



HAL
open science

Improved CRT Algorithm for Class Polynomials in Genus 2

Kristin Lauter, Damien Robert

► **To cite this version:**

Kristin Lauter, Damien Robert. Improved CRT Algorithm for Class Polynomials in Genus 2. ANTS X - Algorithmic Number Theory 2012, Jul 2012, San Diego, United States. pp.437-461, 10.2140/obs.2013.1.437. hal-00734450v2

HAL Id: hal-00734450

<https://hal.science/hal-00734450v2>

Submitted on 17 Apr 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

IMPROVED CRT ALGORITHM FOR CLASS POLYNOMIALS IN GENUS 2

KRISTIN E. LAUTER AND DAMIEN ROBERT

ABSTRACT. We present a generalization to genus 2 of the probabilistic algorithm of Sutherland for computing Hilbert class polynomials. The improvement over the Bröker-Gruenewald-Lauter algorithm for the genus 2 case is that we do not need to find a curve in the isogeny class whose endomorphism ring is the maximal order; rather, we present a probabilistic algorithm for “going up” to a maximal curve (a curve with maximal endomorphism ring), once we find any curve in the right isogeny class. Then we use the structure of the Shimura class group and the computation of (ℓ, ℓ) -isogenies to compute all isogenous maximal curves from an initial one.

This is an extended version of the article published at ANTS X.

1. INTRODUCTION

Cryptographic solutions to provide privacy and security for sensitive transactions depend on using a mathematical group in which the discrete logarithm problem is hard. For example, digital signature schemes or a Diffie-Hellman key exchange may be based on the difficulty of solving the discrete logarithm problem in the group of points on the Jacobian of a genus-2 curve. For this problem to be hard we must ensure that we can choose genus-2 curves over finite fields whose Jacobian have an almost-prime number of points.

One approach to this problem is to construct curves whose Jacobians have a given order using the method of complex multiplication (CM). The CM method works by computing invariants of the curve and then reconstructing the curve using the Mestre-Cardona-Quer [29, 7] algorithm. Invariants are computed by constructing their minimal polynomials, called *Igusa class polynomials*. Computing the invariants is computationally intensive, and there are three known methods for constructing Igusa class polynomials:

- (1) the complex analytic method [35, 39, 40, 36];
- (2) the Chinese remainder theorem method (CRT) [15, 17, 5]; and
- (3) the p -adic lifting method [19, 8, 9].

Currently, the CRT method in genus 2 remains by far the slowest of these three methods, as measured on the small examples that have been computed to date; but the history of the evolution of these three methods in genus 1 gives some hope that the CRT method may be asymptotically competitive with the others. In genus 1, the (explicit) CRT method now holds the record for the best proven bounds on time and space complexity (under GRH), as well as for the size of the largest examples that have been computed [37, 16]. In this paper, we propose numerous improvements to

2000 *Mathematics Subject Classification*. Primary 14K22; Secondary 11Y40, 11Y16, 11G15, 14K02.

Key words and phrases. class field polynomials, CRT, hyperelliptic curve cryptography, isogenies.

the CRT method for computing genus-2 curves, paralleling improvements made by Sutherland [37] to the CRT method in genus 1.

The CRT method works by computing class polynomials modulo many small primes, and then reconstructing the polynomials with rational coefficients (or modulo a much larger prime number) via the Chinese remainder theorem (respectively, the explicit CRT). The CRT method for computing class polynomials in genus 2 was proposed by Eisenträger and Lauter [15]; they gave sufficient conditions on the CRT primes to ensure correctness and included an algorithm for computing endomorphism rings for ordinary Jacobians of genus-2 curves, generalizing Kohel’s algorithm for genus-1 curves. For each small CRT prime p , the algorithm loops through all p^3 possible triples of Igusa invariants of curves, reconstructing the curve and testing for each curve whether it is in the desired isogeny class and whether its endomorphism ring is maximal. The algorithm for computing endomorphism rings from [15] was replaced by a much more efficient probabilistic algorithm in [17], where a number of examples were given for running times of the computations modulo small CRT primes. Bröker, Gruenewald, and Lauter [5] introduced the idea of using computable $(3, 3)$ -isogenies to find other curves in the isogeny class once an initial curve was found, but still searched until finding a curve whose Jacobian has endomorphism ring equal to a maximal order (a *maximal* curve). Another improvement described in [5] was a method to construct other maximal curves using $(3, 3)$ -isogenies once an initial maximal curve is found.

In this paper we present a generalization to genus 2 of the probabilistic Algorithm 1 in Sutherland [37]. The improvement over the genus-2 algorithm presented in [5] is that we do not need to find a maximal curve in the isogeny class; instead, we present a probabilistic algorithm for “going up” to a maximal curve once we find *any* curve in the right isogeny class. Then we use the structure of the Shimura class group and the computation of (ℓ, ℓ) -isogenies to compute all isogenous maximal curves from an initial one. Although we cannot prove that the going-up algorithm succeeds with any fixed probability, it works well in practice, and heuristically it improves the running time of the genus-2 CRT method from p^3 per prime p to $p^{3/2}$ per prime p .

Let K denote a primitive quartic CM field, with real quadratic subfield K^+ and ring of integers \mathcal{O}_K . Let Φ denote a CM-type of K and let K_Φ denote the reflex CM field. Let TN_Φ denote the type norm associated to the CM-type Φ . Informally, the algorithm is as follows; the individual steps will be explained in subsequent sections.

Algorithm 1.

Input: A primitive quartic CM field K with a CM-type Φ , and a collection of CRT primes P_K for K .

Output: Igusa class polynomials $H_i(x)$, $i = 1, 2, 3$, either in $\mathbb{Q}[x]$ or reduced modulo a prime q .

1. Loop through CRT primes $p \in P_K$:
 - (a) Enumerate hyperelliptic curves C of genus 2 over \mathbb{F}_p until a curve in the right \mathbb{F}_p -isogeny class (up to a quadratic twist) is found.
 - (b) Try to go up to a maximal curve from C ; if this step fails, go back to Step 1(a).
 - (c) From a maximal curve C , compute all other maximal curves.
 - (d) Reconstruct the class polynomials $H_i(x)$ modulo p from the Igusa invariants of the set of maximal curves.

2. Recover $H_i(x)$, $i = 1, 2, 3$ in $\mathbb{Q}[x]$ or modulo q using the (explicit) CRT method once we have computed $H_i(x)$ modulo p for enough primes p .

For the dihedral case, one new aspect of our algorithm is that we extend to the CRT setting the idea of computing the class polynomials associated to only one fixed CM-type Φ for K [36, §III.3]. When K is cyclic, this makes no difference, since all isomorphism classes of abelian surfaces with CM by K arise from one CM-type; but when K is dihedral, two CM-types are needed to find all isomorphism classes of CM abelian surfaces. All three previous versions of the CRT algorithm [15, 17, 5] compute the class polynomials classifying all abelian surfaces with CM by \mathcal{O}_K (with either of the two possible CM-types in the dihedral case). The advantage of our approach is that it computes only a factor of half the degree of the whole class polynomial. The drawback of this approach is that in the dihedral case, each factor of the class polynomials is defined over $\mathcal{O}_{K_\Phi^+}$ rather than over \mathbb{Z} . So once we compute the class polynomials modulo \mathfrak{p} as polynomials in $\mathcal{O}_{K_\Phi^+}/\mathfrak{p}$, the CRT step must be performed in $\mathcal{O}_{K_\Phi^+}$.

A CRT prime $\mathfrak{p} \subset \mathcal{O}_{K_\Phi^+}$ is a prime such that all abelian surfaces over \mathbb{C} with CM by (\mathcal{O}_K, Φ) have good reduction modulo \mathfrak{p} . By [34, §III.13], \mathfrak{p} is a CRT prime for the CM-type Φ if and only if there exists an unramified prime \mathfrak{q} in \mathcal{O}_{K_Φ} of degree 1 above \mathfrak{p} of principal type norm (π) with $\pi\bar{\pi} = N_{K/\mathbb{Q}}(\mathfrak{q})$; in particular, this implies that \mathfrak{q} is totally split in the class field corresponding to the abelian surfaces with CM by (\mathcal{O}_K, Φ) . By [20, §3], these surfaces have good reduction modulo \mathfrak{p} , and by a theorem of Tate the isogeny class of the reductions modulo \mathfrak{p} is determined by the characteristic polynomial of $\pm\pi$, at least in the case where $\mathcal{O}_K^* = \{\pm 1\}$. For reasons of efficiency, we will work with CRT primes \mathfrak{p} that are unramified of degree one over $p = \mathfrak{p} \cap \mathbb{Z}$. By [20], the reduction to \mathbb{F}_p of the abelian surfaces with CM by (\mathcal{O}_K, Φ) will then be ordinary. We then make the slight abuse of notation of calling p a CRT prime when there is a CRT prime \mathfrak{p} above it. Note another advantage of restricting to one CM-type: To use p for both CM-types, p needs to split completely into $p = \mathfrak{p}_1\mathfrak{p}_2$ such that both \mathfrak{p}_1 and \mathfrak{p}_2 are CRT primes, and there are fewer p which satisfy this stronger requirement.

In addition to the two main contributions of the paper — the going-up algorithm to find maximal curves, and an improvement to the algorithm to compute maximal curves from maximal curves — we also give improvements to every step of the CRT algorithm. Here we give a brief outline of the paper and a summary of those improvements.

Step 1b of the algorithm (the “going-up” part) is explained in Section 5. We first explain in Section 4 how to compute if a curve is maximal, since this is used in the going-up algorithm. We present some significant improvements over the algorithm from [17]. Step 1c (finding all other maximal curves from one maximal curve) is explained in Section 6. The running times of Steps 1a – c of the algorithm depend greatly on the CRT prime chosen (for instance the size of the isogeny class can vary). Since we have latitude in the choice of CRT primes, we explain our algorithm for sieving them in Section 7.

As for Step 1d, once all maximal abelian surfaces with CM by K are found for a given prime p , it is easy to compute the associated class polynomials modulo p . The class polynomials depend on the choice of Igusa invariants, and we use the invariants recommended in [36, Appendix 3] which give smaller coefficients than those used

in [39, 40, 20]. For the dihedral case the class polynomials must be reconstructed over $\mathcal{O}_{K_\Phi^+}$, and we give more details about this step in Section 8.

As for Step 2, once all maximal abelian varieties with CM by K are found for a given prime p , it is easy to compute the associated class polynomials modulo p . The class polynomials depend on the choice of Igusa invariants, and we use the invariants recommended in [36, Appendix 3] which give smaller coefficients than those used in [39, 40, 20]. For the dihedral case the class polynomials must be reconstructed over $\mathcal{O}_{K_\Phi^+}$, and we give more details about this step in Section 8.

Section 9 gives a complexity analysis, and explains how each improvement affects the final complexity. The final complexity bound, while still not quasilinear, is a significant improvement compared to [5]. Finally, examples demonstrating significantly improved running times are given in Section 10.

This paper is the long version of the paper which appeared at ANTS X, and it differs from the short version in the following ways. This paper contains extra sections detailing the definition of a CRT prime (Section 2), new methods for efficiently finding a curve in the right isogeny class (Section 3), and strategies for sieving CRT primes (Section 7), along with providing additional examples in Section 10.

2. CONDITIONS ON CRT PRIMES

The CRT algorithm computes the Igusa class polynomials modulo primes p , for primes satisfying certain conditions: namely that the interpolation polynomials computed from the Igusa invariants of the abelian varieties over \mathbb{F}_p with CM by \mathcal{O}_K are the reductions modulo p of the Igusa class polynomials. As proved in [15], this is satisfied for primes p such that all abelian varieties with CM by \mathcal{O}_K have ordinary reduction and are defined over \mathbb{F}_p . We call such a prime p a CRT prime (for the field K).

Originally, in [15], the condition given was that CRT primes p should split completely into principal ideals in K_Φ , the reflex field of K . This strong condition implied three consequences: since the prime splits completely in H_{K_Φ} , the Hilbert class field of K_Φ , all abelian surfaces with CM by K are defined over \mathbb{F}_p (not just the ones with CM by Φ); since the prime splits completely in K , the reduction of the abelian surfaces with CM by K is ordinary and the relative norm equation has a solution.

However, this condition is not always necessary, because the CM points in the moduli space may be defined over $K_\Phi(j_1(A))$, the subfield of H_{K_Φ} generated by the Igusa invariants of A , an abelian surface with CM by K .

Moreover, as noted the introduction, in practice we only want to compute the Igusa class polynomials relative to a CM type Φ (which gives a factor of the full Igusa class polynomials in the dihedral case). For efficiency reasons, we also want the reductions to be ordinary. In this case the class polynomials are defined over K_Φ^+ , and we are then looking for CRT primes $\mathfrak{p} \in \mathcal{O}_{K_\Phi^+}$. We call such a prime \mathfrak{p} a CRT prime if all abelian varieties with CM by (\mathcal{O}_K, Φ) have good ordinary reduction to $\mathcal{O}_{K_\Phi^+}/\mathfrak{p}\mathcal{O}_{K_\Phi^+}$. The results from [20] imply that $p = \mathbb{Z} \cap \mathfrak{p}$ is totally split in $\mathcal{O}_{K_\Phi^+}$. In particular, the interpolation polynomials computed over $\mathcal{O}_{K_\Phi^+}/\mathfrak{p}\mathcal{O}_{K_\Phi^+}$ are actually defined over \mathbb{F}_p and by abuse of notation we call p a CRT prime.

We have the following characterization of CRT primes:

Theorem 2. *Let Φ be a CM type and $p \in \mathbb{Z}$ be a prime number. If there exists an unramified prime \mathfrak{p} in \mathcal{O}_{K_Φ} of degree 1 above p of principal typenorm (π) such that $\pi\bar{\pi} = p$, then p is a CRT prime for the CM type Φ .*

If moreover p splits completely in \mathcal{O}_{K_Φ} into primes with principal typenorm, then p is a CRT prime for K .

Proof. Let A/\mathbb{C} be an abelian variety of CM type (K, Φ) . There exists a model of A defined over the subfield H_0 of H_{K_Φ} generated by the Igusa invariants of A . By [34], the condition on \mathfrak{p} implies that it splits completely in this subfield (more precisely, the typenorm of \mathfrak{p} gives the action of the Frobenius in terms of the Shimura class group; the abelian variety $A_{\mathfrak{p}}$ descends to \mathbb{F}_p if and only if this element is trivial). If we fix one of the primes \mathfrak{p}_0 above \mathfrak{p} , then by [20], we have that $A_{\mathfrak{p}_0}$ has ordinary reduction defined over \mathbb{F}_p . All such abelian varieties with CM by Φ have ordinary reduction over \mathbb{F}_p , so p is a CRT prime.

For the second assertion, we have by the hypothesis that p splits completely as $p = \mathfrak{p}_1\mathfrak{p}_2$ in $\mathcal{O}_{K_\Phi^+}$ (which does not depend of the CM type). By the preceding paragraph, since all primes in \mathcal{O}_{K_Φ} above \mathfrak{p}_1 satisfy the condition of the first part, p is a CRT prime for the CM type Φ . This concludes the cyclic case. Now for the dihedral case, let Φ' be the other CM type. We will see in Section 2.2, that the reductions of the abelian varieties with CM by Φ modulo primes in \mathcal{O}_{K_Φ} above \mathfrak{p}_2 correspond exactly to the reductions of the abelian varieties with CM by Φ' modulo primes in $\mathcal{O}_{K_{\Phi'}}$ above \mathfrak{p}_1 . So p is also a CRT prime for Φ' . \square

2.1. The cyclic case. In the cyclic case, there is just one CM type Φ , so the Igusa class polynomials for this CM type are the same as the full Igusa class polynomials for CM by \mathcal{O}_K .

A prime p is then a CRT prime if it splits completely in \mathcal{O}_K as $p = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3\mathfrak{p}_4$ and \mathfrak{p}_1 has a principal typenorm. (Since the other primes are all conjugate to this one, they all also have principal typenorms, and give the same isogeny class.)

2.2. The dihedral case. For the dihedral case, a prime number p is a CRT prime if p splits completely as q_1q_2 in K_Φ^+ , such that q_1 splits completely in K_Φ into primes with principal type norm (π) and $\pi\bar{\pi} = p$. A word of warning: when we compute the resulting interpolating Igusa polynomials over \mathbb{F}_p , they are the reduction of the Igusa class polynomials modulo q_1 and the CRT step will then be computed over $\mathcal{O}_{K_\Phi^+}$, not over \mathbb{Z} as in the cyclic case. We note that if q_2 also splits completely into prime ideals with principal type norm, we could use q_2 in the CRT rather than q_1 . We use the sieving algorithms of Section 7 to select the one most practical for the computation.

Now we explain how to adapt the algorithm to compute directly the Igusa class polynomial with CM by \mathcal{O}_K (not just one fixed CM type). In the dihedral case there are two distinct CM types Φ_1 and Φ_2 (conjugate by an element of the Galois group). Let A_1 (resp. A_2) be an abelian variety with complex multiplication by Φ_1 (resp. Φ_2) and p a prime that splits as q_1q_2 in K_Φ^+ (this field does not depend on the choice of the CM type). Note that since the two CM types are conjugate, this gives that $(A_2)_{q_1}$ ¹ is isogenous to $(A_1)_{q_2}$. Thus both CM types can be handled at

¹Where by abuse of notation we denote by $(A_2)_{q_1}$ the reduction of A_2 modulo a prime above q_1 in the moduli field where A_2 is defined.

once by fixing once and for all the reflex field (corresponding to one fixed CM type Φ) and looking at both primes q_1 and q_2 .

Thus in this case, there are up to two isogeny classes (up to twists) corresponding to maximal curves over \mathbb{F}_p , one for each CM type. So to compute the class polynomial corresponding to all curves with CM by \mathcal{O}_K , we have to look for CRT primes p that split completely in K_Φ^+ ($p = q_1 q_2$) and such that both q_1 and q_2 split completely, and the primes above them have principal type norm. (The corresponding typenorm of the primes above q_1 and q_2 will then give the action of Frobenius on both isogeny classes). So that gives fewer CRT primes to work with. Moreover, in Section 7 we explain how to sieve the CRT primes according to the size of the isogeny class and the ease of endomorphism ring computation and going up. Since we have two isogeny classes here, it is also twice as hard to find good CRT primes p . Thus, there are many drawbacks to working over \mathbb{Z} in the dihedral case: we can't distinguish between the two CM types (this would involve distinguishing reductions modulo q_1 and q_2 , which requires working over $\mathcal{O}_{K_\Phi^+}$ where the q_i are defined), so we have to compute the full class polynomials with CM by \mathcal{O}_K . This class polynomial has bigger coefficients, there are fewer CRT primes, and it is harder to find good CRT primes among them.

2.3. Finding the isogeny class. For a given CRT prime p , the algorithm must find all abelian varieties over \mathbb{F}_p with CM by (\mathcal{O}_K, Φ) , i.e. genus 2 curves H such that $\text{End}(\text{Jac}(H)) \simeq \mathcal{O}_K$. These curves lie in several possible isogeny classes, see [15] for the possible number of group orders in the ordinary case. So the algorithm first tries to find a curve in the corresponding isogeny class, (i.e. such that $\text{End}(\text{Jac}(H)) \otimes \mathbb{Q} \simeq K$). For that we need a characterization of this isogeny class. By a theorem of Tate, the isogeny class is characterized by the zeta function of a curve in it, or equivalently by the characteristic polynomial of the action of the Frobenius $\pi \in K$. So we need to recover $\pi \in K$ and compute its characteristic polynomial.

If $A_{\mathfrak{p}}/\mathbb{F}_p$ is the reduction of an abelian variety with complex multiplication by Φ , then [34] shows that the typenorm of $\mathfrak{p} \cap K_\Phi$ gives the action of the Frobenius. This gives the following algorithm:

Algorithm 3. Characterizing the isogeny class.

Input: An unramified prime \mathfrak{p} (above the prime $p \in \mathbb{Z}$) in \mathcal{O}_{K_Φ} of degree 1 of principal typenorm.

Output: The characteristic polynomial of the Frobenius corresponding to the reduction $A_{\mathfrak{p}}$ of an abelian variety A with CM by (\mathcal{O}_K, Φ) .

1. Compute $TN(\mathfrak{p}) = (\alpha)$.
2. Compute the fundamental unit ξ of K^+ .
3. Choose an embedding $K \hookrightarrow \mathbb{C}$ and let $|\cdot|$ be the corresponding absolute value.
4. Compute $n = \log(p/|\alpha|^2) / \log(|\xi|^2) \in \mathbb{Z}$.
5. Compute $\pi = \alpha \xi^n$.
6. Return the characteristic polynomial of π .

Remark 4. The ideal generated by α has relative norm p (by definition of the type norm), but we need to find a generator π with complex absolute value equal to \sqrt{p} . We adjust α by multiplying by some power of the fundamental unit, ξ in K^+ . In our

algorithm we improve Algorithm 2.1, Step 4, from [17] by computing the complex logarithms of α and ξ , and then directly adjusting α by the correct power of ξ .

Note that $-\pi$ is also of norm \sqrt{p} , so there are actually two isogeny classes corresponding to Jacobians J with $\text{End}(J) \otimes \mathbb{Q} \simeq K$, but the one corresponding to $-\pi$ is given by a quadratic twist of the first class.

3. SEARCHING FOR A CURVE IN THE ISOGENY CLASS

In this section we introduce several ways to speed up the computation when searching for a curve in the right isogeny class modulo a given CRT prime. The standard approach in previous implementations has been to loop through all p^3 triples of Igusa invariants, generating a genus 2 curve for each triple using Mestre's algorithm, and then checking whether it is in the right isogeny class by either counting points on the curve to determine the zeta function or by generating random points on the Jacobian and checking that the points are killed by the group order (corresponding to π or $-\pi$).

But in our new algorithm, we don't need to loop over all the curves, but only to find a curve in the right isogeny class. For each CRT prime, the probability for a random curve to be in a fixed isogeny class is given by the size of the isogeny class divided by the number of isomorphism classes of curves. We can efficiently sample random curves by taking random monic sextic polynomials with coefficients in \mathbb{F}_p . Since there are roughly p^3 isomorphism classes of curves and there are p^6 sextics, the probability that two randomly chosen sextics define isomorphic curves is $1/p^3$. The advantage to this approach is that randomly generating sextics is a much more efficient way to generate curves than generating the curve equation from triples of Igusa invariants by applying Mestre's algorithm for each triple. This gives the following algorithm:

Algorithm 5. Finding a curve in the isogeny class corresponding to $\pm\pi$.

If the characteristic polynomial of π is $X^4 + t_0X^3 + t_1X^2 + t_2X + p^2$ then a curve H is in the isogeny class if and only if $\#H(\mathbb{F}_p) = M_1 := p + 1 + t_0$ and $\#\text{Jac}(H) = N_1 := 1 + t_0 + t_1 + t_2 + p^2$. Likewise we define M_2 to be the cardinality of a curve in the isogeny class corresponding to the twist $-\pi$ and N_2 the cardinality of its Jacobian.

- (1) Take a random hyperelliptic curve $H : y^2 = a_6x^6 + a_5x^5 + a_4x^4 + a_3x^3 + a_2x^2 + a_1x + a_0$.
- (2) Test whether $N_iP = 0$ ($i = 1, 2$) for several random elements $P \in \text{Jac}(H)$. (We adapt dynamically the number of random elements P we take for this step depending on the size of p .)
- (3) If Step 2 succeeds for all P , then test if $\#H = M_i$ and $\#J = N_i$ (for $i = 1$ or for $i = 2$) and return H in case of success. Go back to step 1 otherwise.

Remark 6. Note though, that because the height of the first Igusa class polynomials is smaller than the height of the second and third class polynomials (with the invariants we are using), it can be computed with fewer CRT primes. Once the first polynomial is known, for subsequent CRT primes it may be better to revert to the method of searching for curves using the Igusa invariants: since the $\deg H_1$ possible values of the first invariant are known, we only need to loop through $p^2 \times \deg H_1$ Igusa triples. This is to be compared with the number of curves we expect to loop through before finding a curve in the right isogeny class.

Depending on the CRT field K , there may be more efficient ways to sample random curves than by taking random monic sextic polynomials.

3.1. Rosenhain representation. If the 2-torsion of the Jacobian of any maximal curve in the isogeny class is rational, we can search for curves in Rosenhain form:

Lemma 7. *The following assertions are equivalent:*

- (1) *Every maximal curve can be put in Rosenhain normal form*

$$y^2 = x(x-1)(x-\lambda)(x-\mu)(x-\nu)$$

where $\lambda, \mu, \nu \in \mathbb{F}_p$.

- (2) *The 2-torsion of a Jacobian of a maximal curve is rational over \mathbb{F}_p .*
 (3) $\frac{\pi-1}{2} \in \mathcal{O}_K$.

Proof. It is well known that a hyperelliptic curve of genus 2 can be put into Rosenhain form if and only if the 2-torsion of its Jacobian is rational. So the first assertion implies the second. Now we also have that the 2-torsion on a Jacobian is rational if and only if the Frobenius π acts trivially on it, if and only if $\text{Ker}(\pi-1)$ contains the 2-torsion, if and only if $\frac{\pi-1}{2}$ is an endomorphism. So if one maximal Jacobian A has a rational 2-torsion, then $\frac{\pi-1}{2} \in \text{End}(A) = \mathcal{O}_K$, so the second assertion implies the third. Finally, if $\frac{\pi-1}{2} \in \mathcal{O}_K$, then every maximal Jacobian has $\frac{\pi-1}{2}$ as an endomorphism, so can be put into Rosenhain normal form. \square

It is much easier and faster to loop through curves written in Rosenhain form directly. This approach is also better in the sense that it only loops through curves with rational 2-torsion, thus avoiding generating and computing on curves which could not be maximal curves in our isogeny class. Also, a factor of 6 is saved because the ordering of the roots λ, μ, ν does not matter, and there are 6 ways to permute those 3 roots.

Remark 8. One trade-off we consider in applying the Rosenhain method is the following: while we take random curves over a space one sixth as large as the space of all hyperelliptic curves of genus 2, we can only find curves in the isogeny class in Rosenhain form. We could estimate the size of this intersection by checking whether the element $\frac{\pi-1}{2}$ is in the suborder $R \subset \mathcal{O}_K$ for the various $R \supset \mathbb{Z}[\pi, \bar{\pi}]$ and computing the number of curves with endomorphism ring R , but this is expensive. In practice, we only apply this method when the condition holds for $R = \mathbb{Z}[\pi, \bar{\pi}]$ (and hence for all curves in the isogeny class), which is the case for instance when 2 does not divide the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$.

3.2. Real multiplication. For a given quartic CM field K , write the real quadratic subfield $K^+ = \mathbb{Q}(\sqrt{d})$. For a given CRT prime p , any curve over \mathbb{F}_p with CM by K must also have real multiplication by K^+ . Such curves correspond to CM points associated to K on the Hilbert moduli space associated to K^+ . Such curves are determined by pairs of Gundlach invariants on the Hilbert moduli space instead of triples of Igusa invariants. Thus the algorithm to search for these curves can be improved by looping through Gundlach invariants instead of Igusa invariants. This approach is more efficient because there are only p^2 pairs to search through instead of p^3 triples. Formulas for Gundlach invariants and a method for generating genus 2 curves from Gundlach invariants were given in [26]. Those formulas can be used directly in the search step for each CRT prime.

Here the same trade off has to be considered as the Rosenhain method: if $R \cap K^+ \not\subset \mathcal{O}_{K^+}$ for various orders R such that $\mathbb{Z}[\pi, \bar{\pi}] \subset R \subset \mathcal{O}_K$, we miss all the curves whose Jacobian have endomorphism ring R .

4. CHECKING IF THE ENDOMORPHISM RING IS MAXIMAL

We recall the algorithm described in [15], and describe some improvements. The ideas for computing the endomorphism ring will be used in the going up phase.

4.1. The algorithm of Eisenträger, Freeman and Lauter. Let A/\mathbb{F}_p be an ordinary abelian variety of dimension 2 with CM by K . Let $O = \text{End}(A)$. We know that $\mathbb{Z}[\pi] \subset \mathbb{Z}[\pi, \bar{\pi}] \subset O \subset \mathcal{O}_K$. We want to check if $O = \mathcal{O}_K$. First, the Chinese Remainder Theorem gives us the following proposition:

Proposition 9. *Let $\mathcal{O} = \text{End } A$ and let $\gamma \in \mathcal{O}_K$ be such that $\ell^e \gamma \in \mathbb{Z}[\pi, \bar{\pi}]$. There exists a unique integer polynomial P_γ of degree less than 4 such that $\ell^e p \gamma = P_\gamma(\pi)$, and γ is in \mathcal{O} if and only if $P_\gamma(\pi) = 0$ on $A[\ell^e]$.*

Proof. First note that $[\mathbb{Z}[\pi, \bar{\pi}] : \mathbb{Z}[\pi]] = p$ (see [17, p. 38]), so that $\ell^e p \gamma \in \mathbb{Z}[\pi]$, which means we can write $\ell^e p \gamma = P_\gamma(\pi)$ for a unique $P_\gamma \in \mathbb{Z}[x]$ of degree less than 4. Second, since we are dealing with ordinary abelian surfaces over \mathbb{F}_p , we have $p \nmid [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ by [17, Proposition 3.7], so that $\gamma \in \mathcal{O} \Leftrightarrow p \gamma \in \mathcal{O}$. Lastly, by the universal property of isogenies, we have that $P_\gamma(\pi) = 0$ on $A[\ell^e]$ if and only if $p \gamma \in \mathcal{O}$ (see [15]). Summing up, we only need to check that $P_\gamma(\pi) = 0$ on $A[\ell^e]$ to check that $\gamma \in \mathcal{O}$. \square

Remark 10. Since most of the curves in the isogeny class are not maximal, it is more efficient to check the condition $P_\gamma(\pi) = 0$ on $A[\ell], A[\ell^2], \dots$, rather than directly on $A[\ell^e]$.

4.2. Computing the ℓ^e -torsion. The obvious method of using Proposition 9 to test whether an element of \mathcal{O}_K lies in \mathcal{O} involves computing a basis of the ℓ^e -torsion group. The cost of such a computation depends on the degree of the extension where the ℓ^e -torsion points are defined. We have:

Lemma 11. *Let d be the degree such that the ℓ -torsion points of A are defined over \mathbb{F}_{p^d} . Then $d \leq \ell^4 - 1$. Furthermore, the ℓ^e -torsion is all defined over \mathbb{F}_q with $q = p^{d\ell^{e-1}}$.*

Proof. Let χ_π be the characteristic polynomial of π . Then d is the (multiplicative) order of X in the ring $\mathbb{F}_\ell[X]/\chi_\pi(X)$, so $d \leq \ell^4 - 1$. The second assertion follows from [17, Section 6]. \square

Remark 12. For *maximal* abelian surfaces, [17, Proposition 6.2] gives a better bound for d : In that case we have $d < \ell^3$, and if ℓ is completely split in \mathcal{O}_K we have $d \mid \ell - 1$.

We will use the following algorithm to compute points uniformly in an ℓ -primary group containing $A[\ell^e]$:

Algorithm 13.

Input: An abelian surface A/\mathbb{F}_p and a prime power ℓ^e .

Output: Uniform random points in the group $A(\mathbb{F}_{p^{d\ell^e}})[\ell^\infty]$, defined below.

1. Precomputation:

- (a) Let d be the (multiplicative) order of X in the ring $\mathbb{F}_\ell[X]/\chi_\pi(X)$ and set $d_e = d\ell^{e-1}$.
 - (b) Compute $\chi_{\pi^{d_e}}$ as the resultant in X of $\chi_\pi(Y)$ and $Y^{d_e} - X$, and write $\#A(\mathbb{F}_p^{d_e}) = \chi_{\pi^{d_e}}(1) = \ell^e \gamma$ with γ prime to ℓ .
2. Repeat as needed:
- (a) Take a random point P (uniformly) in $A(\mathbb{F}_p^{d_e})$.
 - (b) Return γP .

Algorithm 4.3 of [17] computes random points in $A(\mathbb{F}_p^{d_e})[\ell^e]$ by taking uniform random points P in $A(\mathbb{F}_p^{d_e})[\ell^\infty]$ and looking at the smallest k such that $\ell^k \cdot P$ is an ℓ^e -torsion point; it generates enough such random points so that the probability that they generate the full ℓ^e -torsion is sufficiently high, and then tests P_γ on these points of ℓ^e -torsion. The algorithm computes how many points are needed so that the probability of generating the full ℓ^e -torsion is greater than $1 - \epsilon$ for some $\epsilon > 0$, so the result is not guaranteed (that is, it is a ‘‘Monte Carlo’’ algorithm). This is very inconvenient in our setting since we need to test a lot of curves across different CRT primes p .

To ensure correctness we can check that the subgroup generated by the points obtained is of cardinality ℓ^{4e} , but this is costly. A more efficient way is as follows: $\{P_1, \dots, P_4\}$ is a basis of the ℓ^e -torsion if and only if $\{\ell^{e-1}P_1, \dots, \ell^{e-1}P_4\}$ is a basis of the ℓ -torsion. But that can be easily checked by computing the 6 Weil pairings $e_\ell(\ell^{e-1}P_i, \ell^{e-1}P_j)$ for $i < j$ and testing whether the corresponding 4×4 matrix is invertible. Since Weil pairings can be computed in time $O(\log(\ell))$, this is much faster. This is our first improvement, yielding a ‘‘Las Vegas’’ algorithm.

The second drawback of the approach of [17] is that, although the random points in $A(\mathbb{F}_p^{d_e})[\ell^\infty]$ are uniform, this is not always the case for the random points in $A(\mathbb{F}_p^{d_e})[\ell^e]$. To have a high probability of generating the full ℓ -torsion then requires taking many random points in $A(\mathbb{F}_p^{d_e})[\ell^\infty]$: If $A(\mathbb{F}_p^{d_e})[\ell^\infty] = \ell^s$, the algorithm requires $\ell^{s-4e}(-\log(\epsilon))^{1/2}$ random points to succeed with probability greater than $1 - \epsilon$. Since generating these points is the most costly part of the algorithm it is best to minimize the number of random points required. Our second improvement is to use an algorithm, due to Couveignes [13] and implemented in the Magma package AVIsogenies [4], to get uniform random points in $A(\mathbb{F}_p^{d_e})[\ell^e]$. Since the full algorithm is described in more detail in [4], we only give an example to illustrate it here.

Suppose that G is an ℓ -primary group generated by a point P of order ℓ^2 and a point Q of order ℓ . Assume that the first random point chosen is $P = R_1$, which gives an ℓ -torsion point $T_1 = \ell P$. The second random point R_2 chosen will be of the form $\alpha P + \beta Q$. In most cases, $\alpha \neq 0$, so the corresponding new ℓ -torsion point is $T_2 = \alpha \ell P$, a multiple of T_1 . However we can correct R_2 by the corresponding multiple: Compute $R'_2 = R_2 - \alpha R_1 = \beta Q$. Thus R'_2 gives the rest of the ℓ -torsion unless $\beta = 0$. In our setting we can use the Weil pairing to express a new ℓ -torsion point in terms of the generating set already constructed (except when we have an isotropic group, in this case we have to compute the ℓ^2 multiples), and we only need $O(1)$ random points to find a basis. The cost of finding a basis of the ℓ^e -torsion is then $O(d_e \log p + \ell^2)$ operations if $\mathbb{F}_p^{d_e}$.

4.3. Reducing the degree. The complexity of finding the basis is closely related to the degree of the extension d_e . Let d_0 be the minimal integer such that $(\pi^{d_0} - 1) \in$

$\ell\mathcal{O}_K$. Then $d_0 \mid d$, and, as remarked in [17], since we only need to check if $\mathcal{O} = \mathcal{O}_K$, we can first check that $(\pi^{d_0} - 1)/\ell$ lies in \mathcal{O} . In other words, we can check that the ℓ -torsion points of A are defined over $\mathbb{F}_{p^{d_0}}$ rather than over \mathbb{F}_{p^d} . If this is the case, the ℓ^e -torsion points are then defined over an extension of degree $d_0\ell^{e-1}$ of \mathbb{F}_p , which allows us to work with smaller extensions.

Another improvement we implemented to reduce the degree is to use twists. Let d'_0 be the minimal integer such that $((-\pi)^{d'_0} - 1) \in \ell\mathcal{O}_K$. Then there are three possibilities: We have either $d'_0 = d_0$, or $d'_0 = 2d_0$, or $d_0 = 2d'_0$. In the third case it is to our advantage to replace A by its twist, because the Frobenius of the twist is represented by $-\pi$, and we can therefore compute the points of ℓ^e -torsion by working over extensions of half the degree.

Example 14. Let H be the curve $y^2 = 80x^6 + 51x^5 + 49x^4 + 3x^3 + 34x^2 + 40x + 12$ of genus 2 over \mathbb{F}_{139} , and let J be the Jacobian of H . By computing the characteristic polynomial of Frobenius for J we find that

$$(\text{End } J) \otimes \mathbb{Q} \cong \mathbb{Q}\left(i\sqrt{13 + 2\sqrt{29}}\right),$$

and we would like to check whether $\text{End } J$ is maximal. In this example, we compute that $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]] = 3^5$, so we need to compute the points in $J[3^5]$, which live over an extension of degree 81. If we had checked the endomorphism ring of the Jacobian of the twist of H , we would have needed to work over an extension of degree 162.

4.4. Reducing the number of endomorphisms to test. One last improvement to the algorithm of [17] is to use the fact that $\text{End } A$ is an order; if we know that $\gamma \in \mathcal{O}$, then we know that the whole ring $\mathbb{Z}[\pi, \bar{\pi}, \gamma]$ is contained in \mathcal{O} . For example, suppose $\{1, \alpha_1, \alpha_2, \alpha_3\}$ is a basis for \mathcal{O}_K and $\alpha_3 = \alpha_1\alpha_2 \pmod{\mathbb{Z}[\pi, \bar{\pi}]^*}$. To check that $\mathcal{O} = \mathcal{O}_K$ we only have to check that α_1 and α_2 are in \mathcal{O} . In fact, since our algorithm works locally at primes ℓ , we only need the relation between α_3 and $\alpha_1\alpha_2$ to hold locally at ℓ .

We use this idea as follows: Suppose that we have checked that $\{\gamma_1, \dots, \gamma_k\}$ are endomorphisms lying in \mathcal{O} , and we want to check if $\gamma \in \mathcal{O}$. Let N_1 be the order of γ in the \mathbb{Z} -module $\mathcal{O}_K/\mathbb{Z}[\pi, \bar{\pi}, \gamma_1, \dots, \gamma_k]$, and N_2 be the order of γ in $\mathcal{O}_K/\mathbb{Z}[\pi, \bar{\pi}]$. If we write $N_2 = \prod \ell_i^{e_i}$, we only have to check that $(N_2/\ell_i^{e_i})\gamma \in \mathcal{O}$ for $\ell_i \mid N_1$. In fact, if the valuation of N_1 at ℓ_i is f_i , then we would only need to check that $(N_1/\ell_i^{f_i})\gamma \in \mathcal{O}$, which means testing if $N_1\gamma = 0$ on the $\ell_i^{f_i}$ -torsion, where $N_1\gamma$ is a polynomial in $\pi, \bar{\pi}$, and the γ_i ($i = 1, \dots, k$). We write this polynomial as $N_1/(pN_2)$ times a polynomial in π , so that we still need to compute the $\ell_i^{e_i}$ -torsion.

Example 15. Let H be the curve $y^2 = 10x^6 + 57x^5 + 18x^4 + 11x^3 + 38x^2 + 12x + 31$ of genus 2 over \mathbb{F}_{59} and let J the Jacobian of H . We have

$$(\text{End } J) \otimes \mathbb{Q} = \mathbb{Q}\left(i\sqrt{29 + 2\sqrt{29}}\right)$$

and we would like to check whether $\text{End } J = \mathcal{O}_K$. The ring \mathcal{O}_K is generated as a \mathbb{Z} -module by $1, \alpha, \beta, \gamma$, where α has order 2 in $\mathcal{O}_K/\mathbb{Z}[\pi, \bar{\pi}]$, β has order 4, and γ has order 40. The algorithm from [17] would require computing the elements of $J[2^3]$ and $J[5]$. But $(\mathcal{O}_K)_2 = \mathbb{Z}_2[\pi, \bar{\pi}, \alpha]$, so we only need to compute in $J[2]$ and $J[5]$.

4.5. The algorithm. Incorporating all these improvements yields the following algorithm:

Algorithm 16. Checking that $\text{End } A$ is maximal.

Input: An ordinary abelian surface A/\mathbb{F}_p with CM by K .

Output: *True* or *false*, depending on whether or not $\text{End } A = \mathcal{O}_K$.

1. Choose a basis $\{1, \alpha_1, \alpha_2, \alpha_3\}$ of \mathcal{O}_K and a basis $\{1, \beta_1, \beta_2, \beta_3\}$ of $\mathbb{Z}[\pi]$ such that $\beta_1 = c_1\alpha_1$, $\beta_2 = c_2\alpha_2$, $\beta_3 = c_3\alpha_3$ and $c_1, c_2, c_3 \in \mathbb{Z}$ with $c_1 \mid c_2 \mid c_3$.
2. (Checking where the ℓ -torsion lives.) For each $\ell \mid [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ do:
 - (a) Let d be the smallest integer such that $\pi^d - 1 \in \ell\mathcal{O}_K$, and d' be the smallest integer such that $(-\pi)^{d'} - 1 \in \ell\mathcal{O}_K$. If $d' < d$, switch to the quadratic twist.
 - (b) Compute a basis of $A[\ell](\mathbb{F}_{p^d})$ using the algorithm from [4].
 - (c) If this basis is of cardinality (strictly) less than 4, return *false*.
 - (d) (Checking the generators of \mathcal{O}_K .) For $i = 1, 2, 3$ do:
 - i. Let N_1 be the order of α_i in $\mathcal{O}_K/\mathbb{Z}[\pi, \bar{\pi}, \alpha_j \mid j < i]$ and N_2 the order of α_i in $\mathcal{O}_K/\mathbb{Z}[\pi, \bar{\pi}]$.
 - ii. If $\ell \mid N_1$, let e be the ℓ -valuation of N_2 and write $pN_2\alpha_i$ as a polynomial $P(\pi)$.
 - iii. Compute a basis of $A(\mathbb{F}_{p^{d\ell e-1}})[\ell^e]$.
 - iv. If $P(\pi) \neq 0$ on this basis, return *false*.
3. Return *true*.

Remark 17. Comments on the algorithm:

- With the way we choose the basis of \mathcal{O}_K , we have $e_1 \leq e_2 \leq e_3$ (for each ℓ dividing the index), so that when we abort early, we may not have the full $A[\ell^{e_3}]$ -torsion to compute. Likewise, rather than going through increasing ℓ , we could go through increasing degrees.
- Since we will apply this algorithm to a lot of different abelian varieties, we can precompute everything that is only related to \mathcal{O}_K and $\mathbb{Z}[\pi, \bar{\pi}]$. Then for each abelian variety A/\mathbb{F}_p we want to test, we only have to compute for all $\ell \mid [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ the subgroups $A[\ell](\mathbb{F}_{p^{d\ell}})$ and then subgroups $A(\mathbb{F}_{p^{d\ell e_i-1}})[\ell^{e_i}]$ for $i = 1, 2, 3$, testing polynomials of the Frobenius on them. Moreover, if ℓ is such that $\ell^2 \nmid [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ then $(\mathcal{O}_K)_\ell/\mathbb{Z}_\ell[\pi, \bar{\pi}]$ is cyclic, so that if $\pi^{d\ell} - 1 \notin \ell\mathbb{Z}[\pi, \bar{\pi}]$, we only need to check that the ℓ -torsion is defined over $\mathbb{F}_{p^{d\ell}}$ (see [17]).

4.6. Complexity. We will measure complexity in terms of operations in the base field \mathbb{F}_p , and we will neglect factors of $\log(p)$. Since the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ is bounded by a polynomial in p by [17, Proposition 6.2], evaluating the polynomials $P(\pi)$ (of degrees at most 3) is done in logarithmic time. The most expensive part of the algorithm is then the computation of $A[\ell^e]$, for the various ℓ dividing the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ where e is at most the ℓ -valuation of the index. According to Lemma 11 and Remark 12, the ℓ^e -torsion points live in an extension of degree at most $d = \ell^{e+3}$. Since $\#A(\mathbb{F}_{p^d}) = p^{2d(1+\epsilon)}$, computing a random point in $A(\mathbb{F}_{p^d})[\ell^e]$ takes $\tilde{O}(d^2)$ operations in \mathbb{F}_p . Correcting this random point requires some pairing computations, and costs at most $O(\ell^2)$ (in case the first points give an isotropic group). Since we need $O(1)$ such random points, the global cost is given by the following proposition (we will only need a very rough bound for the complexity analysis in Section 9):

Proposition 18. *Let $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]] = \prod \ell_i^{e_i}$ be the decomposition of the index into powers of primes. Then checking if an abelian surface in the isogeny class is maximal can be done in time $\sum \tilde{O}(\ell_i^{2e_i+6})$.*

Remark 19. One can compare to [17, Proposition 4.6] to see the speedup we gain in the endomorphism ring computation. We note that our method is exponential in the discriminant, while in [3] one can find a subexponential algorithm to compute the endomorphism ring of an ordinary abelian surface. In ongoing work with Gaetan Bisson, we have developed a method that combines the going-up algorithm of the next section with his endomorphism ring algorithm. Since we still need to take ℓ -isogenies for $\ell \mid [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ in the going-up step, this approach is mainly interesting when the index is divisible by a power of a prime.

5. GOING UP

“Going up” is the process of finding genus-2 curves with maximal endomorphism ring by moving from *any* curve in the isogeny class to a maximal one via isogenies. This is not always possible and we will explain some of the obstructions. One difficulty was already illustrated in [5, Example 8.2], where it was shown that there can be cycles in the isogeny graph involving only nonmaximal curves. Clearly, when trying to “go up”, the algorithm should avoid making cycles in the graph, and we propose one method to avoid that. Further difficulties arise from the fact that the graph of rational (ℓ, ℓ) isogenies can be disconnected, and can even have isolated nodes. This is an important caveat, as this means that our method for going up will not always succeed, so we only have a probabilistic algorithm; furthermore, we cannot currently estimate the probability of failure.

As noted in [17], for the type of fields we can deal with via the CRT method, the cost of going through p^3 Jacobians is dominant compared to checking if the endomorphism ring is maximal. (This imbalance is magnified in our case due to our faster algorithm to compute the ℓ^e -torsion.) In our algorithm, we try to find a random curve in the isogeny class, and we try to select p so that the probability of finding a curve in the right isogeny class is of magnitude $p^{3/2}$. In practice, finding one such curve is still the dominant aspect, which explains why we can afford to spend a lot of effort on going up from this curve.

The algorithm we propose for going up is made possible by the techniques developed in [28, 12, 31] for computing rational (ℓ, ℓ) -isogenies between abelian surfaces over finite fields. If A is an ordinary abelian surface with CM by K , then for each ℓ dividing the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$, we try to find an (ℓ, ℓ) -isogeny path starting from A and going to A' such that $(\mathcal{O}_K)_\ell = (\text{End } A')_\ell$. If this is possible, we let $A = A'$ in the next step (going to the next ℓ). A rather inefficient method for finding A' would be to use the algorithm for computing endomorphism rings which was detailed in the preceding section (modified to handle the case of nonmaximal orders), compute the endomorphism ring of $\text{End } A$ and the (ℓ, ℓ) -isogenous surfaces A' , and keep A' if its endomorphism ring is bigger than that of A . In this section we will describe a more efficient algorithm, which combines the endomorphism ring checks of the preceding section with a going-up phase. Since we are working locally in ℓ , we may as well suppose that we are working over \mathbb{Z}_ℓ .

5.1. Going up for one endomorphism. In this section, we suppose that we have an element $\alpha' \in \mathcal{O}_K$ such that $\alpha := \gamma \ell^e \alpha'$ lies in $\mathbb{Z}[\pi]$ for some $\gamma \in \mathcal{O}_K$ prime to ℓ .

Starting from an abelian surface A in the isogeny class, we want to find an abelian surface A' such that $\alpha/\ell^e \in \text{End } A'$ (or equivalently that $\alpha' \in \text{End } A'$ locally at ℓ).

We saw in Section 4 that α/ℓ^e is in the endomorphism ring of A if and only if $\alpha(A[\ell^e]) = 0$, and we know how to compute this subgroup. More generally, we let $N = \#\alpha(A[\ell^e])$. We think of N as a way to measure the ‘‘obstruction’’ to α/ℓ^e being an element of $\text{End } A$. Our algorithm is as follows: For each (ℓ, ℓ) -isogenous surface A' , we let $N' = \#\alpha(A'[\ell^e])$ and we replace A by A' if $N' < N$. We iterate this process until $N = 1$, in which case we have succeeded, or until we are stuck, in which case we try to find a new random abelian surface in the right isogeny class.

Rather than directly computing the obstruction $N = \#\alpha(A[\ell^e])$, we can compute the partial obstructions $N(\epsilon) := \#\alpha(A[\ell^\epsilon])$ for $\epsilon \leq e$. Starting from $\epsilon = 1$, we take isogenies until we find an abelian surface A with $N(\epsilon) = 1$, which means that $\alpha/\ell^\epsilon \in \text{End } A$. We will now try to take isogenies to reduce the obstruction of higher degree $N(\epsilon + 1)$. Let $k = \alpha(A[\ell^{\epsilon+1}]) \subseteq A[\ell]$. The following lemma helps us select the isogeny we are looking for:

Lemma 20. *With notation and assumptions as above, let A' be an abelian surface isogenous to A such that $\#\alpha(A'[\ell^{\epsilon+1}]) < \#\alpha(A[\ell^{\epsilon+1}])$. Then the kernel of the isogeny $A \rightarrow A'$ intersects nontrivially with $k = \alpha(A[\ell^{\epsilon+1}])$.*

Proof. Let $f: A \rightarrow A'$ be a rational isogeny between A and A' . Then since α is a polynomial in the Frobenius, we have $\alpha \circ f = f \circ \alpha$. In particular, f maps $\alpha(A[\ell^{\epsilon+1}])$ to $\alpha(A'[\ell^{\epsilon+1}])$. If $\#\alpha(A'[\ell^{\epsilon+1}]) < \#\alpha(A[\ell^{\epsilon+1}])$ then there exists $x \in \text{Ker } f \cap \alpha(A[\ell^{\epsilon+1}])$. \square

This gives the following algorithm:

Algorithm 21. Going up for one endomorphism α/ℓ^e .

Input: An ordinary abelian surface A/\mathbb{F}_p with CM by K , a prime power ℓ^e , and an $\alpha \in \ell^e \mathcal{O}_K$.

Output: An abelian surface A'/\mathbb{F}_p isogenous to A such that $\alpha/\ell^e \in \text{End } A'$, or *fail*.

1. Set $\epsilon = 1$.
2. Compute $N(\epsilon) = \#\alpha(A[\ell^\epsilon])$.
3. If $N(\epsilon) = 1$, do:
 - (a) If $\epsilon = e$ then return A .
 - (b) Otherwise, set $\epsilon := \epsilon + 1$, and go back to Step 2.
4. At this point, $N(\epsilon) > 1$. Let \mathcal{L} be the list of all rational maximal isotropic subgroups of $A[\ell]$ which intersect nontrivially with $\alpha(A[\ell^\epsilon])$. For $k \in \mathcal{L}$ do:
 - (a) Compute $A' = A/k$.
 - (b) Let $N'(\epsilon) = \#\alpha(A'[\ell^\epsilon])$.
 - (c) If $N'(\epsilon) < N(\epsilon)$, set $A = A'$ and go back to Step 2.
5. Return *fail*.

Remark 22. As in Section 4 we let d_0 be the minimal integer such that $(\pi^{d_0} - 1) \in \ell \mathcal{O}_K$ and d the minimal integer such that $(\pi^d - 1) \in \ell \mathbb{Z}[\pi]$. Then the ℓ^ϵ -torsion points of A are defined over an extension of degree $d\ell^{\epsilon-1}$. If moreover $(\pi^{d_0} - 1)/\ell \in \text{End } A$ they are actually defined over an extension of degree $d_0\ell^{\epsilon-1}$.

Therefore when we try to go up globally for all endomorphisms α , the first step is to try to go up for the endomorphism $(\pi^{d_0} - 1)/\ell$. During the algorithm, the obstruction N is given by the size of the kernel of $\pi^{d_0} - 1$, whose rank is 4 minus

the rank of the ℓ -torsion points defined over $\mathbb{F}_{p^{d_0}}$. So we compute the size of a basis of $A[\ell](\mathbb{F}_{p^{d_0}})$ and take isogenies, where this size increases until we find the full rank.

5.2. Going up globally. Let $\{1, \alpha_1/\ell^{e_1}, \alpha_2/\ell^{e_2}, \alpha_3/\ell^{e_3}\}$ be a generating set for the maximal order $(\mathcal{O}_K)_\ell$ over the subring $\mathbb{Z}_\ell[\pi, \bar{\pi}]$, where $\alpha_i \in \mathbb{Z}_\ell[\pi, \bar{\pi}]$. Starting from an abelian surface A in the isogeny class, we want to find an abelian surface which is maximal at ℓ .

We could apply Algorithm 21 for each α_i/ℓ^{e_i} , but the algorithm does not guarantee that the endomorphisms already defined on A stay defined during the process, so we would observe loops on nonmaximal abelian surfaces with this method. Moreover we want to reuse the computations of $A[\ell^\epsilon]$, which are the expensive part of the process.

If $N_i = \#\alpha_i(A[\ell^{d_i}])$ for $i = 1, 2, 3$ is the obstruction corresponding to α_i , we define N to be the global obstruction $N = \sum N_i$. We can then adapt the same method: For each (ℓ, ℓ) -isogenous A' , if $N'_i = \#\alpha_i(A'[\ell^{d_i}])$, then we replace A by A' if $\sum N'_i < \sum N_i$. We iterate this process until all the $N_i = 1$, in which case we go to the next ℓ , or until we are stuck, in which case we try to find a new random abelian surface in the right isogeny class.

As before, if $e = \max(e_1, e_2, e_3)$ we first compute $A[\ell^\epsilon]$ and the partial obstructions $N_i(\epsilon) = \#A[\ell^{\min(\epsilon, e_i)}]$ (for $i = 1, 2, 3$). We do the same for the (ℓ, ℓ) -isogenous abelian surfaces, and switch to the new one if $\sum N_i(\epsilon)$ decreases (strictly). This allows working with smaller torsion in the beginning steps.

The level ϵ of the individual obstruction we are working on depends on the endomorphism considered, so if we get stuck on level ϵ , we may have to look at level $\epsilon + 1$ even if not all endomorphisms α_i/ℓ^ϵ are defined yet. For instance, in the case where we are only dealing with two generators, there are examples where $N_1(\epsilon) = 1$, $N_2(\epsilon) \neq 1$ and $N'_1(\epsilon) = 1$, $N'_2(\epsilon) = N_2(\epsilon)$ for all (ℓ, ℓ) -isogenous abelian surfaces A' , so we are stuck on level ϵ . However we can still find an isogenous A' such that $N'_1(\epsilon + 1) < N_1(\epsilon + 1)$.

Finally, as in Remark 22, we first try to go up in a way that increases the size of $A(\mathbb{F}_{p^{d_0}})[\ell]$. If we are unlucky and get stuck, we switch to the computation of the full ℓ -torsion over $\bar{\mathbb{F}}_p$. This method allows working over the smallest extension to compute $A[\ell^\epsilon]$ as soon as possible.

A summary of the algorithm with the notation from above is given below:

Algorithm 23. Going up.

Input: An ordinary abelian surface A/\mathbb{F}_p with CM by K , and a prime ℓ .

Output: An abelian surface A'/\mathbb{F}_p with $\text{End } A = \mathcal{O}_K$ (locally at ℓ), or *fail*.

1. (Special case for the endomorphism $(\pi^{d_0} - 1)/\ell$.) Compute a basis B of $A(\mathbb{F}_{p^{d_0}})[\ell]$. If $\#B < 4$, compute a basis B' of $A'(\mathbb{F}_{p^{d_0}})[\ell]$ for each (ℓ, ℓ) -isogenous abelian surface A' . If $\#B' > \#B$, restart the algorithm with $A' = A$. If $\#B = 4$ or we get stuck, go to the next step.
2. Set $\epsilon = 1$.
3. Compute² $N_i(\epsilon) = \#\alpha_i(A[\ell^{\min(\epsilon, e_i)}])$ for $i = 1, 2, 3$.
4. If $\{N_i : i = 1, 2, 3\} = \{1\}$, do:
 - (a) If $\epsilon = \max(e_i : i = 1, 2, 3)$ then return A .

²The degree of the extension where the full ℓ^ϵ -torsion is defined depends on whether Step 1 succeeded.

- (b) Otherwise, set $\epsilon := \epsilon + 1$ and go back to Step 3.
- 5. Let \mathcal{L} be the list of all rational maximal isotropic kernels of $A[\ell]$ which intersect nontrivially with one of the $\alpha_i(A[\ell^{\min(\epsilon, e_i)}])$. For $k \in \mathcal{L}$ do:
 - (a) Compute $A' = A/k$.
 - (b) Let $N'_i(\epsilon) = \#\alpha_i(A'[\ell^{\min(\epsilon, e_i)}])$.
 - (c) If $\sum N'_i(\epsilon) < \sum N_i(\epsilon)$, restart the algorithm with $A = A'$ (but do not reinitialize ϵ in Step 2).
- 6. If we get stuck and $\epsilon < \max(e_i : i = 1, 2, 3)$, set $\epsilon := \epsilon + 1$ and go back to Step 3.
- 7. Return *fail*.

5.3. Cost of the going-up step. We will see in the examples that the going-up step is a very important part in speeding up the CRT algorithm in practical computations. However, since it is doomed to fail in some cases (see Remark 25), we need to check that it will not dominate the complexity of the rest of the algorithm, so that in theory there will be no drawback to using it. Thus we need to estimate the cost of the going-up step.

The going-up phase is a mix of endomorphism testing and isogeny computations. We already analysed the cost of the endomorphism testing in the preceding section. For the isogeny computation, the points in the kernel of rational (ℓ, ℓ) -isogenies live in an extension of degree at most $\ell^2 - 1$. Transposing the analysis of Section 4.6 to this case shows that the computation of all of the points in these kernels takes at most $\tilde{O}(\ell^4)$ operations in \mathbb{F}_p . There are at most $O(\ell^3)$ such kernels, and each isogeny computation takes at most $\tilde{O}(\ell^4)$ operations in the extension. The final cost is at most $\tilde{O}(\ell^9)$ operations in \mathbb{F}_p for computing all isogenies. For each of the $O(\ell^3)$ isogenous abelian surfaces we do (part of) the endomorphism ring computation, which takes $\tilde{O}(\ell^{2e+6})$ operations, according to Section 4.6. Since the global obstruction computed is of size $O(\ell^e)$, we do at most $O(e)$ steps. The global complexity is then given as follows:

Proposition 24. *Let $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]] = \prod \ell_i^{e_i}$ be the decomposition of the index into prime factors. Then the going-up phase either fails or is done in at most $\tilde{O}(\sum \ell_i^{2e_i+9})$ operations in the base field.*

Remark 25. It is important to note that the going-up phase does not always succeed. We will give some examples of that in Section 10. First, as noted in the introduction of this section, the (ℓ, ℓ) -isogeny graph is not always connected, so if we start with a curve not in the same component as a maximal curve, there is no way to find the maximal curves using only (ℓ, ℓ) -isogenies. Second, even if the curve is in the same component as a maximal curve, finding a maximal curve may involve going through isogenous curves that increase the global obstruction, so the going-up algorithm would not find it.

In practical computations we observed the following behavior: In the very large majority of the cases where we were not able to go up, there actually did not exist any rational (ℓ, ℓ) -isogenies for any curve in the isogeny class. If χ_π is the characteristic polynomial, this can be detected by the fact that χ_π does not factor modulo ℓ as $\chi_\pi = P\bar{P} \pmod{\ell}$ (where \bar{P} is the conjugate of P under the action $\pi \rightarrow p/\pi$, which sends the Frobenius to the Verschiebung). In this situation, there is no way to go

up even locally at ℓ . This gives a criterion for estimating whether one can go up for this ℓ .

6. COMPUTING MAXIMAL CURVES FROM MAXIMAL CURVES

Once a maximal curve in the isogeny class has been found via the random search and going-up steps, we use isogenies to find the other maximal curves. The set of maximal curves in the isogeny class corresponding to a fixed CM-type Φ is a principal homogeneous space under the action of the Shimura class group

$$\mathfrak{C} = \{(I, \rho) \mid I \text{ a fractional } \mathcal{O}_K\text{-ideal with } I\bar{I} = (\rho), \rho \in K^+ \text{ totally positive}\}/K^*,$$

associated to the primitive quartic CM field K , which acts by isogenies (see for instance [5, §3]).

However, using the Magma package AVIsogenies we can only compute isogenies with a maximal isotropic kernel. The lemma below show that in terms of the Shimura class group, this means that we can only compute the action corresponding to (equivalences classes) of elements of the form (I, ℓ) , where I is an ideal in K and ℓ is a prime number.

Lemma 26. *Let (I, ρ) be an element of the Shimura class group \mathfrak{C} and let ℓ be a prime. Then the action of (I, ρ) on a maximal abelian surface A corresponds to an isogeny with maximal isotropic kernel in $A[\ell]$ if and only if $\rho = \ell$ (so if and only if I has relative norm ℓ).*

Proof. This follows from the construction of the action of \mathfrak{C} on the set of maximal abelian surfaces. The action is given by the isogeny $f: \mathbb{C}^2/\Lambda \rightarrow \mathbb{C}^2/I\Lambda$ and moreover the action of \bar{I} corresponds to the dual isogeny \hat{f} (here we identify the abelian surface A with its dual \hat{A} via the principal polarization induced from the CM data). Since ℓ is prime, the isogeny corresponding to I is an (ℓ, ℓ) isogeny if and only if $I\bar{I} = (\rho) = (\ell)$. \square

Therefore to ensure that we can find all other maximal curves using this type of isogeny we make the following heuristic assumption.

Assumption. *There is a polynomial P such that for every primitive quartic CM field K , the Shimura class group associated to K is generated by elements of the form (I, ℓ) , where ℓ ranges over the prime numbers less than $P(\log \Delta)$ and where Δ is the discriminant of K .*

Justification. We have tested this assumption on numerous examples, using the bound $12 \log \Delta'$, where Δ' is the discriminant of the reflex field, which is itself $O(\Delta^2)$. The assumption on the size of the isogenies will be used in the complexity analysis. At worst, we know (under GRH) that the class group of the reflex field is generated by prime ideals of degree one and of norm polynomial in $\log \Delta$ [1, Theorem 1]. But if I is such an ideal of \mathcal{O}_{K_Φ} of norm prime to p , then the element $(\text{TN}(I), N(I))$ will give a horizontal isogeny. So we will at least be able to compute all the maximal curves that are deduced from the first one by an action coming from the type norm. As we will see in the complexity analysis in Section 9, this is sufficient for most discriminants Δ .

Lemma 27. *Let A be an ordinary abelian surface with $(\text{End } A) \otimes \mathbb{Q} = K$, and let $f: A \rightarrow B$ be an isogeny of degree prime to $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$. Then $\text{End } A = \text{End } B$.*

Proof. Let d be the smallest integer that factorizes through f , so $d = f\tilde{f}$ for some isogeny $\tilde{f}: B \rightarrow A$. By assumption d is prime to the index. If $\alpha \in \text{End } A$, then $f \circ \alpha \circ f = d\alpha$ is an endomorphism of B . Since $[\mathcal{O}_K : \text{End } B]$ is prime to d , we have that $\alpha \in \text{End } B$. The same argument shows that $\text{End } B \subseteq \text{End } A$, so $\text{End } A = \text{End } B$. \square

Note that we can precompute generators of the Shimura class group since this data does not depend on the current prime p . We want to find generators of relative norm a prime $\ell \in \mathbb{Z}$ with ℓ as small as possible, since the size of ℓ will directly influence the time spent to find the other maximal curves.

Now for a CRT prime p , there may exist among the generators we have chosen some that divide the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$. We can either find other generators (whose norm will be bigger), or still try to use the precomputed generators. In this case, if such a generator has norm ℓ , then not all new (ℓ, ℓ) -isogenous abelian surfaces will be maximal, so we have to use Algorithm 16 to test which of them is maximal. In that case, after the isogeny is applied, the ℓ^e -torsion (in the notation of Section 5) must again be computed, along with the action of the generators of $(\mathcal{O}_K)_\ell$ over $\mathbb{Z}[\pi, \bar{\pi}]_\ell$. The trade-off depends then on the degree of the extension field required to compute the ℓ^e -torsion for small ℓ dividing the index versus the degree of the field of definition for the points in the kernel of the ℓ -isogeny for ℓ not dividing the index.

Finally, we can also use the group structure of the Shimura class group as follows: Suppose that we have computed maximal curves corresponding to the action of $\alpha_1, \dots, \alpha_t \in \mathfrak{C}$, and we want to find new maximal curves by computing (ℓ, ℓ) -isogeny graphs starting from these curves. Then if $\mathfrak{C}(\ell)$ is the set of elements of the form (I, ℓ) in \mathfrak{C} , then the number of maximal curves that we can find in this way is the cardinality of the subgroup generated by the α_i and $\mathfrak{C}(\ell)$. In particular, as soon as we reach this number, we can stop the computation since it will not yield any new maximal curves. This is particularly useful when ℓ divides the index, because then we avoid some endomorphism tests. In the isogeny graph computation done by AVIsogenies, each node is computed twice since there are two edges between adjacent nodes (corresponding to the isogeny and the dual). Here, since we know the number of nodes, we can abort the computation early.

We thus obtain the following algorithm:

Algorithm 28. Finding all maximal curves from one maximal curve.

Input: An ordinary abelian surface A/\mathbb{F}_p with CM by (\mathcal{O}_K, Φ) .

Output: All abelian surfaces over \mathbb{F}_p with CM by (\mathcal{O}_K, Φ) .

1. Precomputation: Compute a set of generators of the Shimura class group with relative norm ℓ as small as possible. (The set is not chosen to be minimal; on the contrary, we want some redundancy.) For each of the generators, compute the extension degree of the field of definition of the geometric points of the kernel corresponding to this generator.
2. For each generator of (relative) norm ℓ dividing the index, replace the previous degree by the degree of the extension where the ℓ^e -torsion lives. (Usually e is the ℓ -valuation of the index, but the tricks from Section 4 can sometimes reduce it.)
3. Sort the generators by the corresponding degrees to get a list (g_1, \dots, g_n) .
4. For each generator g_i on the list, let ℓ_i be its norm and do:

- (a) Compute the surfaces (ℓ_i, ℓ_i) -isogenous to the one already found. If ℓ_i divides the index, then do an endomorphism ring computation from Section 4 and keep only the maximal curves.
- (b) Repeat until the number of maximal abelian surfaces is equal to $\#\langle \mathfrak{C}(\ell_1), \dots, \mathfrak{C}(\ell_i) \rangle$.

7. STRATEGIES FOR SIEVING CRT PRIMES P

Since we have some latitude in the CRT primes p , we can sieve the primes to use. For instance, we will reject a prime p if the size of the isogeny class is too small, or if computing the endomorphism ring or going up is too costly for this prime. We use a dynamic approach: we reevaluate each discarded CRT prime against the new ones found. In this section, we explain how we estimate the difficulty of the computation associated to one CRT prime.

7.1. Cost of testing if a curve is maximal. Before using a prime for the CRT computation, we first need to check whether testing if a curve is maximal is too expensive. For this, we compute which subgroups $A[\ell^e]$ are needed to compute to test if A is maximal as in Section 4, and in which extensions the points of these subgroups are defined. As already remarked, for ℓ dividing the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$, e is usually the valuation of ℓ in the index, but some of the tricks from Section 4 can reduce it.

If the extension in which we need to do the computation is too large, we exclude the prime p . We will explain later how we estimate whether the computation of the endomorphism ring is too costly compared to the current parameters.

7.2. Size of the isogeny class. For each CRT prime p , the first phase of our algorithm relies on finding a genus 2 curve over \mathbb{F}_p in the right isogeny class. The larger the isogeny class, the larger the probability of finding a curve in the right isogeny class quickly. There are p^3 isomorphism classes of genus 2 curves over \mathbb{F}_p , and since the area of Figure 10.1 in [27] is $32/3$, there are approximately $(32/3)p^{3/2}$ isogeny classes. We could then expect that on average, each isogeny class has roughly $\frac{3p^{3/2}}{32}$ curves.

However it happens that, for a fixed primitive CM field K , for some primes p the isogeny class corresponding to the Frobenius element π can be unfortunately small. In those cases, our algorithm has a lower chance of finding a curve in the isogeny class quickly, so it is most likely more efficient for the algorithm to skip that CRT prime and proceed to another prime where the chance of finding a curve in the right isogeny class is bigger.

To determine whether a potential CRT prime should be skipped or not, we need to estimate the size of the isogeny class. If the estimated size is not at least a certain fraction of $p^{3/2}$ then we skip the prime. The size of the isogeny class is given as $\sum_{\mathcal{O}} \#\mathfrak{C}(\mathcal{O})$, the sum of the sizes of the Shimura class groups associated to all the orders containing $\mathbb{Z}[\pi, \bar{\pi}]$ which are stable under complex conjugation. Of course we do not compute the Shimura class groups in order to make this estimate (even computing only the class group can be too expensive for suborders of large discriminant), but we need to estimate their sizes.

Lemma 6.3 in [27] can be used to calculate the size of the isogeny class exactly. However it requires computing the lattice of suborders of the maximal order \mathcal{O}_K which contain $\mathbb{Z}[\pi, \bar{\pi}]$ and this is already too expensive. In practice, it is enough

to estimate the size of the isogeny class using only the factorization of the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$ and a calculation involving only the order $\mathbb{Z}[\pi, \bar{\pi}]$.

We compute using the proof of Lemma 6.3 in [27]

$$\#\mathfrak{C}(\mathbb{Z}[\pi, \bar{\pi}]) = \frac{c \# \text{Cl}(\mathbb{Z}[\pi, \bar{\pi}]) \text{Reg}(\mathbb{Z}[\pi, \bar{\pi}])}{2 \# \text{Cl}(\mathbb{Z}[\pi + \bar{\pi}]) \text{Reg}(\mathbb{Z}[\pi + \bar{\pi}])}$$

where c is the size of the co-kernel of the norm map from the class group of $\mathbb{Z}[\pi, \bar{\pi}]$ to the narrow class group of $\mathbb{Z}[\pi + \bar{\pi}]$. In the following we will use $c = 1$ in order to have a lower bound for $\#\mathfrak{C}(\mathbb{Z}[\pi, \bar{\pi}])$. Moreover Equation (6.1) of [27] gives us

$$\# \text{Cl}(\mathbb{Z}[\pi, \bar{\pi}]) \text{Reg}(\mathbb{Z}[\pi, \bar{\pi}]) = \# \text{Cl}(\mathcal{O}_K) \text{Reg}(\mathcal{O}_K) [\widehat{\mathcal{O}}_K^* : \widehat{\mathbb{Z}}[\pi, \bar{\pi}]^*].$$

If I is the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$, we can compute $[\widehat{\mathcal{O}}_K^* : \widehat{\mathbb{Z}}[\pi, \bar{\pi}]^*]$ as $[(\mathcal{O}_K/I)^* : (\mathbb{Z}[\pi, \bar{\pi}]/I)^*]$. It is easy to compute $\#(\mathbb{Z}[\pi, \bar{\pi}]/I)^*$; since it is a torsion group and I is prime to $p = [\mathbb{Z}[\pi, \bar{\pi}] : \mathbb{Z}[\pi]]$ it is equal to $\prod_{\ell} \#(\mathbb{Z}[\pi]/I)_{\ell}^*$. Now if ℓ^e is a prime power dividing I , and $\chi_{\pi} = \prod P_i^{e_i}$ the factorization of $\chi_{\pi} \pmod{\ell}$ then

$$\#(\mathbb{Z}[\pi]/I)_{\ell}^* = \prod_i (\ell^{\deg P_i e_i} - \ell^{\deg P_i (e_i - 1)}).$$

Likewise, we can compute $\#(\mathcal{O}_K/I)_{\ell}^*$ by looking at the decomposition of ℓ in \mathcal{O}_K .

Now we use the following estimate: for each divisor d of the index I , the contribution of orders O such that $[O : \mathbb{Z}[\pi, \bar{\pi}]] = d$ to the estimate of the number of curves in the isogeny class is $\#\mathfrak{C}(\mathbb{Z}[\pi, \bar{\pi}])/d$. So we estimate the number of curves as

$$\sum_{d | [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]} \#\mathfrak{C}(\mathbb{Z}[\pi, \bar{\pi}])/d$$

(for d not divisible by an ℓ where we can't go up). We will see in Section 10 how this estimate compares to some real examples.

7.3. Estimating the probability of going up. In practice, we are not interested in the size of the isogeny class, but in the number of curves in the isogeny class from which we can go up. From numerous experiments, we have observed that most of the cases where we can't go up for a particular ℓ arise because no rational (ℓ, ℓ) -isogenies exist at all. But we can easily detect this case by using Proposition 29.

We can thus estimate the number \mathbf{C} of curves from which we can go up as

$$\sum_{d | [\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]} \#\mathfrak{C}(\mathbb{Z}[\pi, \bar{\pi}])/d$$

but where we restrict the divisors d to be such that $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]/d$ is not divisible by a prime ℓ where we cannot go up.

Now if \mathbf{C}_0 is the estimated number of curves in the isogeny class, we estimate that we need \mathbf{C}_0/\mathbf{C} going up tries before succeeding. Now we keep the prime p if \mathbf{C}_0 is not too small, and if the cost of doing all these endomorphism ring computations is at most the time needed to find a curve in the isogeny class where we can go up. (As we will see in the complexity analysis, the endomorphism ring computation is not the dominant phase, so in practice this condition is almost always satisfied).

7.4. A dynamic selection of primes. When we select a prime p , we hope that the size \mathbf{C}_0 of the isogeny class (or more precisely the number of curves \mathbf{C} where we can go up) is approximatively the average size $\Theta(p^{3/2})$ of an isogeny class. However, for small primes p , even if the isogeny class is small, it could be worth it to go through all p^3 curves corresponding to p rather than through $q^{3/2}$ curves for a larger prime q .

So we use a dynamic approach: for each new CRT prime p , we compute \mathbf{C}/p^3 , the expected probability of finding a curve where we can go up. We also look at the corresponding probabilities for the previously discarded CRT primes, and we use the prime giving the maximum probability provided it is more than $1/16p^{-3/2}$.

8. CRT LIFTING OF THE COEFFICIENTS

8.1. Denominators via Bruinier-Yang formula. We note that contrary to the elliptic curve case, the coefficients of the class polynomials are rationals, not integers.

The CRT algorithm terminates when the lifted class polynomials (with denominators cleared) are constant from one CRT prime to the next. The probability that the class polynomials are correct when this happens was estimated in [17, Remark 7.2].

We estimate the denominators using the Bruinier-Yang conjectural formula [6] (proved only for special cases [41, 42]), with minor adjustments from [21], and adapting it to the fact that we use invariants from [36, Appendix 3], which alters the denominator formulas by small powers of 2. A formula for the factorization of the denominators which holds for general primitive quartic CM fields was recently given in [25] which produces a multiple of the denominators when allowing for cancellation with the numerators and the case where K^+ does not have class number 1. As in [15, Theorem 3], we can multiply by the denominators and then use the CRT to reconstruct the polynomials.

We note that since we are using the invariants from [36, Appendix 3], the power of I_{10} in the denominators is 1, 2, 2 rather than 6, 4, 4; but we have to multiply them by the constant factors $2^7, 2^5, 2^{14}$.

8.2. The CRT. In the cyclic case, we compute the class polynomials modulo small integer primes, and we use the CRT to get the result modulo the product P (the “precision”) of these small primes. Once the precision is large enough, we can recover the polynomials over \mathbb{Z} by lifting each coefficient to an integer in the interval $[-P/2, P/2]$.

In the dihedral case, the primes are in $\mathcal{O}_{K_\Phi^+}$, and so is the precision ideal P . Here we explain how to lift a coefficient $x \bmod P$ to $\mathcal{O}_{K_\Phi^+}$. Take the Minkowski embedding of a lift of x , and find the closest vector c_x in the lattice associated to P in the Minkowski embedding. Then c_x corresponds to an element of the ideal P , and our final lift is $x - c_x$. We note that the lattice is of rank 2, so we can directly compute the closest vector rather than doing an LLL approximation.

8.3. Lifting without denominators. We note that in the dihedral case, the denominator from the formulas in [6, 21, 25] is too large, as it takes into account both CM-types. This increases the size of the coefficients we compute, so that using those denominator formulas does not actually give better results than doing a rational reconstruction directly.

With the notation from above, from $x \bmod P$ we want to do a rational lift of x . This time we embed the lattice associated to P into the lattice of rank 3 obtained

by adjoining the vector $[Cx_1, Cx_2, C]$ where x_1 and x_2 are the two real embeddings of (a lift of) x and C is a constant accounting for how skewed we expect the size of the denominator to be compared to the numerators. A minimal vector in this lattice will correspond to an element $N = c + Dx$ where $c \in p$ and D is an integer. We then take N/D as our lift for x .

This solution requires the precision to be the sum of the bit sizes of the numerators and denominator, so it can be even better than using the denominator formulas for small denominators, where there may be cancellation with the numerators.

8.4. Finding irreducible factors. Computing irreducible factors of the class polynomials directly allows to recover them faster since they have smaller coefficients, so that we need less precision.

We know that the orbits under the action of the type norm give the irreducible factors of the class polynomials (or more precisely the irreducible components of the CM locus) [36, Chapter 3]. It is easy to compute these orbits modulo each CRT prime p using the tools from Section 6. However we need to be able to glue the correct orbits together when doing the CRT. For this we can use the “trace trick” from [16].

9. COMPLEXITY

In this section, we give a mostly-heuristic analysis of how Algorithm 23 (the going-up algorithm) and Algorithm 28 (the algorithm to find all maximal curves from one maximal curve) affect the asymptotic complexity of Algorithm 1. We will sometimes call the isogenies we compute in the going-up algorithm *vertical steps*, and the isogenies we compute in Algorithm 28 *horizontal steps*; this is in analogy with the corresponding terminology in the elliptic curve case.

We begin with a quick reminder of the rough complexity analysis of the CRT method in the elliptic curve case, where K is a quadratic imaginary field. In this case there is only one class polynomial H , whose degree is the class number of \mathcal{O}_K , and classical bounds give that $\deg H = \tilde{O}(\sqrt{\Delta})$, where Δ is the discriminant of \mathcal{O}_K . Likewise, the coefficients of H have size $\tilde{O}(\sqrt{\Delta})$. So the whole class polynomial is of size $\tilde{O}(\Delta)$.

Each CRT prime p gives $\log(p)$ bits of information, so neglecting logarithmic factors, we need about $\sqrt{\Delta}$ primes. CRT primes split completely in the Hilbert class field of K , whose Galois group is $\text{Cl}(\mathcal{O}_K)$, so by the Chebotarev theorem the density of CRT primes is roughly $1/\#\text{Cl}(\mathcal{O}_K) \simeq 1/\sqrt{\Delta}$. Neglecting logarithmic factors again, we therefore expect the biggest prime p to be of size $\tilde{O}(\Delta)$.

Now there are $O(p)$ isomorphism classes of elliptic curves, and $\tilde{O}(\sqrt{\Delta})$ maximal curves, so one is found in time $\tilde{O}(p/\sqrt{\Delta}) = \tilde{O}(\sqrt{p})$. Once one maximal curve is found, all others can be obtained using isogenies of degree logarithmic in Δ , so one can recover all maximal elliptic curves over \mathbb{F}_p in time $\tilde{O}(\sqrt{p}) = \tilde{O}(\sqrt{\Delta})$.

We need $\sqrt{\Delta}$ CRT primes, so the total cost is $\tilde{O}(\Delta)$. The CRT reconstruction can be done in quasilinear time too, so in the end the algorithm is quasilinear, even without using a vertical step. If we had not used horizontal steps, the complexity would have been $\tilde{O}(\Delta^{3/2})$.

Now consider the genus-2 case. Let $\Delta_0 = \Delta_{K^+/\mathbb{Q}}$ and $\Delta_1 = N_{K^+/\mathbb{Q}}(\Delta_{K/K^+})$, so that $\Delta = \Delta_{K/\mathbb{Q}} = \Delta_1 \Delta_0^2$. Then the degree of the class polynomials is $\tilde{O}(\Delta_0^{1/2} \Delta_1^{1/2})$,

while the height of their coefficients is bounded by $\tilde{O}(\Delta_0^{5/2}\Delta_1^{3/2})$ (see [36, §II.9] and [20]). In practice, we observe [36, Appendix 3] that the coefficient height is bounded by $\tilde{O}(\Delta_0^{1/2}\Delta_1^{1/2})$, and we will use this observed bound in the following analysis. According to [5, §6.4], the smallest prime is of size $\tilde{O}(\Delta_0\Delta_1)$. We need $\tilde{O}(\Delta_0^{1/2}\Delta_1^{1/2})$ CRT primes, and an analysis using [23], as in [2, §5, Lemma 3], shows that the largest prime is also $\tilde{O}(\Delta_0\Delta_1)$. We remark that the sieving phase does not affect the size of the largest prime (apart from the constant in the big O) as long as we sieve a positive density of CRT primes.

For the horizontal step, the isogeny computation involves primes of size logarithmic in Δ , so the cost of this step is quasilinear in the number $\tilde{O}(\Delta_0^{1/2}\Delta_1^{1/2})$ of maximal curves. This is under the Assumption from Section 6. Without this assumption, what we know is that for each ideal I in \mathcal{O}_{K_Φ} of norm prime to p , the element $(\text{TN}(I), N(I))$ is an element of the Shimura class group whose action is given by a maximally isotropic kernel. In the horizontal step, we can then compute the action of $\text{TN}(\text{Cl}(\mathcal{O}_{K_\Phi}))$ by isogenies of size logarithmic in Δ . By Lemma 6.5 of [5], the cofactor is bounded by $2^{6w(D)+1}$, where $w(D)$ is the number of prime divisors of D . This gives a bound on the number of horizontal isogeny steps we need to take. As remarked in [5, p. 516], we have $w(n) < 2 \log \log n$ outside a density-0 subset of very smooth integers, so the corresponding factor can be absorbed into the \tilde{O} -notation.

In contrast, the complexity of the endomorphism ring computation and the going-up phase involves the largest prime power dividing the index $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]]$. According to Proposition 6.1 of [17] we have that $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]] \leq 16p^2/\sqrt{\Delta}$. For the size of the CRT prime we are considering, we see that $[\mathcal{O}_K : \mathbb{Z}[\pi, \bar{\pi}]] = \tilde{O}(\Delta_0\Delta_1^{3/2})$. We fix $\epsilon = 1/2$. Assuming that the index is uniformly distributed, [14] showed that there is a positive density of CRT primes where the largest prime power dividing the index is $O(\Delta_0^{\epsilon/100}\Delta_1^{\epsilon/100})$. By the complexity analysis of Sections 4.6 and 5.3, we see then that there is a positive density of primes where these algorithms take time at most $O(\Delta_0^\epsilon\Delta_1^\epsilon)$.

We then let $p = \tilde{O}(\Delta_0\Delta_1)$ be a CRT prime. There are $O(\sqrt{p})$ maximal curves, so we expect the isogeny class to be of size $\Theta(p^{3/2})$ (see Heuristic 6.6 in [5]). Up to isomorphism over the algebraic closure, there are p^3 genus-2 curves over \mathbb{F}_p . The original CRT algorithm of [15, 17] looped through all p^3 geometric isomorphism classes of curves and tested whether the corresponding endomorphism ring is maximal. This takes time $\tilde{O}(\Delta_0^3\Delta_1^3) + O(\Delta_0^{3/2+\epsilon}\Delta_1^{3/2+\epsilon})$ per CRT prime. Since $\tilde{O}(\Delta_0^{1/2}\Delta_1^{1/2})$ CRT primes are needed, we find a total cost of $\tilde{O}(\Delta_0^{7/2}\Delta_1^{7/2})$, given our choice of ϵ .

The approach of [5] is to search for only one maximal curve, and then to use horizontal isogenies to find the others. With the improvements proposed in this paper (using *all* horizontal isogenies and not just those coming from the type norm, and the improved endomorphism ring computation), we find a cost of $\tilde{O}(\Delta_0^{5/2}\Delta_1^{5/2}) + O(\Delta_0^{3/2+\epsilon}\Delta_1^{3/2+\epsilon})$ per CRT prime. The total cost is then $\tilde{O}(\Delta_0^3\Delta_1^3)$.

With our method, we need to find a curve in the isogeny class where the going-up step yields a maximal curve. Finding a curve in the isogeny class takes time $O(p^{3/2})$. If X is the number of going-up steps we need to try on average, the cost per CRT prime is then $\tilde{O}(X(\Delta_0^{3/2}\Delta_1^{3/2} + \Delta_0^\epsilon\Delta_1^\epsilon))$. At best, $X = O(1)$, and we have a total cost of $\tilde{O}(\Delta_0^2\Delta_1^2)$ from CRT primes. So at best we have a quasiquadratic complexity, while the CRT itself is quasilinear, and thus negligible. We see that we are still far

from quasilinearity achieved by the analytic method. At worst, $X = O(p)$ (number of random tries in the isogeny class until we find a maximal one directly), and we recover the quasicubic complexity of the previous method.

To improve the complexity, there are two possibilities. The first is to increase the probability of success of the going-up method. This requires an algorithm to compute isogenies with cyclic kernels. But even with that, we achieve at most quasiquadratic complexity because the size of the isogeny class is too small compared to the size of the search space. This is the case because the algorithm computes the class polynomials (a scheme of dimension 0) directly from the moduli space of dimension 3 of all abelian surfaces. In contrast, in the elliptic curve case, the algorithm searches a space of dimension 1 for elements of a space of dimension 0. It would be interesting to find convenient subspaces of the moduli space of smaller dimension, and to work over them. One example would be to use Humbert surfaces, which are of dimension 2, and Gundlach invariants, as proposed in [26].

10. EXAMPLES

10.1. Improvements from the going up phase. We first look at improvements due solely to the going up phase. For that we look at Galois examples with class number one; so there is only one maximal curve and the algorithm from Section 6 is not used. We compare the timings to the three examples in [17, Section 9], so to use (almost) the same CRT primes we deactivate the dynamic approach in the prime selection.

Timings and other information for the old and new algorithms are given in Tables 1, 2 and 3, using a 2.39 GHz AMD Opteron with 4 GB of RAM. The first column gives the CRT prime used; an entry in the timings column indicates whether this CRT prime was used in the calculation. The second column lists the ℓ^d -torsion subgroups required to compute whether a curve is maximal; the entry is in bold if there are no rational (ℓ, ℓ) -isogenies (so that we can't go up for this ℓ), and is in italic if there exist (ℓ, ℓ) -isogenies, but they are too expensive to compute. The third column gives the degree of the field extensions where the points of these subgroups live; the degree is italicized when it is so large that computing the ℓ^d -torsion would be too expensive. The fourth column indicates the total number of curves in the isogeny class, computed via the algorithm from [17]. The fifth gives an estimate, obtained as explained in §7.3, for the number of curves from which we can go up, and, in parentheses, the corresponding estimate for the total number of curves in the isogeny class. The last two columns give the timings, in seconds, of the old and new algorithms, split into “Time exploring curves” + “Time spent computing endomorphism rings/Time spent going up”. The old timings are obtained from [17, Tables 1-3].

Note that much less time is spent exploring curves with the new algorithm, due to the going up algorithm. Also note that, even though the going up phase is more complicated, it is still less costly than the computation of the endomorphism rings in the old algorithm, due to the improvements described in Section 4 and the fact that the new version calls it less often.

The trade-offs in the going up step depend on the discriminant of the CM field K . The more CRT primes we need, the bigger the isogenies and the bigger the degrees in the endomorphism ring computations we allow. Note that computing (ℓ, ℓ) -isogenies requires $O(\ell^2)$ operations in the field where the points of the kernel

p	l^d	α_d	# Curves	Estimate	Time (old)	Time (new)
7	2^2	4	7	8	$0.5 + 0.3$	$0 + 0.2$
17	2	1	39	32	$4 + 0.2$	$0 + 0.1$
23	$2^2, 7$	4, 3	49	51	$9 + 2.3$	$0 + 0.2$
71	2^2	4	7	8	$255 + 0.7$	$5.3 + 0.2$
97	2	1	39	32	$680 + 0.3$	$2 + 0.1$
103	$2^2, 17$	4, 16	119	127	$829 + 17.6$	$0.5 + 1$
113	$2^5, 7$	16, 6	1281	877	$1334 + 28.8$	$0.2 + 1.3$
151	$2^2, 7, 17$	4, 3, 16	-	-	0	0
					3162s	13s

$$H_1 = X - 1836660096, \quad H_2 = X - 28343520, \quad H_3 = X - 9762768$$

TABLE 1. Class polynomials for $K = \mathbb{Q}(i\sqrt{2 + \sqrt{2}})$, $\mathfrak{C}(O_K) = \{0\}$.

p	l^d	α_d	# Curves	Estimate	Time (old)	Time (new)
29	3, 23	2, 264	-	-	-	-
53	3, 43	2, 924	-	-	-	-
61	3	2	9	6	$167 + 0.2$	$0.2 + 0.5$
79	3^3	18	81	54	$376 + 8.1$	$0.3 + 0.9$
107	$3^2, 43$	6, 308	-	-	-	-
113	3, 53	1, 52	159	155	$1118 + 137.2$	$0.8 + 25$
131	$3^2, 53$	6, 52	477	477	$1872 + 127.4$	$2.2 + 44.4$
139	3^5	81	?	486	-	$1 + 36.7$
157	3^4	27	243	164	$3147 + 16.5$	-
					6969s	114s

$$H_1 = X - 268435456, \quad H_2 = X + 5242880, \quad H_3 = X + 2015232.$$

TABLE 2. Class polynomials for $K = \mathbb{Q}(i\sqrt{13 + 2\sqrt{29}})$, $\mathfrak{C}(O_K) = \{0\}$.

are defined when ℓ is congruent to 1 (mod 4), but $\tilde{O}(\ell^4)$ when ℓ is congruent to 3 (mod 4). So in the above example, we computed the (109, 109)-isogenies faster than the (23, 23)-isogenies.

10.2. Dihedral examples. Here we illustrate our new CRT algorithm for dihedral fields, for $K = \mathbb{Q}(X)/(X^4 + 13X^2 + 41)$ with $\mathfrak{C}(K) \simeq \{0\}$.

We first compute the class polynomials over \mathbb{Z} using Spallek's invariants, and obtain the following polynomials in 5956 seconds:

$$H_1 = 64X^2 + 14761305216X - 11157710083200000,$$

$$H_2 = 16X^2 + 72590904X - 8609344200000,$$

$$H_3 = 16X^2 + 28820286X - 303718531500.$$

p	l^d	α_d	# Curves	Estimate	Time (old)	Time (new)
7	-	-	1	1	0.3	0 + 0.1
23	13	84	15	2 (16)	9 + 70.7	0.4 + 24.6
53	7	3	7	7	105 + 0.5	7.7 + 0.5
59	2, 5	1, 12	322	48 (286)	164 + 6.4	1.4 + 0.6
83	3, 5	4, 24	77	108	431 + 9.8	2.4 + 1.1
103	<i>67</i>	<i>1122</i>	-	-	-	-
107	7, 13	3, 21	105	8 (107)	963 + 69.3	-
139	5², 7	60, 2	259	9 (260)	2189 + 62.1	-
181	3	1	161	135	5040 + 3.6	4.5 + 0.2
197	5, 109	24, <i>5940</i>	-	-	-	-
199	5²	60	37	2 (39)	10440 + 35.1	-
223	2, 23	1, 11	1058	39 (914)	10440 + 35.1	-
227	109	<i>1485</i>	-	-	-	-
233	5, 7, 13	8, 3, 28	735	55 (770)	11580 + 141.6	88.3 + 29.4
239	7, 109	6, <i>297</i>	-	-	-	-
257	3, 7, 13	4, 6, 84	1155	109 (1521)	17160 + 382.8	-
313	3, 13	1, 14	?	146 (2035)	-	165 + 14.7
373	5, 7	6, 24	?	312	-	183.4 + 3.8
541	2, 7, 13	1, 3, 14	?	294 (4106)	-	91 + 5.5
571	3, 5, 7	2, 6, 6	?	1111 (6663)	-	96.6 + 3.1
					56585s	776s

$$H_1 = 244140625X - 2614061544410821165056,$$

$$H_2 = 390625X + 14651024316825600,$$

$$H_3 = 390625X + 5076177577574400.$$

TABLE 3. Class polynomials for $K = \mathbb{Q}(i\sqrt{29 + 2\sqrt{29}})$, $\mathfrak{C}(O_K) = \{0\}$.

Next we compute them over the real subfield and use the invariants from [36, Appendix 3]. We get

$$H_1 = 256X - 2030994 + 56133\alpha,$$

$$H_2 = 128X + 12637944 - 2224908\alpha,$$

$$H_3 = 65536X - 11920680322632 + 1305660546324\alpha,$$

where α is a root of $X^2 - 3534X + 177505$, so that $O_{K_0^+} = \mathbb{Z}[\alpha]$. This computation took 1401 seconds, so in this case, the speedup due to using better invariants and computing over the real subfield is more than 4-fold.

10.3. CRT for non principal fields. Here is a bigger example. To our knowledge, this is the first example of class number greater than one computed by the CRT method (we see that we are still far away from the analytic method where the latest version can compute class polynomials with degrees in the thousands):

- $K = \mathbb{Q}(X)/(X^4 + 238X^2 + 833)$ cyclic. $\mathfrak{C}(K) \simeq \mathbb{Z}/2\mathbb{Z}$ is generated by $(7, 7)$ -isogenies.
- Primes used: 19, 59, 67, 83, 149, 191, 223, 229, 239, 257, 349, 463, 557, 613, 661, 733, 859, 1039, 1373, 1613, 1657, 1667, 1733, 1753, 1801, 1871, 1879, 2399, 3449, 3469, 3761, 3931, 4259, 4691, 5347, 5381, 6427, 6571, 6781.
- For $p \approx 6000$, we keep p if we expect more than $\frac{p^{3/2}}{32} \approx 15 \times 10^6$ curves. At this size, it takes around 6 seconds to test 10000 curves, so around 2.5 hours are needed for p .
- Total time: 44062 seconds (not using the dynamic approach).

- We only give the first class polynomial:

$$\begin{aligned}
 H_1(X) = & 168451200633545364243594910146286907316572281862280871005795423612829696.X^2 \\
 & +158582528695513934970693031198523489269724119094630145672062735632518026507497890643968.X \\
 & -2014843977961649893357675219372115899170378669590465187558574259942250352955092541374464.
 \end{aligned}$$

11. APPENDIX: COMPLEXITY OF COMPUTING AN (ℓ, ℓ) -ISOGENY

We recall the following proposition from [12]:

Proposition 29. *The complexity of computing an (ℓ, ℓ) -isogeny between two abelian surfaces is $O(\ell^r)$ operations in the field k where the points of the kernel of the isogeny live. We have $r = 2$ when $\ell \equiv 1 \pmod{4}$ and $r = 4$ when $\ell \equiv 3 \pmod{4}$.*

If A is an abelian variety over a finite field \mathbb{F}_q , the following proposition gives a way to compute the degree of the extension over which the points of a maximal isotropic kernel live.

Proposition 30. *Let χ_π be the quartic polynomial satisfied by the Frobenius element for a smooth irreducible genus 2 curve C over \mathbb{F}_p with simple, ordinary Jacobian $J(C)$. For a prime ℓ with $\ell \neq p$, if there exists an \mathbb{F}_p -rational (ℓ, ℓ) -isogeny, then χ_π factors as $\chi_\pi = P\bar{P} \pmod{\ell}$ (where \bar{P} is the conjugate of P under the action $\pi \rightarrow p/\pi$).*

The order of X in $\mathbb{Z}[X]/(\ell, P)$ gives the degree of the extension in which the points of the corresponding kernel live. In particular, if no such decomposition exists, then there is no (ℓ, ℓ) -isogeny.

Proof. Let $K \subset A[\ell]$ be a maximally isotropic rational kernel. Then since π stabilizes K , if P is the characteristic polynomial of π restricted to K , then P divides χ_π . The cofactor is given by the characteristic polynomial of the action of the Verschiebung p/π on K , so the corresponding factor is \bar{P} . \square

REFERENCES

- [1] Eric Bach, *Explicit bounds for primality testing and related problems*, Math. Comp. **55** (1990), no. 191, 355–380. MR 1023756 (91m:11096)
- [2] Juliana Belding, Reinier Bröker, Andreas Enge, and Kristin Lauter, *Computing Hilbert class polynomials*, in van der Poorten and Stein [38], pp. 282–295. MR 2467854 (2009j:11200)
- [3] Gaetan Bisson, *Endomorphism rings in cryptography*, Ph.D. thesis, Technische Universiteit Eindhoven and Institut National Polytechnique de Lorraine, July 2011.
- [4] Gaetan Bisson, Romain Cosset, and Damien Robert, *AVIsogenies, a library for computing isogenies between abelian varieties*, 2012.
- [5] Reinier Bröker, David Gruenewald, and Kristin Lauter, *Explicit CM theory for level 2-structures on abelian surfaces*, Algebra Number Theory **5** (2011), no. 4, 495–528. MR 2870099
- [6] Jan Hendrik Bruinier and Tonghai Yang, *CM-values of Hilbert modular functions*, Invent. Math. **163** (2006), no. 2, 229–288. MR 2207018 (2008b:11053)
- [7] Gabriel Cardona and Jordi Quer, *Field of moduli and field of definition for curves of genus 2*, in Shaska [33], pp. 71–83. MR 2181874 (2006h:14036)
- [8] Robert Carls, David Kohel, and David Lubicz, *Higher-dimensional 3-adic CM construction*, J. Algebra **319** (2008), no. 3, 971–1006. MR 2379090 (2010e:14042)
- [9] Robert Carls and David Lubicz, *A p -adic quasi-quadratic time point counting algorithm*, Int. Math. Res. Not. IMRN (2009), no. 4, 698–735. MR 2480098 (2010c:14020)
- [10] Jean Chaumine, James Hirschfeld, and Robert Rolland (eds.), *Algebraic geometry and its applications*, Series on Number Theory and its Applications, vol. 5, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2008, Dedicated to Gilles Lachaud on his 60th birthday. MR 2484044 (2009h:14003)

- [11] Alina-Carmen Cojocaru, Kristin Lauter, Rachel Pries, and Renate Scheidler (eds.), *WIN—women in numbers*, Fields Institute Communications, vol. 60, American Mathematical Society, Providence, RI, 2011, Research directions in number theory, Including the proceedings of the Banff International Research Station (BIRS) Workshop held in Banff, AB, November 2–7, 2008. [MR 2777796 \(2012g:11005\)](#)
- [12] Romain Cosset and Damien Robert, *Computing (ℓ, ℓ) -isogenies in polynomial time on Jacobians of genus 2 curves*, Cryptology ePrint Archive, Report 2011/143, 2011.
- [13] J.-M. Couveignes, *Linearizing torsion classes in the Picard group of algebraic curves over finite fields*, J. Algebra **321** (2009), no. 8, 2085–2118. [MR 2501511 \(2010e:14019\)](#)
- [14] K. Dickman, *On the frequency of numbers containing prime factors of a certain relative magnitude*, Ark. Mat. Astr. Fys. **22A** (1930), no. 10, 1–14.
- [15] Kirsten Eisenträger and Kristin Lauter, *A CRT algorithm for constructing genus 2 curves over finite fields*, in Rodier and Vladut [32], preprint version at [arXiv:math/0405305 \[math.NT\]](#), pp. 161–176. [MR 2856565](#)
- [16] Andreas Enge and Andrew V. Sutherland, *Class invariants by the CRT method*, in Hanrot et al. [22], pp. 142–156. [MR 2721418 \(2012d:11246\)](#)
- [17] David Freeman and Kristin Lauter, *Computing endomorphism rings of Jacobians of genus 2 curves over finite fields*, in Chaumine et al. [10], pp. 29–66. [MR 2484047 \(2010a:14042\)](#)
- [18] A. Fröhlich (ed.), *Algebraic number fields: L-functions and Galois properties*, Academic Press [Harcourt Brace Jovanovich Publishers], London, 1977. [MR 0437486 \(55 #10416\)](#)
- [19] P. Gaudry, T. Houtmann, D. Kohel, C. Ritzenthaler, and A. Weng, *The 2-adic CM method for genus 2 curves with application to cryptography*, in Lai and Chen [24], pp. 114–129. [MR 2444631 \(2009j:94110\)](#)
- [20] Eyal Z. Goren and Kristin E. Lauter, *Genus 2 curves with complex multiplication*, Int. Math. Res. Not. IMRN (2012), no. 5, 1068–1142. [MR 2899960](#)
- [21] Helen Grundman, Jennifer Johnson-Leung, Kristin Lauter, Adriana Salerno, Bianca Viray, and Erika Wittenborn, *Igusa class polynomials, embeddings of quartic CM fields, and arithmetic intersection theory*, in Cojocaru et al. [11], pp. 35–60. [MR 2777799](#)
- [22] Guillaume Hanrot, François Morain, and Emmanuel Thomé (eds.), *Algorithmic number theory*, Lecture Notes in Computer Science, vol. 6197, Berlin, Springer, 2010. [MR 2766433 \(2011g:11002\)](#)
- [23] J. C. Lagarias and A. M. Odlyzko, *Effective versions of the Chebotarev density theorem*, in Fröhlich [18], pp. 409–464. [MR 0447191 \(56 #5506\)](#)
- [24] Xuejia Lai and Kefei Chen (eds.), *Advances in cryptology—ASIACRYPT 2006*, Lecture Notes in Computer Science, vol. 4284, Berlin, Springer, 2006. [MR 2452006 \(2009e:94091\)](#)
- [25] Kristin Lauter and Bianca Viray, *An arithmetic intersection formula for denominators of Igusa class polynomials*, 2012. [arXiv:1210.7841 \[math.NT\]](#)
- [26] Kristin Lauter and Tonghai Yang, *Computing genus 2 curves from invariants on the Hilbert moduli space*, J. Number Theory **131** (2011), no. 5, 936–958. [MR 2772480 \(2012e:11110\)](#)
- [27] H. W. Lenstra, Jr., J. Pila, and Carl Pomerance, *A hyperelliptic smoothness test. II*, Proc. London Math. Soc. (3) **84** (2002), no. 1, 105–146. [MR 1863397 \(2003f:11190\)](#)
- [28] David Lubicz and Damien Robert, *Computing isogenies between abelian varieties*, 2010, to appear in *Compositio Mathematica*. [arXiv:1001.2016 \[math.AG\]](#)
- [29] Jean-François Mestre, *Construction de courbes de genre 2 à partir de leurs modules*, in Mora and Traverso [30], pp. 313–334. [MR 1106431 \(92g:14022\)](#)
- [30] Teo Mora and Carlo Traverso (eds.), *Effective methods in algebraic geometry*, Progress in Mathematics, vol. 94, Birkhäuser Boston Inc., Boston, MA, 1991, Papers from the symposium (MEGA-90) held in Castiglione della Pescaia, April 17–21, 1990. [MR 1106410 \(91m:14003\)](#)
- [31] Damien Robert, *Fonctions θ et applications à la cryptographie*, Ph.D. thesis, Université Henri Poincaré — Nancy 1, July 2010.
- [32] François Rodier and Serge Vladut (eds.), *Arithmetics, geometry, and coding theory (AGCT 2005)*, Séminaires et Congrès [Seminars and Congresses], vol. 21, Société Mathématique de France, Paris, 2010, Papers from the conference held in Marseilles, September 26–30, 2005. [MR 2866893 \(2012h:14002\)](#)
- [33] Tanush Shaska (ed.), *Computational aspects of algebraic curves*, Lecture Notes Series on Computing, vol. 13, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2005, Papers from the conference held at the University of Idaho, Moscow, ID, May 26–28, 2005. [MR 2182657 \(2006e:14003\)](#)

- [34] Goro Shimura, *Abelian varieties with complex multiplication and modular functions*, Princeton Mathematical Series, vol. 46, Princeton University Press, Princeton, NJ, 1998. [MR 1492449 \(99e:11076\)](#)
- [35] Anne-Monika Spallek, *Kurven vom Geschlecht 2 und ihre Anwendung in Public-Key-Kryptosystemen*, Ph.D. thesis, Universität Gesamthochschule Essen, July 1994.
- [36] Theodorus Cornelis Streng, *Complex multiplication of abelian surfaces*, Ph.D. thesis, Universiteit Leiden, June 2010.
- [37] Andrew V. Sutherland, *Computing Hilbert class polynomials with the Chinese remainder theorem*, *Math. Comp.* **80** (2011), no. 273, 501–538. [MR 2728992 \(2011k:11177\)](#)
- [38] Alfred J. van der Poorten and Andreas Stein (eds.), *Algorithmic number theory*, Lecture Notes in Computer Science, vol. 5011, Berlin, Springer, 2008. [MR 2467865 \(2009h:11002\)](#)
- [39] Paul van Wamelen, *Examples of genus two CM curves defined over the rationals*, *Math. Comp.* **68** (1999), no. 225, 307–320. [MR 1609658 \(99c:11079\)](#)
- [40] Annegret Weng, *Constructing hyperelliptic curves of genus 2 suitable for cryptography*, *Math. Comp.* **72** (2003), no. 241, 435–458 (electronic). [MR 1933830 \(2003i:14029\)](#)
- [41] Tonghai Yang, *An arithmetic intersection formula on Hilbert modular surfaces*, *Amer. J. Math.* **132** (2010), no. 5, 1275–1309. [MR 2732347 \(2012a:11078\)](#)
- [42] ———, *Arithmetic intersection on a Hilbert modular surface and the Faltings height*, 2010. [arXiv:1008.1854 \[math.NT\]](#)

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WA 98052, USA.
E-mail address: klauter@microsoft.com

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WA 98052, USA.
Current address: INRIA Bordeaux–Sud-Ouest, 200 avenue de la Vieille Tour, 33405 Talence cedex, France
E-mail address: damien.robert@inria.fr