminates, and that its expected running time is no more than twice the expected time when the first choice of $s$ works (this factor approaches 1 as $n$ increases).

REmark 2. For practical implementation one may prefer to fix $s=e_{1}$ (or $s=-e_{1}$ ) and instead change the choice of $G$ if $s=e_{1}$ does not work. Using $s=e_{1}$ simplifies the implementation and can reduce the bound $B$ significantly. Empirically, for large values of $|D|$ and $q$, using $s=e_{1}$ is very likely to work with every choice of $G$. Alternatively, one may start with $s=e_{1}$ and then switch to a random $s$ if necessary. In all our examples, including all the computations in $\S 6$, and in the heuristic analysis of $\S 5.4$, we use $s= \pm e_{1}$, but our mathematical results (Propositions $1-3$ ) all assume a random $s$, as specified in Algorithm 1.

We note three immediate generalizations of Algorithm 1. First, other class invariants may be treated with suitable modifications to step 2 , see $[\mathbf{2 4}]$. Second, it is not necessary for $q$ to be prime; $q$ may be a prime power $q_{0}^{e}$ satisfying $4 q=t^{2}-v^{2} D$ with $t \not \equiv 0 \bmod q_{0}$. One then computes $V$ and the $W_{k} \bmod q_{0}$, and performs the root-finding operations in $\mathbf{F}_{q}$. Third, we can compute $V$ and the $W_{k}$ over $\mathbf{Z}$ using the standard CRT: replace $q$ by the product of the primes in $S$ and lift the results of step 4 to $\mathbf{Z}$. Steps 7 and 8 may then later be applied to any $q$ that splits completely in the ring class field.

### 3.1. Example

Let us find a root of $H_{D} \bmod q$ using Algorithm 1, with $D=-971$ and $q=1029167$. The class group is cyclic of order 15 , and the optimal polycyclic presentation ${ }^{\dagger}$ has norms $\ell_{1}=3$

[^1]and $\ell_{2}=5$ and relative orders $r_{1}=5$ and $r_{2}=3$. We choose $G$ to be the subgroup of order $n=5$, which is conveniently of the form (9), so we need not distinguish directions when enumerating $\mathrm{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$. For convenience we set $s=-e_{1}$, so that $\left.y_{i}=\theta_{i, n-1}\right)$.

Computing the bound $B$ as in $\S 5$, we have $b=\log _{2} B \approx 340$ bits, and select a set of primes in $\mathcal{P}_{D}$ whose product exceeds $4 B$. As described in [42, § 3], we choose primes that optimize the search for $j_{1} \in \operatorname{Ell} l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$, obtaining a set of 27 primes:

$$
S=\{263,353,1871, \ldots, 38677,43237,62873\}
$$

We then precompute parameters for the explicit CRT $(\bmod q)$, using [42, Algorithm 2.3].
For $p=263$ we find $j_{1}=252$, and enumerate $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ from $j_{1}$ as,


The horizontal arrows denote 3 -isogenies and the vertical arrows are 5 -isogenies. The three $G$-orbits of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ are $\{252,38,151,121,258\}, \quad\{70,112,182,198,140\}$, and $\{202,130,183,196,136\}$, corresponding to the rows in the diagram above. Note that these orbits do not depend on the choice of direction made at the start of each line of isogenies, nor do they depend on the choice of $j_{1}$.

Continuing with $p=263$, we compute the products

$$
\begin{aligned}
& P_{1}(X) \equiv_{p}(X-252)(X-38)(X-151)(X-121)(X-258), \\
& P_{2}(X) \equiv_{p}(X-70)(X-112)(X-182)(X-198)(X-140), \\
& P_{3}(X) \equiv_{p}(X-202)(X-130)(X-183)(X-196)(X-136) .
\end{aligned}
$$

Each $\theta_{i k}$ is obtained as the coefficient of $X^{k}$ in the polynomial $P_{i}$ :

$$
\begin{aligned}
& P_{1}(X) \equiv_{p} X^{5}+232 X^{4}+32 X^{3}+159 X^{2}+208 X+158 \\
& P_{2}(X) \equiv_{p} X^{5}+87 X^{4}+252 X^{3}+139 X^{2}+103 X+21 \\
& P_{3}(X) \equiv_{p} X^{5}+205 X^{4}+86 X^{3}+113 X^{2}+121 X+116
\end{aligned}
$$

We then set $y_{1}=\theta_{14}=232, y_{2}=\theta_{24}=87$, and $y_{3}=\theta_{34}=205$. Using the values $y_{i}$ and the $\theta_{i k}$, we compute

$$
\begin{aligned}
V(Y) & =\left(Y-y_{1}\right)\left(Y-y_{2}\right)\left(Y-y_{3}\right) \equiv_{p} Y^{3}+2 Y^{2}+104 Y+59, \\
W_{0}(Y) & =\sum \theta_{i 0} V(Y) /\left(Y-y_{i}\right) \equiv_{p} 32 Y^{2}+259 Y+152, \\
W_{1}(Y) & =\sum \theta_{i 1} V(Y) /\left(Y-y_{i}\right) \equiv_{p} 169 Y^{2}+41 Y+153, \\
W_{2}(Y) & =\sum \theta_{i 2} V(Y) /\left(Y-y_{i}\right) \equiv_{p} 148 Y^{2}+117 Y+277, \\
W_{3}(Y) & =\sum \theta_{i 3} V(Y) /\left(Y-y_{i}\right) \equiv_{p} 107 Y^{2}+115 Y+244 .
\end{aligned}
$$

We do not need $W_{4}$ because $y_{i}=\theta_{i, n-1}$ implies $W_{n-1}\left(y_{i}\right)=y_{i} V^{\prime}\left(y_{i}\right)$, hence the coefficient of $X^{n-1}$ in $U\left(X, y_{i}\right)$ is just $y_{i}$. We complete our work for $p=263$ by updating the CRT coefficient data $(\bmod q)$ for the 15 nontrivial coefficients above, using [42, Algorithm 2.4]. The same procedure is then applied for each $p \in S$.

Having processed all the primes in $S$, a small postcomputation step [42, Algorithm 2.5] yields $V$ and the $W_{k}$ modulo $q$ :

$$
\begin{aligned}
V(Y) & \equiv_{q} Y^{3}+947907 Y^{2}+829791 Y+760884 \\
W_{0}(Y) & \equiv_{q} 975377 Y^{2}+130975 Y+363724 \\
W_{1}(Y) & \equiv_{q} 240332 Y^{2}+135971 Y+616131 \\
W_{2}(Y) & \equiv_{q} 126738 Y^{2}+479879 Y+908580 \\
W_{3}(Y) & \equiv_{q} 340801 Y^{2}+1000285 Y+68659
\end{aligned}
$$

The roots of $V \bmod q$ are $y_{1}=336976, y_{2}=898530$, and $y_{3}=904088$, none of which are roots of $V^{\prime} \bmod q$. Using $y=y_{1}$ we compute

$$
\begin{aligned}
U(X, y) & =X^{5}+y X^{4}+\frac{1}{V^{\prime}(y)}\left(W_{3}(y) X^{3}+W_{2}(y) X^{2}+W_{1}(y) X+W_{0}(y)\right) \\
& \equiv_{q} X^{5}+336976 X^{4}+556976 X^{3}+849678 X^{2}+363260 X+95575
\end{aligned}
$$

and we then find that $x=590272$ is a root of $U(X, y)$, and hence of $H_{D}$, modulo $q$.
This completes the execution of Algorithm 1 on the inputs $D=-971$ and $q=1029167$. If we now set $k=x /(1728-x) \equiv_{q} 638472$, then the elliptic curve $E / \mathbf{F}_{q}$ defined by

$$
Y^{2}=X^{3}+3 k X+2 k \equiv_{q} X^{3}+886249 X+247777
$$

has complex multiplication by the quadratic order with discriminant -971 .

### 3.2. Complexity

The running time of Algorithm 1 has have four principal components. Let us define $T_{\text {find }}$, $T_{\text {enum }}$, and $T_{\text {build }}$ (respectively) as the average expected time, over $p \in S$, to: find $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ (step 5a), enumerate the $G$-orbits of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ (step 5b), and build the polynomials $V$ and $W_{k}$ modulo $p$ (steps 5 c and 5 d ). Additionally, let $T_{\text {root }}$ be the expected time to find a root $y$ of $V \bmod q$ and to compute and find a root of $U(X, y) \bmod q($ steps 7 and 8$)$. As shown in [42], the cost of the precomputation in step 4 and the total cost of all CRT computations are negligible, as are steps $1-3$. Thus the total expected running time is

$$
\begin{equation*}
O\left(|S|\left(T_{\text {find }}+T_{\text {enum }}+T_{\text {build }}\right)+T_{\text {root }}\right) . \tag{10}
\end{equation*}
$$

When $G$ is trivial, or equal to $\operatorname{cl}(\mathcal{O})$, Algorithm 1 reduces to the standard CM method, computing $H_{D} \bmod q$ as in [42, Algorithm 2]. In this case, under the GRH we have the following bounds:

$$
\begin{aligned}
& |S|=O\left(|D|^{1 / 2} \log \log |D|\right), \\
& T_{\text {find }}=O\left(h \log ^{5+\epsilon} h\right), \\
& T_{\text {enum }}=O\left(h \log ^{5+\epsilon} h\right), \\
& T_{\text {build }}=O\left(h \log ^{3+\epsilon} h\right), \\
& T_{\text {root }}=O\left(h \log h \log ^{2+\epsilon} q\right) .
\end{aligned}
$$

The first four bounds are proved in [42, Lemmas 7 and 8].
The last bound is based on the standard probabilistic root-finding algorithm of [8, § 7], using fast arithmetic in $\mathbf{F}_{p}[x]$. With Kronecker substitution, multiplying two polynomials of degree $d$ in $\mathbf{F}_{q}[x]$ uses $O(\mathrm{M}(d \log q))$ bit operations, and this yields an $O(\mathrm{M}(d \log q) \log q)$ bound on the expected time to find a single root of a polynomial in $\mathbf{F}_{q}[x]$ of degree $d$. Here $\mathbf{M}(n)$ denotes the time to multiply two $n$-bit integers [45, Definition 8.26], which we assume to be superlinear, and which satisfies the Schönhage-Strassen bound M $(n)=O(n \log n \log \log n)$; see [39].

When $n=|G|$ properly divides the class number $h$, the times for $T_{\text {build }}$ and $T_{\text {root }}$ may change; these are analyzed below. The times $T_{\text {find }}$ and $T_{\text {enum }}$ are independent of the choice of $G$, as is the space complexity. Depending on the bound $B$, the size of $S$ may also be reduced. This can lead to a substantial practical improvement, depending on the choice of $G$ and the value of $s$, but we defer this issue to $\S 5$. For the moment we simply note that the bound on $|S|$ above holds for any choice of $G$ and for every $s$.

We now consider the time $T_{\text {build }}$. Computing the $\theta_{i k}$ in step 5 c via (5) involves building $m$ polynomials $P_{i}$ of degree $n$ in $\mathbf{F}_{p}[X]$ as products of their linear factors. Let $\mathrm{M}(n)$ denote the time to multiply two $n$-bit integers [45, Definition 8.26], which we assume to be superlinear. With Kronecker substitution, multiplying two polynomials of degree $n$ in $\mathbf{F}_{p}[x]$ uses $O(\mathrm{M}(n \log p))$ bit operations. Using a product tree for each $P_{i}$ yields the bound ${ }^{\dagger}$

$$
\begin{equation*}
O(m \mathrm{M}(n \log p) \log n) \subset O\left(h \log ^{2} n \log ^{1+\epsilon} p\right) \tag{11}
\end{equation*}
$$

on the cost of computing the $\theta_{i k}$. Here we have used $4 p>|D|>h \geqslant n$ and the $O(n \log n \log \log n)$ bound [39] for $\mathrm{M}(n)$, which we note applies to the algorithms used in our implementation. The cost of computing the $y_{i}$ is also dominated by the time to build $m$ polynomials of degree $n$ as products of their linear factors: for each $G$-orbit $G_{i}$ we compute the polynomial $\prod_{j \in G_{i}}(X-j) \in \mathbf{F}_{p}[X]$ whose coefficients are the values of the symmetric functions $e_{1}, \ldots, e_{n}$ on $G_{i}$ (up to a sign), from which we compute the linear combination $y_{i}=s$.

The cost of building $V$ as a product of its linear factors is $O(\mathrm{M}(m \log p) \log m)$, which is dominated by the cost of computing the $n$ polynomials $W_{k}$ as linear combinations of $V(Y) /\left(Y-y_{i}\right)$ with coefficients $\theta_{i k}$. Using a recursive algorithm to compute each $W_{k}$ as in [45, Algorithm 10.9], we obtain a total cost of

$$
\begin{equation*}
O(n \mathbf{M}(m \log p) \log m) \subset O\left(h \log ^{2} m \log ^{1+\epsilon} p\right) \tag{12}
\end{equation*}
$$

for each iteration of step 5 d . The sum of (11) and (12) is $O\left(h \log ^{2} h \log ^{1+\epsilon} p\right)$, essentially the same as the cost of building $H_{D} \bmod p$ from its linear factors. There is an improvement in the implicit constants, but asymptotically we gain at most a factor of 2 in the time $T_{\text {build }}$.

We may gain much more in the time $T_{\text {root }}$. The total expected time for steps 7 and 8 of Algorithm 1 is bounded by

$$
\begin{equation*}
T_{\text {root }}=O(\mathrm{M}(m \log q) \log q+h \mathrm{M}(\log q)+\mathrm{M}(n \log q) \log q) . \tag{13}
\end{equation*}
$$

The three terms in (13) reflect the time to: (a) find root $y$ of $V \bmod q$ that is not a root of $V^{\prime} \bmod q(\mathrm{~b})$ compute $U_{y}(X)=U(X, y) \bmod q$ and (c) find a root of $U_{y} \bmod q$. In (a), any common roots of $V$ and $V^{\prime}$ are first removed from $V$ via repeated division by the GCD, which takes negligible time. We may bound (13) by

$$
\begin{equation*}
T_{\text {root }}=O\left(h \log ^{1+\epsilon} q+(m+n) \log h \log ^{2+\epsilon} q\right), \tag{14}
\end{equation*}
$$

improving $T_{\text {root }}$ by a factor of $\min (\log h \log q, m n /(m+n))$ compared to the time to find a root of $H_{D} \bmod q$. This can reduce the cost of root-finding dramatically, as may be seen in $\S 6$.

## 4. A second algorithm

Algorithm 1 obtains a root of $H_{D} \bmod q$ as a root of $U(X, y) \bmod q$, where $y$ is a root of $V \bmod q$, and $U$ is defined by

$$
\begin{equation*}
U(X, Y)=\frac{1}{V^{\prime}(Y)} \sum_{k=0}^{n} W_{k}(Y) X^{k} \tag{8}
\end{equation*}
$$

[^2]To compute $U(X, y)$ we need to evaluate each $W_{k}$ at $y \in \mathbf{F}_{q}$. We observe that it is not necessary to know the coefficients of $W_{k}$ to do this, we could instead use

$$
\begin{equation*}
W_{k}(Y)=\sum_{i=1}^{m} \theta_{i k} \frac{V(Y)}{\left(Y-y_{i}\right)}, \tag{7}
\end{equation*}
$$

provided that we know $V$ and the $\theta_{i k}$ (which determine the $y_{i}$ ).
Unfortunately we do not know the $\theta_{i k}$ in $\mathbf{F}_{q}$. Algorithm 1 computes the $\theta_{i k}$ in $\mathbf{F}_{p}$, for each prime $p \in S$, but we cannot readily export this knowledge to $\mathbf{F}_{q}$ because the $\theta_{i k}$ do not correspond to the reductions of rational integers ${ }^{\dagger}$. Indeed, the entire reason for using the polynomials $V$ and $W_{k}$ is that they have coefficients in $\mathbf{Z}$ and are thus defined over any field.

However, there is nothing to stop us from using (7) to evaluate $W_{k}$ in $\mathbf{F}_{p}$. Given any $z \in \mathbf{Z}$, we can certainly compute $W_{k}(z) \bmod p$, and if we do this for sufficiently many $p$ we can apply the explicit CRT $(\bmod q)$ to obtain $W_{k}(z) \bmod q$.

In particular, we can apply this to a lift of $y \in \mathbf{F}_{q} \cong \mathbf{Z} / q \mathbf{Z}$ to $\mathbf{Z}$. Explicitly, let $\varphi=\phi \pi$, where $\pi$ is the unique field isomorphism from $\mathbf{F}_{q}$ to $\mathbf{Z} / q \mathbf{Z}$ and $\phi$ maps each residue class in $\mathbf{Z} / q \mathbf{Z}$ to its unique representative in the interval $[0, q-1]$. We then have $W_{k}(\varphi(y)) \equiv \varphi\left(W_{k}(y)\right) \bmod q$.

Thus it suffices to compute $w_{k}=W_{k}(\varphi(y)) \bmod q$, and this can be accomplished by comput$\operatorname{ing} w_{k} \bmod p$ for sufficiently many primes $p$. Note that while $y$ is a root of $V \bmod q$, when we reduce $\varphi(y)$ modulo $p$, we should not expect to get a root of $V \bmod p$, nor do we need to; we are simply evaluating the integer polynomial $W_{k}(Y)$ at the integer $\varphi(y)$, modulo many primes $p$.

This leads to our second algorithm, which proceeds in two stages. The first stage computes $V \bmod q$ and finds a root $y$. The second stage computes the values $w_{k} \bmod q$, then computes $U(X, y) \bmod q$ and finds a root $x$. The second stage requires a bound on $\left|w_{k}\right|$, for which we may use $B m q^{m-1}$, where $B$ bounds the coefficients of the $W_{k}$. For the sake of simplicity we use a single bound for both stages ( $m q^{m-1}$ times the bound used in Algorithm 1), but in practice one may compute separate bounds for each stage (in stage 1 it is only necessary to bound the coefficients of $V$ ). In order to achieve the best space complexity, certain steps are intentionally repeated, and some may require a more careful implementation, see $\S 4.3$ for details.
As before, the choice of $G \subset \operatorname{cl}(\mathcal{O})$ in step 1 determines $n=|G|$ and $m=h / n$, and the values $y_{i}, \theta_{i k}$ and $W_{k}$ are defined for $1 \leqslant i \leqslant m$ and $0 \leqslant k<n$, where the $y_{i}$ depend on the symmetric function $s$ constructed in step 1 .

Algorithm 2. Given $D$ and $q \in \mathcal{P}_{D}$, compute a root $x$ of $H_{D} \bmod q$ as follows.

1. Select $G$, generate a random $s$, and compute $B$, as in Algorithm 1, then set $B \leftarrow m q^{m-1} B$.
2. Compute a polycyclic presentation $\Gamma$ for $\operatorname{cl}(\mathcal{O})$.
3. Use $B$ to select $S \subset \mathcal{P}_{D}$, and perform CRT precomputation $(\bmod q)^{\ddagger}$.
4. For each $p \in S$ :
(a) Find $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ as in [42, Alg. 1], and cache $j_{1}(p)=j_{1}$.
(b) Enumerate the $G$-orbits $G_{i}$ of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ using $j_{1}$ and $\Gamma$.
(c) Compute the $y_{i}$ and $V \bmod p$.
(d) Update CRT data for $V \bmod q$.
5. Perform CRT postcomputation to obtain $V \bmod q$.
6. Find a root $y$ of $V \bmod q$ that is not a root of $V^{\prime} \bmod q$.

If no such root exists then return to step 1 .
7. For each $p \in S$ :
(a) Let $j_{1}=j_{1}(p)$ be the element of $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ computed in step 4a.

[^3](b) Enumerate the $G$-orbits $G_{i}$ of $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ using $j_{1}$ and $\Gamma$.
(c) Compute the values $\theta_{i k}, y_{i}, \phi(y)$ and $V \bmod p$, and then use them to compute the values $w_{k}=W_{k}(\varphi(y)) \bmod p$, via the formula in $(7)^{\dagger}$.
(d) Update CRT data for the $w_{k} \bmod q$.
8. Perform CRT postcomputation to obtain the $w_{k} \bmod q$.
9. Compute $U_{y}(X)=U(X, y) \bmod q$ and output a root $x$ of $U_{y} \bmod q$.

Computing the $w_{k}$ in step 7 c via (7) is faster than computing the coefficients of $W_{k}$, by a factor of $\log ^{2} m$. For a suitable choice of $G$ (with $n=|G|$ small, say polylogarithmic in $h$ ), this may improve the time complexity of the entire algorithm, as shown in § 4.2. However, it is often better to choose $G$ to optimize the bound $B$, as described in $\S 5$, which will tend to make $G$ large.

More significantly, computing the scalars $w_{k}$ rather than the polynomials $W_{k}$ reduces the space complexity to $O(h \log h+(m+n) \log q)$, which may be much better than the $O(h \log q)$ space complexity of Algorithm 1 when $q$ is large. This can even be improved to $O((m+n) \log q)$ using a more intricate implementation, but this may increase the time complexity slightly. The details are given in §4.3, where the space complexity of Algorithm 2 is analyzed.

### 4.1. Example

Let us revisit the example of $\S 3.1$, with $D=-971$ and $q=1029167$. As before, the class group is cyclic of order $h=15$, we let $G$ be the subgroup of order $n=5$, and set $s=-e_{1}$ so that $y_{i}=\theta_{i, n-1}$.

The first stage of Algorithm 2 proceeds as in Algorithm 1, except that our height bound is now $\log _{2}\left(m q^{m-1} B\right) \approx 366$ bits, so we select a slightly larger set $S$, with 29 primes whose product exceeds $4 m q^{m-1} B$. For $p=263$ we again find that the three $G$-orbits of $\operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ are $\{252,38,151,121,258\},\{70,112,182,198,140\}$, and $\{202,130,183,196,136\}$. But instead of computing all the $\theta_{i k}$ as in Algorithm 1, we only need to compute

$$
\begin{aligned}
& y_{1} \equiv_{p}-(252+38+151+121+258) \equiv_{p} 232, \\
& y_{2} \equiv_{p}-(70+112+182+198+140) \equiv_{p} 87, \\
& y_{3} \equiv_{p}-(202+130+183+196+136) \equiv_{p} 205,
\end{aligned}
$$

which yields

$$
V(Y)=\left(Y-y_{1}\right)\left(Y-y_{2}\right)\left(Y-y_{3}\right) \equiv_{p} Y^{3}+2 Y^{2}+104 Y+59
$$

After computing $V \bmod p$ for all $p \in S$ we obtain, via the explicit CRT $(\bmod q)$,

$$
V(Y) \equiv_{q} Y^{3}+947907 Y^{2}+829791 Y+760884
$$

and we again find that $y=336976$ is a root of $V$ and not $V^{\prime}$. We now lift $y$ from $\mathbf{F}_{q} \cong \mathbf{Z} / q \mathbf{Z}$ to the integer $\varphi(y)=336976$ and begin the second stage.

For $p=263$ we recompute the three $G$-orbits as above, and this time we compute all of the $\theta_{i k}$, as we did in Algorithm 1. Setting $y_{1}=\theta_{14}=49, y_{2}=\theta_{24}=87$, and $y_{3}=\theta_{34}=205$, we recompute $V \bmod p$ as above.

We now let $z$ be the reduction of $\varphi(y) \bmod p$ and find that $z=73$ (which is not a root of $V \bmod p$, as expected). We then compute

$$
\begin{aligned}
& z_{1}=V(z) /\left(z-y_{1}\right)=\left(z-y_{2}\right)\left(z-y_{3}\right) \equiv_{p} 7, \\
& z_{2}=V(z) /\left(z-y_{2}\right)=\left(z-y_{1}\right)\left(z-y_{3}\right) \equiv_{p} 211, \\
& z_{3}=V(z) /\left(z-y_{3}\right)=\left(z-y_{1}\right)\left(z-y_{2}\right) \equiv_{p} 122 .
\end{aligned}
$$

[^4]The $z_{i}$ can be computed in two ways: (a) evaluate $V(z)$, invert the $\left(z-y_{i}\right)$, and compute the $z_{i}$ 's as products or (b) simultaneously compute the $z_{i}$ as the complements of the $\left(z-y_{i}\right)$ using a product tree, as in $[42, \S 6.1]$. Both use $O(m)$ field operations, but (b) is faster and works even when $z$ is a root of $V \bmod p$.

We now compute the $w_{k}$ as linear combinations of the $z_{i}$ with coefficients $\theta_{i k}$ :

$$
\begin{aligned}
& w_{0}=W_{0}(z)=\theta_{10} z_{1}+\theta_{20} z_{2}+\theta_{30} z_{3} \equiv_{p} 227, \\
& w_{1}=W_{1}(z)=\theta_{11} z_{1}+\theta_{21} z_{2}+\theta_{31} z_{3} \equiv_{p} 79, \\
& w_{2}=W_{0}(z)=\theta_{12} z_{1}+\theta_{22} z_{2}+\theta_{32} z_{3} \equiv_{p} 44, \\
& w_{3}=W_{0}(z)=\theta_{13} z_{1}+\theta_{23} z_{2}+\theta_{33} z_{3} \equiv_{p} 242 .
\end{aligned}
$$

When evaluated at $z$, the polynomials $W_{k} \bmod p$ computed by Algorithm 1 yield the same values $w_{k}$ above. But here we obtained the $w_{k}$ without computing the $W_{k}$, using just $O(h)$ operations to compute them directly from the $y_{i}$ and the $\theta_{i k}$. Most importantly, the CRT data used to compute the $w_{k} \bmod q$ only consumes $O(m \log q)$ space, versus $O(h \log q)$ for the $W_{k} \bmod q$.

After computing the $w_{k} \bmod p$ for all $p \in S$, the explicit CRT $(\bmod q)$ yields:

$$
w_{0} \equiv_{q} 180694, \quad w_{1} \equiv_{q} 270105, \quad w_{2} \equiv_{q} 92440, \quad w_{3} \equiv_{q} 110998 .
$$

We then evaluate $V^{\prime}(y) \bmod q$ and use the $w_{k}$ to compute

$$
\begin{aligned}
U(X, y) & =X^{5}+y X^{4}+\frac{1}{V^{\prime}(y)}\left(w_{3} X^{3}+w_{2} X^{2}+w_{1} X+w_{0}\right) \\
& \equiv_{q} X^{5}+336976 X^{4}+556976 X^{3}+849678 X^{2}+363260 X+95575
\end{aligned}
$$

and find that $x=590272$ is a root of $U(X, y)$, and hence of $H_{D}$, modulo $q$.

### 4.2. Time complexity

To simplify our analysis we shall initially assume that

$$
\begin{equation*}
m \log q=O\left(|D|^{1 / 2} \log |D|\right) . \tag{15}
\end{equation*}
$$

Depending on $m=h /|G|$, this may allow $\log q$ to be exponentially larger than $\log |D|$, but here we have in mind the case where $\log q$ is polynomial in $\log |D|$; see § 4.4 for an approach better suited to large $q$. Assuming (15), our bound on $|S|$ is the same as in our analysis of Algorithm 1 in §3.2. Under the GRH we have

$$
\begin{equation*}
|S|=O\left(|D|^{1 / 2} \log \log |D|\right), \tag{16}
\end{equation*}
$$

which bounds the total expected number of iterations in steps 4 and 7 .
As with Algorithm 1, the expected running time of Algorithm 2 is bounded by

$$
\begin{equation*}
O\left(|S|\left(T_{\text {find }}+T_{\text {enum }}+T_{\text {build }}\right)+T_{\text {root }}\right), \tag{17}
\end{equation*}
$$

where $T_{\text {find }}, T_{\text {enum }}, T_{\text {build }}$, and $T_{\text {root }}$ are as defined in $\S 3.2$. The term $T_{\text {find }}$ is the same for both algorithms, and the term $T_{\text {enum }}$ is doubled in Algorithm 2. The bound $O\left(h \log ^{2} h \log ^{1+\epsilon} \log p\right)$ on $T_{\text {build }}$ in Algorithm 1 becomes

$$
\begin{equation*}
T_{\text {build }}=O\left(m \log ^{2} m \log ^{1+\epsilon} p+h \log ^{2} n \log ^{1+\epsilon} p\right) \tag{18}
\end{equation*}
$$

for Algorithm 2, with $p=\max S$. The first term in (18) is the time to build $V$, while the second is the time to compute the $\theta_{i k}$ and $w_{k}$. For suitable $n$, say $n=\log ^{c} h$ for some $c>2$, Algorithm 2 effectively reduces $T_{\text {build }}$ by a factor of $\log ^{2} h$. There is also a minor improvement in $T_{\text {root }}$; the bound given in (14) becomes

$$
\begin{equation*}
T_{\text {root }}=O\left((m+n) \log h \log ^{2+\epsilon} q\right) . \tag{19}
\end{equation*}
$$

In both the GRH-based and heuristic complexity analyses of [42, § 7], the bound for $T_{\text {enum }}$ dominates the sum $T_{\text {find }}+T_{\text {enum }}+T_{\text {build }}$. Thus a better bound on $T_{\text {build }}$ does not improve the worst-case complexity. However, the worst-case scenario is atypical, and for almost all $D$ (a set of density 1) this sum is dominated by $T_{\text {build. }}$. This is heuristically argued in [42, § 7], and it can be proven using [6] and assuming the GRH (but we will not do so here).

To remove the worst-case impact of $T_{\text {enum }}$ (and also $T_{\text {find }}$ ), let us consider the performance of Algorithm 2 on a suitably restricted set of discriminants. Fix positive constants $c_{1}, \ldots, c_{7}$, and $\epsilon$. For a positive real parameter $\alpha$, let $\mathcal{D}(\alpha)$ denote the set of negative discriminants $D$ with the following properties.
(i) The set of integers $p \leqslant c_{1}|D| \log ^{1+\epsilon}|D|$ satisfying $4 p=t^{2}-v^{2} D$ with $v \geqslant c_{2} \log ^{1 / 2+\epsilon / 3}|D|$ contains at least $c_{3}|D|^{1 / 2} \log ^{\epsilon / 2}|D|$ primes.
(ii) There are at least two primes $\ell \leqslant c_{4} \log ^{1 / 2}|D|$ for which $\left(\frac{D}{\ell}\right)=1$.
(iii) There is a divisor of $h=h(D)$ in the interval $\left[c_{5} \log ^{\alpha} h, \exp \left(c_{6} \log ^{3 / 4} h\right)\right]$.

Conditions (i) and (ii) ensure that for $D \in \mathcal{D}(\alpha)$, both $T_{\text {find }}$ and $T_{\text {enum }}$ are bounded by $O\left(|D|^{1 / 2} \log ^{5 / 2+\epsilon}|D|\right)$; we refer to $[42, \S 7]$ for details. Provided that $\alpha>1 / 2$, for $D \in \mathcal{D}(\alpha)$ we can choose $G$ so that

$$
c_{5} \log ^{1 / 2} h \leqslant m \leqslant \exp \left(c_{6} \log ^{3 / 4} h\right) \quad \text { and } \quad n \leqslant h /\left(c_{5} \log ^{1 / 2} h\right),
$$

where $n=|G|$ and $m n=h$. Applying (18), we see that $T_{\text {build }}$ is then also bounded by $O\left(|D|^{1 / 2} \log ^{5 / 2+\epsilon}|D|\right)$.

To ensure that our assumption in (15) is satisfied, let us define

$$
\mathcal{P}_{D}(\alpha)=\left\{q \in \mathcal{P}_{D}: \log q \leqslant c_{7} \log ^{1+\alpha}|D|\right\} .
$$

This definition is more restrictive than necessary, but for simplicity we impose a uniform bound. For $q \in \mathcal{P}_{D}(\alpha)$ we then have $T_{\text {root }}=O\left(|D|^{1 / 2+\epsilon}\right)$, which is a negligible component of (17). This yields the following proposition.

Proposition 1. Assume the GRH. For all $\alpha>1 / 2, D \in \mathcal{D}(\alpha)$, and $q \in \mathcal{P}_{D}(\alpha)$ there is a choice of $G$ for which the expected running time of Algorithm 2 is $O\left(|D| \log ^{5 / 2+\epsilon}|D|\right)$.

Proposition 2. For all $\alpha<\log 2$, the set $\mathcal{D}(\alpha)$ has density 1 in the set of all imaginary quadratic discriminants.

Proof. It suffices to show that each of the properties (i), (ii), and (iii) hold for a set of discriminants $D$ with density 1 (in the set of all imaginary quadratic discriminants). For (i), this follows from [6, Theorem 2].

For (ii), consider just the odd primes $\ell_{1}, \ldots, \ell_{k}$ less than $\log \log |D| \leqslant c_{4} \log ^{1 / 2}|D|$, for sufficiently large $|D|$. The proportion of congruence classes modulo $L=\prod \ell_{i}$ corresponding to integers $x$ for which $\left(\frac{x}{\ell_{i}}\right) \neq 1$ for all but at most one $\ell_{i}$ is equal to

$$
\prod \frac{\ell_{i}+1}{2 \ell_{i}}+\sum_{i} \frac{\ell_{i}-1}{2 \ell} \prod_{j \neq i} \frac{\ell_{j}+1}{2 \ell_{j}}<(2 / 3)^{k}+(k / 2)(2 / 3)^{k-1}=o(1) .
$$

Thus the number of discriminants in the interval $[-2 D,-D]$ that do not satisfy property (ii) is $o(|D|)$.

For (iii), recall that for any $\epsilon>0$ and almost all integers $n$ there are at least $(1-\epsilon) \log \log n$ distinct prime divisors of $n$; see [29, Theorem 431]. Therefore almost all discriminants $D$ have at least $k=(1-\epsilon) \log \log |D|-1$ distinct odd prime factors, and for all such discriminants, $h=h(D)$ is divisible by $2^{k-1}$; see [20, Lemma 5.6.8]. As shown by Siegel, $\log h=(1 / 2+$ $o(1)) \log |D|$, thus for all $\alpha<\log 2$ we can choose $\epsilon>0$ so that $2^{k-1}>c_{5} \log ^{\alpha} h$ for all sufficiently large $|D|$.

Propositions 1 and 2 together imply that, under the GRH, for almost all discriminants $D<0$ one can find a root of $H_{D} \bmod q$ in $O\left(|D| \log ^{5 / 2+\epsilon}|D|\right)$ expected time, for all $q \in \mathcal{P}_{D}(\alpha)$ with $1 / 2<\alpha<\log 2$.

We note that the time required to identify and select a suitable $G \subset \operatorname{cl}(\mathcal{O})$ is negligible by comparison. When $D$ is fundamental we can obtain a set of generators for the class group $\mathrm{cl}(\mathcal{O})$ using ideal class representatives of prime norm bounded by $6 \log ^{2}|D|$, under the GRH [5], and for non-fundamental $D$ we also include ideals of prime-power norm for primes dividing the conductor, of which there are $O(\log |D|)$. Given a generating set $S$ for $\mathrm{cl}(\mathcal{O})$ of size $O\left(\log ^{2}|D|\right)$, we can apply generic algorithms to compute the group structure and an explicit basis for $\operatorname{cl}(\mathcal{O})$ consisting of elements of prime-power order, in time $O\left(h^{1 / 2+\epsilon}\right)$, where $h=|\operatorname{cl}(\mathcal{O})|=h(D)$, via [43, Proposition 4$]^{\dagger}$. Once this has been done, it is easy to construct an explicit basis for a subgroup $G$ of any desired order $n$ dividing $h$.

Remark 3. It is usually better to choose $G$ to optimize the height bound $B$ rather than making the choice required by Proposition 1. Heuristically, this yields a better improvement in the time complexity, a factor of nearly $\log |D|$ rather than $\log ^{1 / 2}|D|$, see Heuristic Claim 1 in § 5 .

### 4.3. Space complexity

For convenience we assume the GRH and the restriction (15) on the size of $q$. An alternative approach with a weaker restriction on $q$ is given in the next section. Under these assumptions, $\log p \sim \log |D|$ for all $p \in S$, and $|S|=O\left(|D|^{1 / 2} \log \log |D|\right)$, as in (16). A straightforward implementation of Algorithm 2 yields a space complexity of

$$
\begin{equation*}
O(|S| \log |D|+|S| \log q+h \log |D|+(m+n) \log q) \tag{20}
\end{equation*}
$$

The first term of (20) represents storage for the set $S$, the second term is storage for precomputed data used to apply the explicit CRT $(\bmod q)$, the third term is storage for $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$, and the last term is storage for the $n$ values $w_{k}$ and the $m$ coefficients of $V$ that are computed via the explicit CRT $(\bmod q)$. The second term dominates, and we can immediately bound (20) by $O\left(|D|^{1 / 2} \log \log |D| \log q\right)$, which is the same as the space complexity of Algorithm 1.

However, the only essential term in (20) is the last one, which may be as small as $O\left(|D|^{1 / 4} \log q\right)$. We now show how to eliminate the first three terms in (20). In practice the most useful term to eliminate (or reduce) is the second one, which requires only a very minor change, but we consider each in turn.
4.3.1. The first term. We wish to avoid storing the entire set $S$ at any one time. Our strategy is to process $S$ in batches of size $O\left(|D|^{c}\right)$, for some positive $c<1 / 4$.

In [42] the set $S$ is chosen by enumerating a larger subset $S_{z} \subset \mathcal{P}_{D}$ defined by a parameter $z$ that depends on the bound $B$. The primes in $S_{z}$ satisfy the norm equation $4 p=t^{2}-v^{2} D$, where $v$ is $O\left(\log ^{3+\epsilon}|D|\right)$, and the bound on $t$ depends on $v, h(D)$, and $z$, but is in any case $O\left(|D|^{1 / 2+\epsilon}\right)$. As an alternative to the sieving approach of [42, Algorithm 2.1], we enumerate $S_{z}$ by running through all the integers of the form $\left(t^{2}-v^{2} D\right) / 4$, with $v$ and $t$ suitably bounded, and applying a polynomial-time primality test [2] to each. This takes $O\left(|D|^{1 / 2+\epsilon}\right)$ time and negligible space.

Each prime $p$ in $S_{z}$ is assigned a 'rating' $r(p)$, that reflects the expected cost of finding $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$. The details are not important here, but we may assume that the positive real numbers $r(p)$ are distinctly represented using a precision of $O(\log |D|)$ bits. Let $r_{\max }$ be the

[^5]largest (worst) rating among the primes in $S_{z}$. For a suitable $r \in\left[0, r_{\max }\right]$, we then let $S$ consist of the primes in $S_{z}$ with ratings bounded by $r$, where $r$ is chosen so that $S$ has the appropriate size. We can determine such an $r$ space-efficiently using a binary search on the interval [ $0, r_{\text {max }}$ ], enumerating $S_{z}$ at most $\left\lceil\log _{2}\left|S_{z}\right|\right\rceil$ times. Similarly, we can partition $S$ into batches of size $|D|^{c}$ by partitioning the interval $[0, r]$ into subintervals whose endpoints are determined using a binary search. The total cost of computing this partition is $O\left(|D|^{1-c+\epsilon}\right)$, and this also bounds the cost of enumerating all the batches in $S$. Thus the time complexity is negligible and we use $o\left(|D|^{1 / 4}\right)$ space.
4.3.2. The second term. Let $M$ be the product of the primes $p_{i}$ in $S$. As described in $\S 2.3$, when applying the explicit CRT $(\bmod q)$, we precompute integers $M_{i}=M / p_{i} \bmod q$ and $a_{i} \equiv\left(M / p_{i}\right)^{-1} \bmod p_{i}$; this would normally be done in step 3 of Algorithm 2. The total size of the $M_{i}$ is $O(|S| \log q)$, which accounts for the second term of (20). Rather than computing the $M_{i}$ in step 3 , we just compute $M \bmod q$ in step 3 , and when we need $M_{i}$ during a CRT update step for $p_{i}($ steps 4 d and 7 d$)$, we invert $p_{i} \bmod q$ and multiply to obtain $M_{i}=\left(M / p_{i}\right) \bmod q$. This reduces the second term of (20) from $O(|S| \log q)$ to $O(|S| \log |D|+\log q)$, which is the space used by the $a_{i}$ and $M \bmod q$.

With the changes described in $\S 4.3 .1$, we compute the primes $p_{i}$ in $S$ in batches of size $|D|^{c}$, where $c<1 / 4$. We now do the same with the $a_{i}$, computing the $a_{i}$ for each batch of primes $p_{i}$ as follows. If $N$ is the product of the primes in the batch, then it suffices to compute the integer $(M / N)$ modulo $N$, and then simultaneously compute the product of $(M / N)$ and $\left(N / p_{i}\right)$ modulo $p_{i}$, for all the primes $p_{i}$ in the batch, to obtain the $a_{i}$. The $a_{i}$ (and the $p_{i}$ ) for a given batch are stored only as long as it takes to process the batch; this means that they are computed twice in Algorithm 2: once in step 4 and then again in step 7. Using a product tree for each batch, it takes $O\left(|D|^{1-c+\epsilon}\right)$ time to compute the $a_{i}$ for all the batches, which is negligible. With this change, the second term of (20) is reduced to $o\left(|D|^{1 / 4}\right)+O(\log q)$.

The only data that is retained once the processing of a given batch of primes has been completed are two values associated to each coefficient $c$ that is to be computed modulo $q$ (the $m$ coefficients of $V \bmod q$ in step 5 , and the $n$ values $w_{k} \bmod q$ ). These two values are the partial sums $\sum c_{i} a_{i} M_{i}$ and $\sum c_{i} a_{i} / p_{i}$, where the first is computed modulo $q$ and the latter is stored with a precision of $O(\log q)$ bits, as described in $\S 2.3$. These two values are updated as each prime $p_{i}$ is processed (across all the batches), and never require more than $O(\log q)$ bits. There are a total of $m+n$ coefficients, which accounts for the fourth term $O((m+n) \log q)$ listed in (20), which also bounds the space required to perform the CRT postcomputation; using the accumulated partial sums, the postcomputation is effectively just a single subtraction to evaluate (3).
4.3.3. The third term. Steps 4b and 7b of Algorithm 2 both enumerate the $G$-orbits of $j_{1} \in \operatorname{Ell}_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$, using the polycyclic presentation $\Gamma$. Let $j_{i k}$ denote the $k$ th element of the $G$ orbit $G_{i}=\left(j_{i 1}, \ldots, j_{i n}\right)$, where $i$ ranges from 1 to $m$ and $k$ ranges from 1 to $n$. There are a total of $h=m n$ values $j_{i k}$, and these account for the $O(h \log |D|)$ third term in (20), using the GRH bound $\log p=O(\log |D|)$. In order to reduce the space required, we need to process the $j_{i k}$ as they are enumerated, rather than storing them all. Our basic strategy is to order the enumeration so that we compute the $G$-orbits one by one and discard each $G$-orbit once it has been processed (in fact, we need to enumerate the $G$-orbits twice in step 7 in order to achieve this). Depending on the presentation $\Gamma$, enumerating the $G$-orbits efficiently may present some complications. These will be addressed below. We first consider how to process the $G$-orbits in a space-efficient manner as they are enumerated.

Let us recall the values we must derive from the $j_{i k}$. In each iteration of step 4 we compute $m$ values $y_{i} \bmod p$, each of which is a linear combination of elementary symmetric functions applied to $G_{i}=\left(j_{i 1}, \ldots, j_{i n}\right)$. We then compute the polynomial $V \bmod p$ whose roots are
the $y_{i}$. In each iteration of step 7 we again compute the $y_{i} \bmod p$, and also the coefficients $\theta_{i k}$ of the polynomials

$$
P_{i}(X)=\prod_{k=1}^{n}=\left(X-j_{i k}\right)=\sum_{k=0}^{n} \theta_{i k} X^{k}
$$

defined in (5). Up to a sign, the $\theta_{i k}$ are just the elementary symmetric functions of $j_{i 1}, \ldots, j_{i n}$, so we can derive each $y_{i}$ from the $\theta_{i k}$. The $y_{i}$ are then used to compute $V \bmod p$ (again), and also the values $z_{i}=V(z) /\left(z-y_{1}\right)$, where $z=\varphi(y) \bmod p$ is the reduction of the integer $\varphi(y)$ corresponding to the root $y$ of $V \bmod q$ computed in step 6 . Finally, we compute the values $w_{k}=\sum_{i=1}^{m} \theta_{i k} z_{i}$, as described in §4.1.

The space required by the $y_{i}$, the $z_{i}$, the $w_{k}$ and the polynomials $V$ and $V^{\prime}$ is just $O((m+n) \log p)$, which is within our desired complexity bound. We now explain how to process the $G$-orbits in a way that achieves this space complexity. For each $G_{i}$ we compute the polynomial $P_{i}$ with coefficients $\theta_{i k}$, use the $\theta_{i k}$ to compute $y_{i}$, and then discard $G_{i}$ and the $\theta_{i k}$. In step 4 this is all that is required; once all the $G$-orbits have been processed we compute $V(Y)=\prod_{i=1}^{m}\left(Y-y_{i}\right)$. In step 7 we proceed as in step 4, and after computing $V$ and the $y_{i}$ we compute $V^{\prime}$ and the values $z_{i}=V(z) /\left(z-y_{i}\right)$. We then enumerate the $G$-orbits a second time, recompute the $\theta_{i k}$ as above, and then update each of $n$ partial sums $w_{k}=\theta_{i k} z_{i}$ by adding the term $\theta_{i k} z_{i}$. The space required to process a $G$-orbit is $O(n \log p)$ for the $j_{i k}$ and the $\theta_{i k}$, which are then discarded, plus a total of $O((m+n) \log p)$ space used to store the value $y_{i}$ and $z_{i}$, and the partial sums $w_{k}$ that are retained. Thus we can process all the $G$-orbits using $O((m+n) \log p)$ space.

We now consider the enumeration of the $G$-orbits. If the polycyclic presentation $\Gamma$ for $\operatorname{cl}(\mathcal{O})$ uses the sequence of ideal classes $\left[\mathfrak{l}_{1}\right], \ldots,\left[\mathfrak{l}_{r}\right]$ and the chosen subgroup $G$ is generated by a prefix of $\Gamma$, say

$$
\begin{equation*}
G=\left\langle\left[\mathfrak{l}_{1}\right], \ldots,\left[\mathfrak{l}_{d}\right]\right\rangle, \tag{21}
\end{equation*}
$$

for some $d<r$, then the elements $j_{i 1}, \ldots, j_{\text {in }}$ of the $G$-orbit $G_{i}$ will appear consecutively in the enumeration of $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ obtained using $\Gamma$ and no special processing is required. But the situation in (21) is very special, much more so than condition (9) given in §2.6. Indeed, it forces $G$ to be trivial if $r=1$. We can ensure that $G$ has the form in (21) if we construct $\Gamma$ by computing a polycyclic presentation for $G$ and extending it to a polycyclic presentation for $\operatorname{cl}(\mathcal{O})$. Unfortunately, the norms arising in a polycyclic presentation for a proper subgroup of $\operatorname{cl}(\mathcal{O})$ may be very large: $\Omega\left(|D|^{1 / 3}\right)$ in the counterexample of [42, § 5.3].

To address this, we compute the action of ideals with uncomfortably large norm by representing them as a product of ideals with small norms. Assuming the GRH, it follows from [16, Theorem 2.1] that every element $[\mathfrak{a}]$ of $\operatorname{cl}(\mathcal{O})$ can be expressed in the form $[\mathfrak{a}]=\left[\mathfrak{p}_{1} \ldots \mathfrak{p}_{t}\right]$, where the $\mathfrak{p}_{i}$ are ideals of prime norm bounded by $\log ^{c}|D|$, for any $c>2$, and $t=\lceil C \log h / \log \log |D|\rceil=O(\log |D|)$, for some constant $C$ that depends on $c$. We can find such a representation in $O\left(h^{1+\epsilon}\right)$ expected time (using negligible space), by simply generating random products of the form $\left[\mathfrak{p}_{1} \ldots \mathfrak{p}_{t}\right]$ until we find $[\mathfrak{a}]^{\dagger}$. After precomputing such a representation in step 2, we can compute the action of any element of $\operatorname{cl}(\mathcal{O})$ on any element of $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ in $O\left(\log ^{6+\epsilon}|D|\right)$ expected time, allowing us to efficiently handle arbitrarily large norms in $\Gamma$.

In contrast to the first two terms, removing the third term from (20) may increase the overall running time significantly. The time complexity is increased by a logarithmic factor, but in return, the space complexity may be reduced by an exponential factor. We did not make this trade-off in our implementation, as the space complexity of $O(h \log h+(m+n) \log q)$

[^6]achieved by simply optimizing the second term is already more than sufficient for the range of $D$ used in the examples of $\S 6$. However for larger computations, increasing the running time by a polylogarithmic factor in order to reduce the space by an exponential factor may be very attractive, especially in a parallel implementation.

Eliminating the first three terms of (20) leads to the following proposition, which was summarized in the introduction. The constraints on $m$ and $q$ ensure that (15) is satisfied.

Proposition 3. Assume the GRH, and fix $\delta \geqslant 0$. Let $D$ be a discriminant with class number $h=m n$ and let $q \in \mathcal{P}_{D}$, such that $m \leqslant O\left(|D|^{1 / 2-\delta}\right)$ and $\log q=O\left(|D|^{\delta} \log |D|\right)$. With modifications 4.3.1-3, the expected running time of Algorithm 2 on inputs $D$ and $q$ is $O\left(|D| \log ^{6+\epsilon}|D|\right)$, using $O((m+n) \log q)$ space.

### 4.4. Handling large $q$

When $q$ is not constrained by (15), the bound $m q^{m-1} B$ used by Algorithm 2 to determine the size of $S$ may lead to a significant increase in the running time of Algorithm 2 relative to Algorithm 1, which just uses the bound $B$. To better handle large $q$ in a space-efficient manner, we make a minor modification to Algorithm 2 that allows us to use the bound $m q B$ instead. Unless $q$ is extraordinarily large, this will not be significantly different than using the bound $B$. We are forced to give up the improved bound on $T_{\text {build }}$ in (18), but the resulting algorithm will still have a running time that is at worst twice that of Algorithm 1, and will typically be only slightly slower.

The key is to avoid exponentiating $\varphi(y)$ in $\mathbf{F}_{p}$. Instead we compute all of the powers $y, y^{2}, y^{3}, \ldots, y^{m-1}$ in $\mathbf{F}_{q}$, and then lift these to integers $\varphi(y), \ldots, \varphi\left(y^{m-1}\right)$ in the interval $[0, q-1]$, which can be reduced modulo $p$. This is done between steps 6 and 7 of Algorithm 2. We now modify step 5 c to compute the coefficients of $W_{k}=\sum a_{i k} Y^{i} \bmod p$ as in Algorithm 1, and then compute values $w_{k}^{\prime} \bmod p$ via

$$
w_{k}^{\prime}=a_{m-1, k} \varphi\left(y^{m-1}\right)+a_{m-2, k} \varphi\left(y^{m-2}\right)+\ldots+a_{1, k} \varphi(y)+a_{0, k},
$$

that we use instead of $w_{k}=W_{k}(\varphi(y)) \bmod p$. Note that the integers $w_{k}^{\prime}$ are not equal to the integers $W_{k}(\varphi(y))$, but we have $w_{k}^{\prime} \equiv W_{k}(\varphi(y)) \bmod q$, which is all that is required, even though $w_{k}$ and $w_{k}^{\prime}$ will typically be distinct modulo $p$.

We may combine this approach with any of the space optimizations considered in the previous section. When $q$ is very large, the third term of (20) is likely to be dominated by the others, so we only optimize the first two terms of (20). Provided $\log q=O\left(|D|^{1 / 2}\right)$, the analysis in $[42, \S 7]$ then yields an upper bound on the running time of this modified version of Algorithm 2.

Proposition 4. Assume the GRH. Let $D$ be a discriminant with class number $h=m n$, and let $q \in \mathcal{P}_{D}$ satisfy $\log q=O\left(|D|^{1 / 2}\right)$. With modifications 4.3.1-2 and 4.4, the expected running time of Algorithm 2 on inputs $D$ and $q$ is $O\left(|D| \log ^{5+\epsilon}|D|\right)$, using $O((m+n) \log q+h \log h)$ space.

## 5. Height bounds

In this section we derive an upper bound $B$ on the absolute values of the integer coefficients of $V$ and $W_{k}$ defined in $\S 2.5$. More precisely, we compute a height bound $b$ on the maximum bit-length of any coefficient occurring in $V$ or $W_{k}$. This is used by Algorithms 1 and 2 to choose a set of CRT primes $S$ whose product exceeds $4 B=2^{b+2}$.

We first derive a general height bound $b_{\max }$ that depends only on $D$ and can be used with any choice of the subgroup $G \subset \operatorname{cl}(\mathcal{O})$ and the random symmetric function
$s=e_{1}+c_{2} e_{2}+\ldots+c_{n} e_{n}$ constructed in Algorithms 1 and 2. Under the GRH, we have $b_{\max }=$ $O\left(|D|^{1 / 2} \log |D| \log \log |D|\right)$, which is all that is needed for the proofs of Propositions 1-4.
We then fix $s=e_{1}$ and derive height bounds that depend not only on $D$, but also on the subgroup $G$. As may be seen in Table 2 , the actual heights can vary significantly with $G$. An optimal choice of $G$ may improve the performance of both Algorithms 1 and 2 substantially, provided that we have a height bound that accurately reflects the impact of $G$. Heuristically, we expect to reduce $b$ by a factor of $\log h / \log \log h$, on average, by computing a customized $b$ for each candidate subgroup $G$ and choosing the best one. Of course the optimal choice of $G$ depends not only on $b$, but also on how the size of $G$ impacts the complexity of building polynomials and finding roots, but when $\log q \ll h$ the height bound is usually the most critical factor.

Throughout this section we work with $j$-invariants, which allows us to use rigorous (and quite accurate) bounds on their size. Heuristic bounds for other class invariants can be obtained by scaling linearly, as discussed in $\S 6$. We note that all the bounds we derived in this section hold unconditionally; the GRH and the heuristic analysis in § 5.4 are only used to obtain asymptotic growth estimates.

### 5.1. Height bound derivations

As in $\S 2.5$, let $G=\left\{\beta_{1}, \ldots, \beta_{n}\right\}$ be a subgroup of $\operatorname{cl}(\mathcal{O})$ with cosets $\alpha_{1} G, \ldots, \alpha_{m} G$, so that every element of $\operatorname{cl}(\mathcal{O})$ is of the form $\alpha_{i} \beta_{k}$ with $1 \leqslant i \leqslant m$ and $1 \leqslant k \leqslant n$. We may uniquely represent $\alpha_{i} \beta_{k}$ by a primitive reduced binary quadratic form $A_{i k} x^{2}+B_{i k} x y+C_{i k} y^{2}$ with discriminant $D$. If $\tau_{i k}$ denotes the complex number $\left(-B_{i k}+\sqrt{D}\right) /\left(2 A_{i k}\right)$, we have

$$
H_{D}(X)=\prod_{i, k}\left(X-j\left(\tau_{i k}\right)\right)
$$

where $j(z)$ is the classical modular function. We assume without loss of generality that $\alpha_{1}$ and $\beta_{1}$ are the identity element, and fix $x=j\left(\tau_{11}\right)$ so that $\left[\alpha_{i} \beta_{k}\right] x=j\left(\tau_{i k}\right)$. We shall use the explicit bound

$$
\begin{equation*}
\left|j\left(\tau_{i k}\right)\right| \leqslant \exp \left(\pi \sqrt{|D|} / A_{i k}\right)+2114.567 \tag{22}
\end{equation*}
$$

proved in [21, p. 1094], and define $b_{i k}$ to be the logarithm in base 2 (denoted ' ${ }^{\prime}$ ') of the RHS of (22), and we note that $b_{i k}>0$. We define the height $\mathrm{ht}(F)$ of a nonzero polynomial $F(X)=$ $\sum c_{j} X^{j}$ in $\mathbf{C}[X]$ by $\operatorname{ht}(F)=\lg \max \left|c_{j}\right|$. We seek an upper bound $b$ on $\max \left\{\operatorname{ht}(V), \operatorname{ht}\left(W_{k}\right)\right\}$ in terms of the $b_{i k}$.

Table 2. Actual heights for various $G \subset \operatorname{cl}(\mathcal{O})$ with $D=-221606831$.

| $\|G\|$ | Bits | $\|G\|$ | Bits | $\|G\|$ | Bits | $\|G\|$ | Bits |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 1 | 1983568 | 35 | 1017514 | 182 | 639986 | 910 | 395909 |
| 2 | 1737305 | 39 | 955880 | 195 | 672404 | 1001 | 444642 |
| 3 | 1600984 | 42 | 959237 | 210 | 649274 | 1155 | 413905 |
| 5 | 1464042 | 55 | 879633 | 231 | 607751 | 1365 | 392521 |
| 6 | 1430692 | 65 | 841574 | 273 | 603539 | 1430 | 402990 |
| 7 | 1354754 | 66 | 780769 | 286 | 540873 | 2002 | 414360 |
| 10 | 1286551 | 70 | 877290 | 330 | 531985 | 2145 | 422627 |
| 11 | 1235548 | 77 | 791884 | 385 | 522887 | 2310 | 401968 |
| 13 | 1202022 | 78 | 760840 | 390 | 540120 | 2730 | 409766 |
| 14 | 1188816 | 91 | 756960 | 429 | 525472 | 3003 | 436780 |
| 15 | 1195102 | 105 | 773983 | 455 | 430383 | 4290 | 471475 |
| 21 | 1093207 | 110 | 677448 | 462 | 487746 | 5005 | 507403 |
| 22 | 962794 | 130 | 720919 | 546 | 492453 | 6006 | 549648 |
| 26 | 1010539 | 143 | 697728 | 715 | 452019 | 10010 | 756598 |
| 30 | 1006310 | 154 | 616795 | 770 | 429293 | 15015 | 1039684 |
| 33 | 998157 | 165 | 678832 | 858 | 437618 | 30030 | 1983568 |

The largest $b_{i k}$ is $b_{11}$, since $\alpha_{1} \beta_{1}$ is the identity and we therefore have $A_{11}=1$. The bound (22) is nearly tight (one can prove a similar lower bound), and this implies that $b_{11} \approx \lg (e) \pi \sqrt{|D|}$. This is about ten times the typical value of $h$, so we shall not be concerned with optimizing terms that are asymptotically smaller than $h$, which we may assume includes both $m$ and $n$, and even $m \lg n$ and $n \lg m$.

As in $\S 2.5$, the $\theta_{i k}$ are defined via

$$
\begin{equation*}
P_{i}(X)=\prod_{k=1}^{n}\left(X-j\left(\tau_{i k}\right)\right)=\sum_{k=1}^{n} \theta_{i k} X^{k} \tag{23}
\end{equation*}
$$

where $\prod_{i=1}^{m} P_{i}=H_{D}$. This yields the bound

$$
\begin{equation*}
\operatorname{ht}\left(P_{i}\right) \leqslant n+\sum_{k=1}^{n} b_{i k} \tag{24}
\end{equation*}
$$

For $V(Y)=\prod_{i=1}^{m}\left(Y-y_{i}\right)$, with $y_{i}=s\left(j\left(\tau_{i 1}\right), \ldots, j\left(\tau_{i n}\right)\right)$, we assume that $s$ is of the form $s= \pm\left(e_{1}+c_{2} e_{2}+\ldots+c_{n} e_{n}\right)$ with $c_{k} \geqslant 0$, and define $z_{i}=\left|s\left(2^{b_{i 1}}, \ldots, 2^{b_{i n}}\right)\right|$. We then have $\left|y_{i}\right| \leqslant z_{i}$, with $z_{i}>1$. Thus

$$
\begin{equation*}
\operatorname{ht}(V) \leqslant m+\sum_{i=1}^{m} \lg z_{i} \tag{25}
\end{equation*}
$$

For the nonzero polynomials $W_{k}(Y)=\sum_{i=1}^{m} \theta_{i k} \prod_{\hat{\imath} \neq i}\left(Y-y_{\hat{\imath}}\right)$, we note that

$$
\sum_{i=1}^{m}\left|\theta_{i k}\right| 2^{m} \prod_{\hat{\imath} \neq i} z_{\hat{\imath}}
$$

is positive, and bounds every coefficient of $W_{k}$. We have

$$
\begin{align*}
\operatorname{ht}\left(W_{k}\right) & \leqslant \lg \left(\sum_{i=1}^{m}\left|\theta_{i k}\right| 2^{m} \prod_{\hat{\imath} \neq i} z_{\hat{\imath}}\right) \leqslant \lg \left(m \max _{i}\left|\theta_{i k}\right| 2^{m} \prod_{\hat{\imath} \neq i} z_{\hat{\imath}}\right) \\
& \leqslant m+\lg m+\max _{i}\left(\operatorname{ht}\left(P_{i}\right)+\sum_{\hat{\imath} \neq i} \lg z_{\hat{\imath}}\right) . \tag{26}
\end{align*}
$$

This expression does not depend on $k$, thus it applies to every nonzero $W_{k}$.
5.1.1. A general height bound. For any choice of $s=e_{1}+c_{2} e_{2}+\ldots+c_{n} e_{n}$ with $0 \leqslant c_{k}<$ $2 m^{2}$, we have $z_{i} \leqslant 2 m^{2} n 2^{n-1} 2^{b_{i 1}+\ldots+b_{i n}}$, using $\binom{n}{k} \leqslant 2^{n-1}$, and therefore

$$
\begin{equation*}
\lg z_{i} \leqslant n+\lg n+2 \lg m+\sum_{k=1}^{n} b_{i k} \tag{27}
\end{equation*}
$$

From (25) we obtain

$$
\begin{equation*}
\operatorname{ht}(V) \leqslant m+m n+m \lg n+2 m \lg m+\sum_{i, k} b_{i k} \leqslant 3 h+2 h \lg h+\sum_{i, k} b_{i k} \tag{28}
\end{equation*}
$$

where we have used $h=m n \geqslant m \lg n$. Using the crude bounds $\max _{i} \sum_{k} b_{i k} \leqslant \sum_{i, k} b_{i k}$ and $\sum_{\hat{\imath} \neq i} \lg z_{\hat{\imath}} \leqslant \sum_{i} \lg z_{i}$, and applying (24) and (27) to (26), we obtain

$$
\begin{align*}
\operatorname{ht}\left(W_{k}\right) & \leqslant m+\lg m+n+m n+m \lg n+2 m \lg m+2 \sum_{i, k} b_{i k} \\
& \leqslant 5 h+2 h \lg h+2 \sum_{i, k} b_{i k} \tag{29}
\end{align*}
$$

Since the bound in (29) dominates the bound in (28), we define

$$
\begin{equation*}
b_{\max }=5 h+2 h \lg h+2 \sum_{i, k} b_{i k}, \tag{30}
\end{equation*}
$$

where $i$ runs from 1 to $m$ and $k$ runs from 1 to $n$.
Recall that, under the GRH, we have the bounds $h=O\left(|D|^{1 / 2} \log \log |D|\right)$ and $\sum_{i, k} 1 / A_{i k}=$ $O(\log |D| \log \log |D|)$, as noted in $\S 2.4$. These imply

$$
\begin{equation*}
b_{\max }=O\left(|D|^{1 / 2} \log |D| \log \log |D|\right) \tag{31}
\end{equation*}
$$

5.1.2. Optimized height bounds. We now fix $s=e_{1}$, which implies $z_{i}=\sum_{k=1}^{n} b_{i k}$. This choice of $s$ minimizes our height bound and simplifies the calculations.

We have

$$
\begin{equation*}
\lg z_{i} \leqslant \lg \left(\sum_{k=1}^{n} 2^{b_{i k}}\right) \leqslant \lg n+\max _{k} b_{i k} \tag{32}
\end{equation*}
$$

and from (25) we find that

$$
\begin{equation*}
\operatorname{ht}(V) \leqslant m+m \lg n+\sum_{i=1}^{m} \max _{k} b_{i k} . \tag{33}
\end{equation*}
$$

Applying (24) and (32) to (26) yields

$$
\begin{aligned}
\mathrm{ht}\left(W_{k}\right) & \leqslant \lg m+m+\max _{i}\left(n+\sum_{k} b_{i k}+\sum_{\hat{\imath} \neq i}\left(\lg n+\max _{k} b_{\hat{i} k}\right)\right) \\
& \leqslant \lg m+m+n+m \lg n+\max _{i}\left(\sum_{k} b_{i k}+\sum_{\hat{\imath}} \max _{k} b_{\hat{\imath} k}-\max _{k} b_{i k}\right) \\
& \leqslant \lg m+m+n+m \lg n+\sum_{i} \max _{k} b_{i k}+\max _{i}\left(\sum_{k} b_{i k}-\max _{k} b_{i k}\right),
\end{aligned}
$$

where $i$ runs from 1 to $m$ and $k$ runs from 1 to $n$. This bound dominates the bound in (33), so it bounds $\operatorname{ht}(V)$ as well as $\operatorname{ht}\left(W_{k}\right)$. Thus we define

$$
\begin{equation*}
b=\lg m+m+n+m \lg n+\sum_{i} \max _{k} b_{i k}+\max _{i}\left(\sum_{k} b_{i k}-\max _{k} b_{i k}\right) \tag{34}
\end{equation*}
$$

as our height bound for $G$, which we typically round up to the nearest integer.

### 5.2. Example

Returning to our example with $D=-971$ and $h(D)=15$, let us compute $b$ for the subgroup $G \subset \operatorname{cl}(\mathcal{O})$ of order $n=5$. We can use the same polycyclic presentation $\left[\mathrm{l}_{1}\right],\left[\mathrm{l}_{2}\right]$ for $\operatorname{cl}(\mathcal{O})$ as before, where the ideals $\mathfrak{l}_{1}$ and $\mathfrak{l}_{2}$ have norms $\ell_{1}=3$ and $\ell_{2}=5$ and $\left[\mathfrak{l}_{1}\right]$ generates $G$, but now we compute directly in the class group, using composition of binary quadratic forms [14] rather than computing isogenies. In the notation of $\S 5.1$, we have $\beta_{k}=\left[{ }_{1}^{k-1}\right]$ and $\alpha_{i}=\left[l_{2}^{i-1}\right]$. Enumerating $\operatorname{cl}(\mathcal{O})$ yields the approximate values

$$
\begin{array}{lllll}
b_{11}=141.23, & b_{12}=47.08, & b_{13}=15.75, & b_{14}=15.75, & b_{15}=47.08, \\
b_{21}=28.25, & b_{22}=11.45, & b_{23}=20.18, & b_{24}=11.96, & b_{25}=11.45, \\
b_{31}=11.45, & b_{32}=11.96, & b_{33}=20.18, & b_{34}=11.45, & b_{35}=28.25,
\end{array}
$$

where each row corresponds to a coset of $G$. The value of $b_{22}$, for example, is computed using $A_{22}=15$, since $\alpha_{2} \beta_{2}=\left[l_{2} \mathfrak{l}_{1}\right]$ is represented by the reduced binary quadratic form $15 X^{2}+13 X Y+19 Y^{2}$, and we have

$$
b_{22}=\lg (\exp (\pi \sqrt{971} / 15)+2114.567) \approx 11.45
$$

To compute $b$ we just need the sum $s_{i}$ and maximum $t_{i}$ of $b_{i k}$ over $k$. These can be computed during the enumeration and stored using $O(m \log |D|)$ space. We have $s_{1}=266.89, s_{2}=83.28$, $s_{3}=83.38$, and $t_{1}=141.23, t_{2}=28.25, t_{3}=28.25$. The $t_{i}$ sum to 197.73 , the maximum of $s_{i}-t_{i}$ is 125.65. Applying (34) yields

$$
b \approx \lg 3+5+3+3 \lg 5+197.73+125.65 \leqslant 340,
$$

and our height bound is 340 bits.
As noted earlier, we may compute $V$ and the $W_{k}$ over $\mathbf{Z}$ using Algorithm 1, by computing them modulo some $q \geqslant 2^{b+1}$. We find that $\max \left\{\operatorname{ht}(V), \operatorname{ht}\left(W_{k}\right)\right\}$ is in fact about 324, within five percent of our computed bound $b$.

For comparison, if we instead choose $G$ to be the subgroup of order 3, we obtain $s_{1}=$ $165.15, s_{2}=55.45, s_{3}=78.71, s_{4}=78.71, s_{5}=55.45$, and $t_{1}=141.23, t_{2}=28.25, t_{3}=47.08$, $t_{4}=47.08, t_{5}=28.25$. The $t_{i}$ sum to 291.89, the maximum of $s_{i}-t_{i}$ is 31.63 , and (34) becomes

$$
b \approx \lg 5+3+5+5 \lg 3+291.89+31.63 \leqslant 342 .
$$

If we make $G$ trivial, we have $t_{i}=s_{i}=b_{i 1}$, the sum of the $t_{i}$ is 421.51 , the maximum of $s_{i}-t_{i}$ is zero, and we get $b=438$, which is nearly the same as the value 434 one obtains from [42, Lemma 8], which uses a more careful analysis than we do here.

### 5.3. Computational complexity of optimizing the height bound

We can compute a basis $\left(\gamma_{1}, \ldots, \gamma_{r}\right)$ for the class $\operatorname{group} \operatorname{cl}(\mathcal{O})$ in $O\left(|D|^{1 / 4+\epsilon}\right)$ time and space, under the GRH [14, Proposition 9.7.16]. We may assume that each $\gamma_{i}$ has prime power order $\ell_{i}^{e_{i}}$. We then consider subgroups $G$ of the form

$$
\begin{equation*}
G=\left\langle\gamma_{1}^{\ell_{1}^{d_{1}}}\right\rangle \times \ldots \times\left\langle\gamma_{r}^{\ell_{r}^{d_{r}}}\right\rangle \tag{35}
\end{equation*}
$$

with $0 \leqslant d_{i} \leqslant e_{i}$, which includes subgroups of every possible order $n$ dividing $h=h(\mathcal{O})$. There are at most $h$ such subgroups $G$, and enumerating the cosets of $G$ takes $O\left(h \log ^{1+\epsilon} h\right)$ time, using fast composition of forms [38].

Thus we can compute a height bound $b$ for every subgroup $G$ of the form in (35) in $O\left(h^{2} \log ^{1+\epsilon} h\right)$ time. This is $O\left(|D| \log ^{1+\epsilon}|D|\right)$ under the GRH, which is dominated by our bounds for the running times of Algorithms 1 and 2. In fact, there are only $O\left(h^{\epsilon}\right)$ distinct orders that can arise among the candidate subgroups $G$, and if we restrict our attention to subgroups of the form in (9), there may be even fewer $G$ to consider. In practice, the time spent optimizing $b$ is completely negligible (and well worth the effort in any case).

As noted in the example, we only need $O(m \log |D|)$ space to compute the height bound for a given subgroup $G$, which is within the complexity bound of Proposition 2.

### 5.4. Heuristic analysis

Ignoring the minor terms in (34), the value

$$
b^{*}=\sum_{i=1}^{m} \max _{k} b_{i k}+\max _{i}\left(\sum_{k=1}^{n} b_{i k}-\max _{k} b_{i k}\right)
$$

closely approximates $b$. When $m=h$ or $n=h$ we have $b^{*}=\sum_{i, k} b_{i k}$, and in any case $b^{*}$ is never greater than this sum. As with $b_{\max }$, this yields an asymptotic bound for $b^{*}$ of $O\left(|D|^{1 / 2} \log |D| \log \log |D|\right)$, under the GRH, which then also bounds $b$. This is all that can be said in general, since $h$ could be prime.

But $h$ is rarely prime. This can occur only when $|D|$ is prime, and even then it is unlikely (by Cohen-Lenstra [17]). Let us consider the typical situation, where we are more or less free to
choose the size of $G$, at least up to a constant factor ${ }^{\dagger}$. Of course $b^{*}$ depends on the particular choice of $G$, not just its order $n$, but to simplify matters we focus on $n$, and proceed to derive a heuristic estimate for $b^{*}$ as a function of $n$ and $m=h / n$.

Let us assume that the cosets $G_{i}$ of $G$ are ordered so that $\min _{k} A_{i k}$ is increasing with $i$ (thus $G_{1}=G$ contains the identity element $\alpha_{1} \beta_{1}$ with $A_{11}=1$ ). As a heuristic, let us suppose that, on average, we have

$$
\sum_{i, k} \frac{1}{A_{i k}} \approx \sum_{i, k} \frac{1}{i+(k-1) m}=\sum_{t=1}^{h} \frac{1}{t} \approx \log h .
$$

That is, we view $\sum 1 / A_{i k}$ as an approximation to a harmonic sum in which the terms corresponding to the $i$ th coset of $G$ appear at positions $i, i+m, \ldots, i+(n-1) m$. This heuristic is based on empirical data collected during the construction of Table 3, which involved analyzing the subgroups of more than 10000 distinct class groups $\operatorname{cl}(\mathcal{O})$, with discriminants ranging from $10^{5}$ to $10^{16}$. We should emphasize that for any particular choice of $\operatorname{cl}(\mathcal{O})$ and $G$, the actual situation may deviate quite significantly from this idealized scenario, but if one averages over a large set of class groups and a large sample of their subgroups, one finds, for example, that the average rank of $\min _{k} A_{i k}$ among all the $A_{i k}$ is approximately $i$, and we note that the approximation $\sum_{i, k}\left(1 / A_{i k}\right) \approx \log h$ is correct to within an $O(\log \log h)$ factor, under the GRH. In any case, our primary justification for this heuristic is that it yields predictions that work well in practice.

Applying the heuristic yields

$$
\sum_{i=1}^{m} \max _{k} \frac{1}{A_{i k}}=\sum_{i=1}^{m} \frac{1}{i} \approx \log m
$$

and

$$
\max _{i}\left(\sum_{k=1}^{n} \frac{1}{A_{i k}}-\max _{k} \frac{1}{A_{i k}}\right)=\sum_{k=1}^{n-1} \frac{1}{m k+1} \approx \frac{\log n}{m},
$$

Table 3. Height bound optimization. (Note: Each row summarizes data collected for the first 1000 fundamental discriminants $|D| \geqslant N$. The value $b_{h}$ is the unoptimized height bound, corresponding to $|G|=h$, while $b_{n}$ is the optimized height bound, attained when $|G|=n$. Bars denote mean values.)

| $N$ | $\bar{h}$ | $\bar{n}$ | $\bar{m}$ | $\bar{b}_{h}$ | $\bar{b}_{n}$ | $\bar{b}_{h} / \bar{b}_{n}$ | $\max b_{h} / b_{n}$ |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | :---: |
| $10^{5}$ | 147 | 39 | 3.7 | 7626 | 4486 | 1.7 | 3.0 |
| $10^{6}$ | 459 | 99 | 4.6 | 28387 | 14892 | 1.9 | 3.6 |
| $10^{7}$ | 1470 | 254 | 5.8 | 105184 | 48667 | 2.2 | 4.2 |
| $10^{8}$ | 4632 | 671 | 6.9 | 377174 | 157603 | 2.4 | 4.9 |
| $10^{9}$ | 14640 | 1740 | 8.4 | 1339636 | 509688 | 2.6 | 5.5 |
| $10^{10}$ | 46434 | 4849 | 9.6 | 4709013 | 1644023 | 2.9 | 5.6 |
| $10^{11}$ | 146598 | 14777 | 9.9 | 16338099 | 5374105 | 3.0 | 6.1 |
| $10^{12}$ | 462979 | 41189 | 11.2 | 56202741 | 17182753 | 3.3 | 6.6 |
| $10^{13}$ | 1460465 | 114560 | 12.7 | 191932881 | 54720882 | 3.5 | 7.1 |
| $10^{14}$ | 4644982 | 377059 | 12.3 | 656497242 | 179083436 | 3.7 | 7.4 |
| $10^{15}$ | 14608895 | 964998 | 15.1 | 2211596515 | 560192565 | 4.0 | 9.2 |
| $10^{16}$ | 46276481 | 2695634 | 17.2 | 7462636834 | 1742205583 | 4.3 | 9.2 |

[^7]which implies that $b^{*}$ is within a constant factor of
$$
\left(\log m+\frac{\log n}{m}\right)|D|^{1 / 2}
$$

This suggests that if we wish to minimize $b^{*}$, then we should make $n$ exponentially larger than $m$. If we let $m \approx \log h$ and $n=h / m \approx h / \log h$, then we expect to have $b^{*}=$ $O\left(|D|^{1 / 2} \log \log |D|\right)$, improving our worst-case bound by a factor of $\log |D|$, and improving the average case, where $\sum_{i, k}\left(1 / A_{i k}\right) \approx \log h$, by a factor of $\log |D| / \log \log |D|$.

Using $m=\log h$ allows us to satisfy the bound (15) used to analyze Algorithm 2 whenever $\log q=O\left(|D|^{1 / 2}\right)$, which is a very mild restriction. This choice of $m$ precludes the improvement attained by Algorithm 2 under Proposition 1, since it makes $n$ too big, but it does lead to the following claim.

Heuristic Claim 1. Assuming $\log q=O\left(|D|^{c}\right)$ for some $c<1 / 2$, the average-case running time of both Algorithms 1 and 2 using $s=e_{1}$ is $O\left(|D|^{1 / 2} \log ^{2+\epsilon}|D|\right)$.

Empirical data supporting the heuristic analysis above can be found in Table 3. Each row of the table gives data for 1000 fundamental discriminants of approximately the same size. We note that for the choice of $G$ that minimizes $b$, the number of cosets $m$ is quite close to $\log h$, on average, as expected. The two rightmost columns list, respectively, the average and best-case improvement achieved by optimizing the height bound. The growth rate is consistent with the $\Omega(\log h / \log \log h)$ prediction. In practice, the actual speedup is substantially better than the height-bound improvement would suggest, for reasons that will be explained in the next section where we analyze practical computations that amply demonstrate the benefit of optimizing the height bound.

## 6. Computational results

This section presents performance data and computational results. In order to handle a wider range of discriminants, and to give the most practically relevant examples, we use class invariants derived from various modular functions to which the CRT method has been adapted. These include, among others, the Weber $f$-function, double $\eta$-quotients, and the Atkin functions $A_{N}$. We refer to $[\mathbf{2 4}, \S 3]$ for definitions of these invariants and a detailed discussion of their implementation using the CRT method. Here we briefly summarize some key properties of the class invariants we use.

### 6.1. Class invariants

Let $\mathcal{O}=\mathbf{Z}[\tau]$ be an imaginary quadratic order with discriminant $D$, for some $\tau$ in the upper halfplane. The $j$-invariant $j(\tau)$ is a root of the Hilbert class polynomial $H_{D}$ and generates the ring class field $K_{\mathcal{O}}$. Let $f(z)$ by a modular function of level $N$ related to $j(z)$ by $\Psi_{f}(f(z), j(z))=0$, where $\Psi_{f}(F, J)$ is a polynomial with integer coefficients. The value $f(\tau)$ is an algebraic integer, and when $f(\tau)$ lies in $K_{\mathcal{O}}$ we call it a class invariant ${ }^{\dagger}$. A given modular function typically yields class invariants only for a restricted subset of discriminants; for example, the primes dividing $N$ must not be inert in $\mathbf{Q}(\sqrt{D})$.

We then define the class polynomial $H_{D}[f]$ by

$$
H_{D}[f](X)=\prod_{\alpha \in \operatorname{cl}(\mathcal{O})}(X-[\alpha] f(\tau))
$$

[^8]For the functions we consider, $H_{D}[f]$ has integer coefficients, and the techniques we have developed to find a root of $H_{D} \bmod q$ apply equally well to $H_{D}[f] \bmod q$. Having found a root $f_{0}$ of $H_{D}[f]$, we may obtain a root $j_{0}$ of $H_{D}$ as a solution to $\psi(Y)=\Psi_{f}\left(f_{0}, Y\right)=0$. Since the degree of $\psi$ does not depend on $D$ or $q$, we may bound it by $O(1)$, where the implicit constant depends on $f$. Thus deriving $j_{0}$ from $f_{0}$ takes just $O\left(\log ^{2+\epsilon} q\right)$ time.

### 6.2. Heuristic height bounds

The key reason to consider alternative class invariants is that $H_{D}[f]$ may have much smaller coefficients than $H_{D}$. Let us define the height factor of $f$ as $c(f)=\operatorname{deg}_{F} \Psi_{f} / \operatorname{deg}_{J} \Psi_{f}$. Asymptotically, we have

$$
\operatorname{ht}\left(H_{D}[f]\right)=\frac{h t\left(H_{D}\right)}{c(f)}+O(1),
$$

where the constant $c(f)$ may be as large as 72 . If $b$ bounds the height of $H_{D}$, we regard $b / c(f)$ as an approximate bound on the height of $H_{D}[f]$, but add a small constant (say 256 bits) to account for the $O(1)$ term. We treat the optimized height bound $b$ computed in $\S 5$ in the same way.

This heuristic approach may, in rare cases, yield a bound that is too small. In practice this is easy to detect. The correct polynomial $V(Y)$ must split completely into linear factors in $\mathbf{F}_{q}[Y]$, and any sort of random error is extremely likely to yield a polynomial that does not. Verifying that $V(Y)$ splits into linear factors can easily be incorporated into the root-finding step at no additional cost $^{\dagger}$. If $V(Y)$ is found to be incorrect, we may then either retry with a larger height bound, or simply revert to $f=j$ and use the rigorous bound proven in $\S 5$.

In most practical applications of the CM method, we seek an elliptic curve $E / \mathbf{F}_{q}$ with prescribed order $N$, where the prime factorization of $N$ is known (or provisionally assumed). In this situation we can test whether we have constructed a suitable curve in time $O\left(\log ^{2+\epsilon} q\right)$, via [42, Lemma 6], which is negligible. Although it is usually unnecessary, one can also verify the endomorphism ring of the constructed curve, provided that we know the factorization of the integer $v$ in the norm equation $4 q=t^{2}-v^{2} D$. Using the algorithm in [11, Algorithm 2], this takes time subexponential in $\log |D|$, under heuristic assumptions, which is also negligible.

### 6.3. Implementation

Our tests were performed on a small network of quad-core AMD Phenom II 945 CPUs, each clocked at 3.0 GHz . The computation of class polynomials (or decompositions thereof) was distributed across up to 48 cores, depending on the size of the test, with essentially linear speedup, while all root-finding operations were performed on a single core. For consistency we report total CPU times, summed over all threads.

The software was implemented using the gmp [27] and zn_poly [30] libraries, with the gcc compiler [25]. Polynomial arithmetic modulo the small primes $p \in S$ was handled via zn_poly, while polynomial arithmetic modulo large primes $q$ used the cache-friendly truncated FFT approach described in [31], layered on top of the gmp library. In order to simplify the implementation, when selecting the subgroup $G$ to optimize the height bound, only subgroups of the form (9) in $\S 2.6$ were considered. Additionally, of the various space optimizations described in $\S 4.3$ that may be applied to Algorithm 2, only the changes necessary to achieve a space complexity of $O(h \log h+(m+n) \log q)$ were used (see $\S 4.3 .2$ ). A more complete implementation would improve some of the results presented here.

[^9]As noted in Remark 2, in our implementation we fixed $s=e_{1}$. This choice of $s$ worked in every large $\left(|D|>10^{6}, \log q>160\right)$ example that we tested, which included more than a million different combinations of $D$ and $G$. We conjecture that $s=e_{1}$ always works when using $j$-invariants, but note that it can fail for other class invariants in rare cases (the handful of exceptions we found all involved very small discriminants, and in each such case switching to $s=e_{2}$ worked).

### 6.4. Accelerated CM computations with Algorithm 1

We applied Algorithm 1 to several examples that have previously appeared in the literature. The examples in Table 4 are taken from [42, Table 2] where they appear as representatives of a large set of computations to construct elliptic curves suitable for pairing-based cryptography. These examples are also used in [24, Table 1] with the class invariants we use here ${ }^{\dagger}$. The first five discriminants in Table 5 originally appeared in [21, Table 1], and can also be found in [42, Table 4] and [24, Table 2]. The remaining discriminants are from [42, Table 4].

The time $T_{\text {poly }}$ listed in Tables 4-7 is the total time spent computing the polynomial $V$ and the polynomials $W_{k}$, in the case of Algorithm 1 (Steps 1-6), and the total time spent computing the polynomial $V$ and the values $w_{k}$ in the case of Algorithm 2 (steps 1-5 and 7-8), including all precomputation. The time $T_{\text {root }}$ is the time spent on root-finding operations (steps $7-8$ in Algorithm 1 and steps 6 and 9 in Algorithm 2). For the smaller examples, these are averages over 10 runs; with the large examples there is very little variance in the root-finding times.

The 'Standard' computations listed in Tables 4 and 5 correspond directly to the computations in [24], and are equivalent to running Algorithm 1 with $G=\operatorname{cl}(\mathcal{O})$. The 'Accelerated' computations used Algorithm 1 with $G$ chosen to minimize $T_{\text {tot }}$, based on heuristic formulas for $T_{\text {poly }}$ and $T_{\text {root }}$ extrapolated from empirical data. In most cases this minimizes the corresponding height bound, but not always; in the $h=100000$ example of Table 5, using $n=5000$ rather than $n=2500$ improves the ratio $b_{h} / b_{n}$ from 4.93 to 5.84 and reduces $T_{\text {poly }}$ by 10 s , but it increases $T_{\text {root }}$ by 21 s , so this subgroup was not chosen.

The $T_{\text {poly }}$ times listed in Table 4 are in each case slightly greater than the quantity $|S|\left(T_{\text {find }}+T_{\text {enum }}+T_{\text {build }}\right)$, due to time spent updating the CRT data in step 3e of Algorithm 1, which is included in $T_{\text {poly }}$. This difference is only a few percent for the values of $q$ used in these examples, but becomes more significant when $q$ is very large (see $\S 6.5$ ).

The third example in Table 4 illustrates four ways in which Algorithm 1 can reduce the time required to apply the CM method using the CRT approach.
(1) The height bound $b_{n}=50180$ is about 3 times smaller than $b_{h}=151359$, which reduces $|S|$ similarly, from 4477 to 1519 .
(2) The average time $T_{\text {find }}$ spent finding an element of $E l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$ is reduced from 42 to 28 ms , because the primes that remain in $S$ are those for which it is easier to find curves in $E l l_{\mathcal{O}}\left(\mathbf{F}_{p}\right)$.
(3) The average time $T_{\text {build }}$ spent building polynomials from their roots (or computing linear combinations) is reduced from 63 to 22 ms , because the degrees of the polynomials involved are $m=22$ and $n=2458$ rather than $h=54076$.
(4) Working with polynomials of lower degree reduces the time $T_{\text {root }}$ spent finding roots dramatically: from 171 s to 5 s .

As may be seen in Tables 4 and 5, the speedup achieved by Algorithm 1 is typically better than the height bound ratio $b_{h} / b_{n}$, for the reasons noted above. In the last example of Table 5, with discriminant $D=-170868609071$ and class number $h=1000000$, computing an optimized

[^10]height bound $b_{n}$ with $n=25000$ improves the height bound by a factor of about 7 , but $T_{\text {poly }}$ is reduced by nearly a factor of 14 and $T_{\text {root }}$ is reduced even more.

The discriminant $D=-170868609071$ also appears in [42, Table 4], which lists a time equivalent to 150 CPU days on our current test platform to compute the Hilbert class polynomial $H_{D}$ modulo a 256 bit prime $q$. Here we instead use the Weber f-function, with a height factor of 72 , and are able to further improve the height bound by a further factor

Table 4. Example CM constructions.

|  | Example 1 | Example 2 | Example 3 |
| :--- | ---: | ---: | ---: |
| Discriminant $\|D\|$ | 13569850003 | 11039933587 | 12901800539 |
| Field size $\lceil\lg q\rceil$ | 177 | 231 | 172 |
| Class number $h$ | 20203 | 11280 | 54076 |
| Presentation $\ell_{1}^{r_{1}}, \ldots, \ell_{k}^{r_{k}}$ | $7^{20203}$ | $17^{1128}, 19^{10}$ | $3^{27038}, 5^{2}$ |
| Modular function $f$ | $A_{71}$ | $A_{47}$ | $A_{71}$ |
| Height factor $c(f)$ | 36 | 24 | 36 |
| Standard |  |  |  |
| Subgroup size $\|G\|=h$ | 20203 | 11280 | 54706 |
| Height bound $b_{h}$ | 63127 | 56631 | 151939 |
| Number of primes $\|S\|$ | 1993 | 1783 | 4477 |
| $T_{\text {find }}(\mathrm{ms})$ | 48 | 110 | 42 |
| $T_{\text {enum }}(\mathrm{ms})$ | 33 | 48 | 23 |
| $T_{\text {build }}(\mathrm{ms})$ | 15 | 7 | 63 |
| $T_{\text {poly }}(\mathrm{s})$ | 197 | 295 | 597 |
| $T_{\text {root }}(\mathrm{s})$ | 54 | 171 |  |
| $T_{\text {tot }}(\mathrm{s})$ | $\mathbf{5 6}$ | $\mathbf{3 4 7}$ | $\mathbf{7 6 8}$ |
| Accelerated | $\mathbf{2 5 3}$ |  |  |
| Subgroup size $\|G\|=n$ |  | 1128 | 2458 |
| Height bound $b_{n}$ | 227 | 30957 | 50180 |
| Number of primes $\|S\|$ | 35115 | 994 | 1519 |
| $T_{\text {find }}(\mathrm{ms})$ | 115 | 105 | 28 |
| $T_{\text {enum }}(\mathrm{ms})$ | 44 | 47 | 23 |
| $T_{\text {build }}(\mathrm{ms})$ | 33 | 3 | 22 |
| $T_{\text {poly }}(\mathrm{s})$ | 6 | 155 | 118 |
| $T_{\text {root }}(\mathrm{s})$ | 95 | $\mathbf{4}$ | 5 |
| $T_{\text {tot }}(\mathrm{s})$ | 0 | $\mathbf{1 5 9}$ | $\mathbf{1 2 3}$ |

TABLE 5. CM constructions using the Weber $\mathfrak{f}$-function and $q \approx 2^{256}$.

| $\|D\|$ | $h$ | $n$ | $b_{h} / b_{n}$ | Standard |  | Accelerated |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | $T_{\text {poly }}$ | $T_{\text {root }}$ | $T_{\text {poly }}$ | $T_{\text {root }}$ |
| 6961631 | 5000 | 250 | 3.63 | 1.0 | 25 | 0.2 | 0.7 |
| 23512271 | 10000 | 250 | 3.65 | 3.9 | 58 | 0.8 | 0.7 |
| 98016239 | 20000 | 625 | 4.10 | 21 | 126 | 3.5 | 2.1 |
| 357116231 | 40000 | 625 | 4.82 | 90 | 282 | 11 | 2.2 |
| 2093236031 | 100000 | 2500 | 4.93 | 750 | 812 | 88 | 11 |
| 8364609959 | 200000 | 4000 | 6.34 | 3590 | 1805 | 301 | 16 |
| 17131564271 | 300000 | 6250 | 6.47 | 9070 | 2890 | 708 | 19 |
| 30541342079 | 400000 | 12500 | 6.31 | 16900 | 3910 | 1380 | 71 |
| 42905564831 | 500000 | 15625 | 7.11 | 28300 | 4410 | 2300 | 85 |
| 170868609071 | 1000000 | 25000 | 7.06 | 123000 | 9260 | 8840 | 159 |

of 7 using a decomposition of the class polynomial $H_{D}[\mathfrak{f}]$. We eventually obtain a root of the original polynomial $H_{D} \bmod q$, and it takes only 2.5 CPU hours to do so, an overall speedup by nearly a factor of 1500 .

### 6.5. Optimizing space with Algorithm 2

As noted in Remark 2, choosing $G$ to optimize the height bound may negate any performance advantage Algorithm 2 might have over Algorithm 1. Indeed, Algorithm 1 is usually faster, due to the larger height bound required by Algorithm 2, and the fact that Algorithm 2 repeats the enumeration step in its second stage. However when $q$ is large, Algorithm 2 may use much less space than Algorithm 1, which can actually lead to a better running time. In this scenario we use the modified form of Algorithm 2 described in $\S 4.4$, which makes the height bound increase negligible, and to optimize space we choose $G$ so that $m=h / n$ is approximately equal to but no larger than $n=|G|$.

Table 6 compares the time and space required by Algorithms 1 and 2 for a fixed discriminant $D=-300000504611$ and increasingly large primes $q \approx 2^{k}$. The class number is $h=2^{18}$, and in each case we choose $G$ so that $m=n=2^{9}$. In addition to the times $T_{\text {find }}, T_{\text {enum }}$, and $T_{\text {build }}$ listed in Table 4, we also list the average time $T_{\text {crt }}$ spent updating CRT data for each of the primes $p \in S$. This time is negligible when $q$ is of moderate to cryptographic size (say, up to 1024 bits), but when $q$ is very large, as may occur in elliptic curve primality proving [3, 36], the time Algorithm 1 spends updating its CRT data becomes quite significant ${ }^{\dagger}$.

TABLE 6. Algorithms 1 and 2 with $n=|G|=512$ and $q \approx 2^{k}$. (Note: $D=-300000504611$ with $h(D)=262144$ using $A_{71}$.)

| $\lg k$ | alg | $b$ <br> $(\mathrm{bits})$ | $\|S\|$ | $T_{\text {find }}$ <br> $(\mathrm{ms})$ | $T_{\text {enum }}$ <br> $(\mathrm{ms})$ | $T_{\text {build }}$ <br> $(\mathrm{ms})$ | $T_{\text {crt }}$ <br> $(\mathrm{ms})$ | $T_{\text {poly }}$ <br> $(\mathrm{s})$ | $T_{\text {root }}$ <br> $(\mathrm{s})$ | CRT data <br> $(\mathrm{MB})$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | ---: | ---: |
| 7 | 1 | 378315 | 10013 | 165 | 107 | 166 | 25 | 4660 | 1 | 8.2 |
|  | 2 | 378452 | 10016 | 166 | 213 | 168 | 0 | 5510 | 1 | 0.03 |
| 8 | 1 | 378315 | 10013 | 165 | 107 | 166 | 26 | 4660 | 3 | 12.6 |
|  | 2 | 378350 | 10020 | 166 | 213 | 169 | 0 | 5510 | 3 | 0.05 |
| 9 | 1 | 378315 | 10013 | 165 | 107 | 165 | 28 | 4670 | 12 | 21.0 |
|  | 2 | 378836 | 10026 | 166 | 213 | 169 | 0 | 5520 | 12 | 0.08 |
| 10 | 1 | 378315 | 10013 | 165 | 107 | 166 | 33 | 4270 | 37 | 37.7 |
|  | 2 | 379348 | 10039 | 166 | 213 | 169 | 0 | 5530 | 37 | 0.15 |
| 11 | 1 | 378315 | 10013 | 165 | 107 | 166 | 43 | 4820 | 142 | 71.3 |
|  | 2 | 380372 | 10066 | 166 | 213 | 169 | 0 | 5540 | 142 | 0.28 |
| 12 | 1 | 378315 | 10013 | 165 | 107 | 166 | 73 | 5120 | 697 | 138 |
|  | 2 | 382420 | 10119 | 166 | 213 | 169 | 0 | 5590 | 697 | 0.54 |
| 13 | 1 | 378315 | 10013 | 165 | 107 | 166 | 129 | 5690 | 3420 | 273 |
|  | 2 | 386516 | 10225 | 167 | 213 | 169 | 0 | 5630 | 3420 | 1.06 |
| 14 | 1 | 378315 | 10013 | 165 | 107 | 166 | 225 | 6700 | 16510 | 541 |
|  | 2 | 394708 | 10437 | 168 | 214 | 170 | 1 | 5810 | 16510 | 2.11 |
| 15 | 1 | 378315 | 10013 | 165 | 107 | 166 | 461 | 9100 | 81100 | 1078 |
|  | 2 | 411902 | 10859 | 170 | 214 | 171 | 2 | 6060 | 81100 | 4.21 |

[^11]One can see the two disadvantages of Algorithm 2 in Table 6 ; it requires a slightly larger $S$, and $T_{\text {enum }}$ is doubled. However Algorithm 2 needs much less space for its CRT data, and spends negligible time updating it. In our implementation both Algorithms 1 and 2 use $O(h \log |D|)$ space for the computations performed modulo each prime $p \in S$, about 10 MB in this example, but Algorithm 1 requires $O(h \log q)$ space for its CRT data, regardless of the choice of $G$, whereas Algorithm 2 only requires $O((m+n) \log q)$ space. As shown in the last two rows of Table 6, for $q \approx 2^{32768}$ Algorithm 1 uses more than 1 GB of CRT data, compared to about 4 MB for Algorithm 2, which leads to a significant time advantage for Algorithm 2. Note that the memory required by Algorithm 2 to store its CRT data is actually half the size listed in Table 6 , since with $m=n$ the CRT data is evenly split across the 2 stages and the CRT data for the first stage can be discarded before the second stage begins.

Due to its superior space complexity, Algorithm 2 is able to effectively handle a broader range of $|D|$ and $q$ than Algorithm 1. The next section gives an example of a computation with $|D| \approx 10^{15}$ and $q \approx 10^{10000}$ that is easily handled by Algorithm 2 but would be impractical to compute on our test platform using Algorithm 1, or any algorithm that requires space proportional to the size of $H_{D} \bmod q$.

### 6.6. Some large examples

We also tested Algorithms 1 and 2 with some larger discriminants, beginning with $D=$ -1000000013079299 , which has class number $h(D)=10034174$. This discriminant appears in [24], where it was used to construct an elliptic curve over a 256 bit prime field via a class invariant derived from the Atkin function $A_{71}$. As noted in [24], this set of parameters was chosen so that the level $N=71$ is ramified in $\mathbf{Q}(\sqrt{D})$, which allows us to work with the square root of the class polynomial $H_{D}\left[A_{71}\right]$, reducing both the degree and the height bound by a factor of two. The decomposition techniques described here can be applied directly to the polynomial $\sqrt{H_{D}\left[A_{71}\right]}$, allowing both Algorithms 1 and 2 to take advantage of this situation.

Table 7 gives timings for five computations that constructed elliptic curves modulo a 256 bit prime $q$ by obtaining a root of the polynomial $\sqrt{H_{D}\left[A_{71}\right]} \bmod q$. The first row corresponds to the original computation in [24]. The next two rows give timings for Algorithms 1 and 2 when the subgroup $G$ is chosen to optimize the running time of Algorithm 1, with $n=|G|=44399$. This reduced the total CPU time by nearly a factor of 5 , allowing the entire computation to be completed in less than a day of elapsed time on 48 cores. The portion of CPU time spent on root-finding was cut dramatically, from more than a day to under five minutes. This improvement is particularly helpful in a distributed implementation, as rootfinding is not as easy to parallelize as the other steps and is most conveniently performed on a single CPU.

The last two rows of Table 7 give timings for Algorithms 1 and 2 when $G$ is chosen to optimize the space used by Algorithm 2. This increases the running time by about $15 \%$, but requires

TABLE 7. CM computations with $|D|=10^{15}+13079299$ and $q \approx 2^{256}$.

| alg | $n$ | $b$ <br> $(\mathrm{bits})$ | $\|S\|$ | $T_{\text {find }}$ <br> $(\mathrm{ms})$ | $T_{\text {enum }}$ <br> $(\mathrm{ms})$ | $T_{\text {build }}$ <br> $(\mathrm{ms})$ | $T_{\text {crt }}$ <br> $(\mathrm{ms})$ | $T_{\text {poly }}$ <br> $(\mathrm{s})$ | $T_{\text {root }}$ <br> $(\mathrm{s})$ |
| :--- | ---: | ---: | :--- | ---: | ---: | ---: | ---: | ---: | ---: |
| - | - | 21533401 | 438700 | 17500 | 1580 | 25000 | 531 | 19400000 | 95600 |
| 1 | 44399 | 8315747 | 170112 | 12700 | 1580 | 9210 | 531 | 4120000 | 237 |
| 2 | 44399 | 8344202 | 150662 | 12700 | 3160 | 8350 | 5 | 4190000 | 237 |
| 1 | 3277 | 11130011 | 227504 | 13700 | 1580 | 7180 | 535 | 5260000 | 0 |
| 2 | 3277 | 11518641 | 235482 | 14400 | 3170 | 2510 | 0 | 4780000 | 0 |

less than 2 MB of CRT data, compared to about 250 MB for the original computation (and Algorithm 1). This reduced the total memory usage from around 500 MB to about 100 MB .

As noted in $\S 6.5$, the reduced space required by Algorithm 2 becomes critical for larger values of $q$. To demonstrate this, we performed a sixth computation with the discriminant $D=-1000000013079299$, this time using $q \approx 2^{33220}$. The total running time for Algorithm 2 was about 5800000 s (including root-finding), just a $20 \%$ increase over the 256 bit computation, and the size of the CRT data was about 25 MB , yielding a total memory usage under 200 MB. The 10000 -digit prime $q$ and the coefficients of the constructed curve are too large to conveniently print here, but they are available at http://math.mit.edu/~drew.

By contrast, Algorithm 1, and the algorithm of [24], requires more than 20 GB of CRT data for this example, and this data needs to be updated for every prime $p \in S$. This makes it infeasible to even attempt this computation with Algorithm 1 on our test platform, whereas Algorithm 2 was easily able to address this example.

Finally, we performed two record-setting computations, one with $h(D)>5 \cdot 10^{7}$ and the other with $|D|>10^{16}$, again using the polynomial $\sqrt{H_{D}\left[A_{71}\right]}$. First, we used the discriminant $D=-506112046263599$ with class number $h(D)=50666940$ to construct an Edwards curve of the form $x^{2}+y^{2}=1+c x^{2} y^{2}$, where

$$
c=3499565016101407566774046926671095877424725326083135202080143113943636512545
$$

over the 256 bit prime field $\mathbf{F}_{q}$ with
$q=28948022309329048855892746252171986268338819619472424415843054443714437912893$.
The trace of this curve is

$$
t=340282366920938463463374607431768266146
$$

and the group order $q+1-t$ is 4 times a prime. This computation took approximately 200 days of CPU time (about 5 days elapsed time) using Algorithm 2, which in this case was faster than Algorithm 1.

Next we used $D=-10000006055889179$ with class number $h(D)=25459680$ to construct an elliptic curve with Weierstrass equation $y^{2}=x^{3}-3 x+c$, where
$c=15325252384887882227757421748102794318349518712709487389817905929239007568605$,
over the 256 bit prime field $\mathbf{F}_{q}$ with
$q=28948022309329048855892746252171992875431396939874100252456123922623314798263$.
This curve has trace

$$
t=-340282366920938463463374607431768304979
$$

and the group order is prime. This computation took about 400 days of CPU time (under 8 days elapsed time) using Algorithm 1, which was faster than Algorithm 2 for this discriminant.

Acknowledgement. I am grateful to Andreas Enge and François Morain for providing further details of the algorithms in $[\mathbf{2 3}, \mathbf{2 8}]$ and to David Harvey for his assistance with zn_poly. I would also like to sincerely thank the anonymous referee, whose careful reading and comprehensive feedback greatly improved the clarity and rigor of this article.

## References

1. A. Agashe, K. Lauter and 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 41 (eds A. J. van der Poorten and A. Stein; American Mathematical Society, 2004) 1-17.
2. M. Agrawal, N. Kayal and N. Saxena, 'PRIMES is in P', Ann. Math. (2) 160 (2004) 781-793.
3. A. O. L. Atkin and F. Morain, 'Elliptic curves and primality proving', Math. Comp. 61 (1993) 29-68.
4. E. BACH, 'Analytic methods in the analysis and design of number-theoretic algorithms', ACM Distinguished Dissertation 1984 (MIT Press, 1985).
5. E. Bach, 'Explicit bounds for primality testing and related problems', Math. Comp. 55 (1990) no. 191, 355-380.
6. S. BaIER and L. ZhaO, 'On primes in arithmetic progressions', Int. J. Number Theory 5 (2009) no. 6, 1017-1035.
7. J. Belding, R. Bröker, A. Enge and K. Lauter, 'Computing Hilbert class polynomials', Algorithmic Number Theory Symposium-ANTS VIII, Lecture Notes in Computer Science 5011 (eds A. J. van der Poorten and A. Stein; Springer, 2008) 282-295.
8. E. R. Berlekamp, 'Factoring polynomials over large finite fields', Math. Comp. 24 (1970) no. 111, 713-735.
9. D. J. Bernstein, 'Detecting perfect powers in essentially linear time, and other studies in computational number theory', PhD Thesis, University of California at Berkeley, 1995.
10. D. J. Bernstein and J. P. Sorenson, 'Modular exponentiation via the explicit Chinese Remainder theorem', Math. Comp. 76 (2007) 443-454.
11. G. Bisson and A. V. Sutherland, 'Computing the endomorphism ring of an ordinary elliptic curve over a finite field', J. Number Theory 113 (2011) 815-831.
12. R. Bröker, 'A p-adic algorithm to compute the Hilbert class polynomial', Math. Comp. 77 (2008) 2417-2435.
13. R. Bröker, K. Lauter and A. V. Sutherland, 'Modular polynomials via isogeny volcanoes', Math. Comp. 81 (2012) 1201-1231.
14. J. BUChMANN and U. Vollmer, Binary quadratic forms: an algorithmic approach, Algorithms and Computations in Mathematics 20 (Springer, Berlin, 2007).
15. J. Chao, O. Nakamura, K. Sobataka and S. Tsujii, 'Construction of secure elliptic cryptosystems using CM tests and liftings', Advances in cryptology-ASIACRYPT'98, Lecture Notes in Computer Science 1514 (Springer, 1998) 95-109.
16. A. M. Childs, D. Jao and V. Soukharev, Constructing elliptic curve isogenies in quantum subexponential time, Preprint, 2011, http://arxiv.org/abs/1012.4019v2.
17. H. Cohen and H. W. Lenstra Jr., 'Heuristics on class groups of number fields', Number Theory, Noordwijkerhout 1983, Lecture Notes in Mathematics 1068 (Springer, 1984) 33-62.
18. J.-M. Couveignes and T. Henoce, 'Action of modular correspondences around CM points', Algorithmic Number Theory Symposium-ANTS V, Lecture Notes in Computer Science 2369 (eds C. Fieker and D. R. Kohel; Springer, 2002) 234-243.
19. D. A. Cox, Primes of the form $x^{2}+n y^{2}$ : Fermat, class field theory, and complex multiplication (John Wiley and Sons, 1989).
20. R. Crandall and C. Pomerance, Prime numbers: a computational perspective, 2nd edn (Springer, 2005).
21. A. Enge, 'The complexity of class polynomial computation via floating point approximations', Math. Comp. 78 (2009) 1089-1107.
22. A. Enge and F. Morain, 'Comparing invariants for class fields of imaginary quadratic fields', Algorithmic Number Theory Symposium-ANTS V, Lecture Notes in Computer Science 2369 (eds C. Fieker and D. R. Kohel; Springer, 2002) 252-266.
23. A. Enge and F. Morain, 'Fast decomposition of polynomials with known Galois group', Applied algebra, algebraic algorithms, and error correcting codes - 2003, Lecture Notes in Computer Science 2643 (Springer, 2003) 254-264.
24. A. Enge and A. V. Sutherland, 'Class invariants for the CRT method', Algorithmic Number Theory Symposium-ANTS IX, Lecture Notes in Computer Science 6197 (eds G. Hanrot, F. Morain and E. Thomé; Springer, 2010) 142-156.
25. Free software foundation, 'GNU compiler collection', version 4.4.3, 2010, available at http://gcc.gnu.org/.
26. A. Gee and P. Stevenhagen, 'Generating class fields with Shimura reciprocity', Algorithmic Number Theory Symposium-ANTS III, Lecture Notes in Computer Science 1423 (Springer, 1998) 442-453.
27. T. GRanlund et al., GNU multiple precision arithmetic library, September 2010, version 5.0.1, available at http://gmplib.org/.
28. G. Hanrot and F. Morain, 'Solvability by radicals from an algorithmic point of view', International Conference on Symbolic and Algebraic Computation-ISSAC 2001 (ACM, 2001) 175-182.
29. G. H. Hardy and E. M. Wright, An introduction to the theory of numbers, 5th edn (Oxford Science Publications, 1979).
30. D. HARVEY, zn_poly: a library for polynomial arithmetic, version $0.9,2008$, http://cims.nyu.edu/~harvey/zn_poly.
31. D. Harvey, 'A cache-friendly truncated FFT', Theoret. Comput. Sci. 410 (2009) 2649-2658.
32. S. Ionica and A. Joux, 'Pairing the volcano', Algorithmic Number Theory Symposium-ANTS IX, Lecture Notes in Computer Science 6197 (eds G. Hanrot, F. Morain and E. Thomé; Springer, 2010) 201-218.
33. J. C. Lagarias and A. M. Odlyzko, 'Effective versions of the Chebotarev density theorem', Algebraic number fields: L-functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975) (Academic Press, 1977) 409-464.
34. S. LaNG, Elliptic functions, 2nd edn (Springer, 1987).
35. J. E. Littlewood, 'On the class-number of the corpus $P(\sqrt{-k})$ ', Proc. Lond. Math. Soc. 27 (1928) 358-372.
36. F. Morain, 'Primality proving using elliptic curves: an update', Algorithmic Number Theory SymposiumANTS III, Lecture Notes in Computer Science 1423 (Springer, 1998) 111-127.
37. K. Rubin and A. Silverberg, 'Choosing the correct elliptic curve in the CM method', Math. Comp. 79 (2010) 545-561.
38. A. SChÖNHAGE, 'Fast reduction and composition of binary quadratic forms', International Symposium on Symbolic and Algebraic Computation-ISSAC'91 (ed. S. M. Watt; ACM, 1991) 128-133.
39. A. Schönhage and V. Strassen, 'Schnelle Multiplikation großer zahlen', Computing 7 (1971) 281-292.
40. J.-P. Serre, 'Complex multiplication', Algebraic number theory (eds J.W.S. Cassels and A. Fröhlich; Academic Press, 1967).
41. A. V. Sutherland, 'Order computations in generic groups', PhD Thesis, MIT, 2007, http://groups.csail.mit.edu/cis/theses/sutherland-phd.pdf.
42. A. V. Sutherland, 'Computing Hilbert class polynomials with the Chinese remainder theorem', Math. Computation 80 (2011) 501-538.
43. A. V. Sutherland, 'Structure computation and discrete logarithms in finite abelian p-groups', Math. Comp. 80 (2011) 477-500.
44. B. Leendert and van der Waerden, Algebra, vol. I (Springer, 1991). Originally published in German as Moderne algebra in 1930-1931.
45. J. von zur Gathen and J. Gerhard, Modern computer algebra, 2nd edn (Cambridge University Press, 2003).
46. H. Weber, Lehrbuch der algebra, 3rd edn, vol. III (Chelsea, 1961).

Andrew V. Sutherland<br>Department of Mathematics<br>Massachusetts Institute of Technology<br>77 Massachusetts Avenue, Cambridge, MA 02139<br>USA

drew@math.mit.edu


[^0]:    ${ }^{\dagger}$ We assume throughout that fast probabilistic methods are used to find roots of polynomials over finite fields.

[^1]:    ${ }^{\dagger}$ The symmetry of the $\ell_{i}$ and $r_{i}$ in this small example is entirely coincidental.

[^2]:    ${ }^{\dagger}$ Our bounds count bit operations and hold for all constants $\epsilon>0$ (and often for $\epsilon=o(1)$ ).

[^3]:    ${ }^{\dagger}$ The $\theta_{i k}$ are integers of $G$ 's fixed field $L \subset K_{\mathcal{O}}$. They do not all lie in $\mathbf{Q}$ unless $G=\operatorname{cl}(\mathcal{O})$.
    ${ }^{\ddagger}$ To optimize space this step may need to be performed in an amortized fashion, see § 4.3.

[^4]:    ${ }^{\dagger}$ It is possible to do this without explicitly computing $V \bmod p$, see §4.1.

[^5]:    ${ }^{\dagger}$ Note that the exponent $\operatorname{lcm}_{\alpha \in S}|\alpha|$ of $\mathrm{cl}(\mathcal{O})$ can be computed in time $O\left(h^{1 / 2+\epsilon}\right)$; see [41].

[^6]:    ${ }^{\dagger}$ This can be improved to $O\left(h^{1 / 2+\epsilon}\right)$ time and $O\left(h^{1 / 2-\epsilon}\right)$ space using a birthday-paradox approach with a time/space trade-off, but we don't need to do this.

[^7]:    ${ }^{\dagger}$ Under the random bisection model, a random integer $N$ in some large interval will have prime-power factors whose logarithms approximate a geometric progression [4]. One then has divisors of $N$ in most intervals of the form $[M, c M] \subset[1, N]$, for a suitable constant $c$. We heuristically assume that the same applies to $h$.

[^8]:    ${ }^{\dagger}$ We do not require $f(\tau)$ to generate $K_{\mathcal{O}}$; we can obtain a generator as a root of $\Psi_{f}(f(\tau), Y)$.

[^9]:    ${ }^{\dagger}$ The first step of the standard root-finding procedure computes the polynomial $\operatorname{gcd}\left(Y^{q}-Y, V(Y)\right)$ whose degree is the number of distinct roots of $V$; duplicate roots can be accounted for by taking gcds with derivatives of $V$.

[^10]:    ${ }^{\dagger}$ The timings listed here for the standard computations are slightly better (about $5 \%$ ) than those in [24] due to a more recent version of gmp.

[^11]:    ${ }^{\dagger}$ As discussed in $[42, \S 6.3]$, we should eventually transition from the explicit CRT to a standard CRT approach as $q$ grows, but here $\lg q$ is still much smaller than the height bound $b$.

