THE OPEN BOOK SERIES 2

ANTS XIII
Proceedings of the Thirteenth Algorithmic Number Theory Symposium

Counting points on genus-3 hyperelliptic curves with explicit real multiplication

Simon Abelard, Pierrick Gaudry, and Pierre-Jean Spaenlehauer

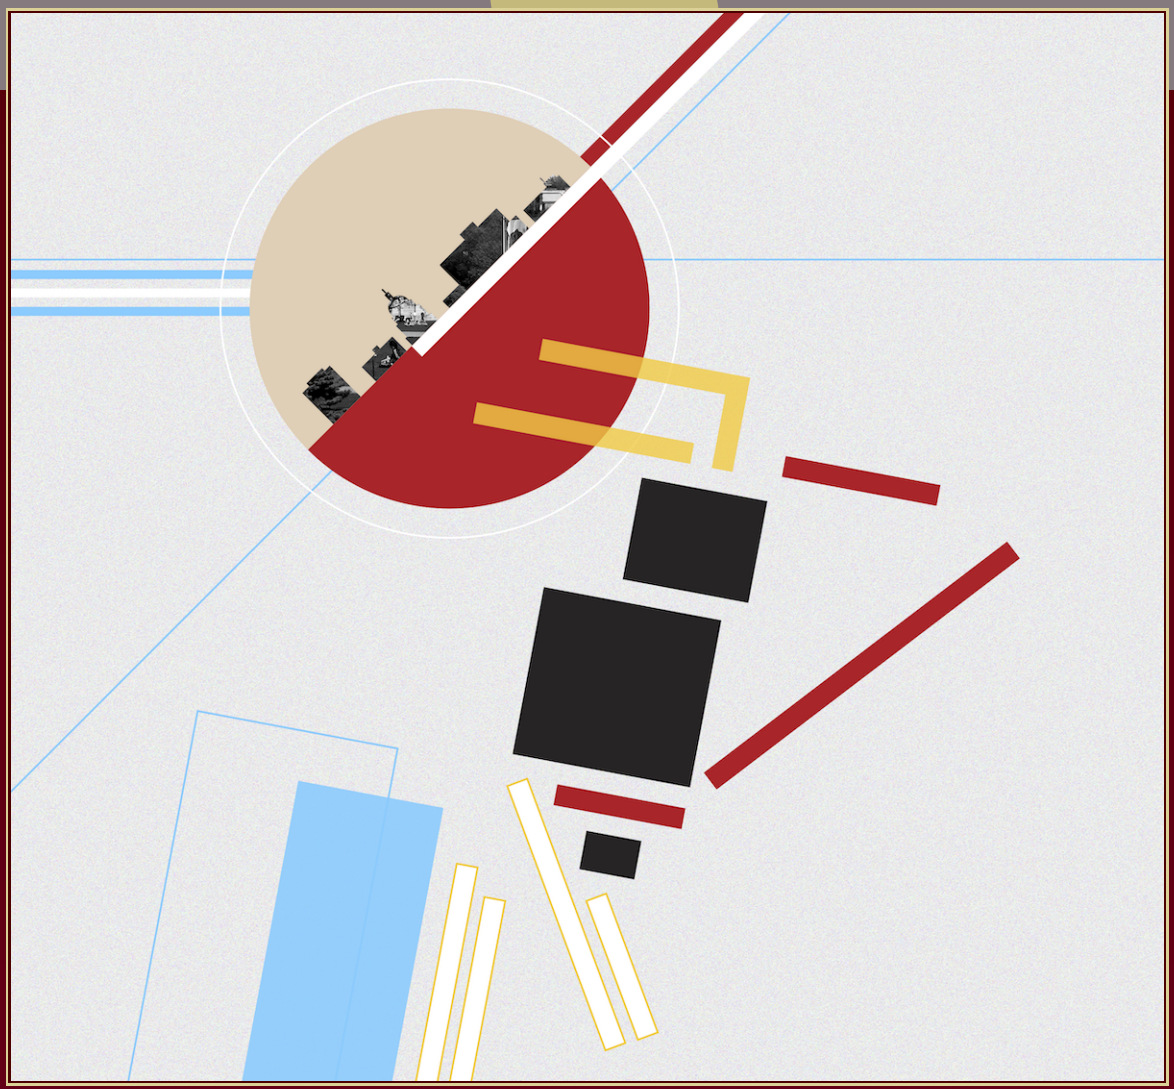

msp

THE OPEN BOOK SERIES **2** (2019)
Thirteenth Algorithmic Number Theory Symposium
dx.doi.org/10.2140/obs.2019.2.1

msp

Counting points on genus-3 hyperelliptic curves with explicit real multiplication

Simon Abelard, Pierrick Gaudry, and Pierre-Jean Spaenlehauer

We propose a Las Vegas probabilistic algorithm to compute the zeta function of a genus-3 hyperelliptic curve defined over a finite field $\mathbb{F}_q$, with explicit real multiplication by an order $\mathbb{Z}[\eta]$ in a totally real cubic field. Our main result states that this algorithm requires an expected number of $\widetilde{O}((\log q)^6)$ bit-operations, where the constant in the $\widetilde{O}(\,)$ depends on the ring $\mathbb{Z}[\eta]$ and on the degrees of polynomials representing the endomorphism η. As a proof-of-concept, we compute the zeta function of a curve defined over a 64-bit prime field, with explicit real multiplication by $\mathbb{Z}[2\cos(2\pi/7)]$.

1. Introduction

Since the discovery of Schoof's algorithm [25], the problem of efficiently computing zeta functions of curves defined over finite fields has attracted a lot of attention, as its applications range from the construction of cryptographic curves to testing conjectures in number theory. We focus on the problem of computing the zeta function of a hyperelliptic curve $\mathcal{C}$ of genus 3 defined over a finite field $\mathbb{F}_q$ using ℓ-adic methods, in the spirit of Schoof's algorithm and its generalizations [23; 17; 2]. Although these methods are polynomial with respect to $\log q$, the exponents in the best known complexity bounds grow quickly with the genus. Another line of research is to use p-adic methods [18; 24; 7; 14], which are polynomial in the genus but exponential in the size of the characteristic of the underlying finite field. Variants of these methods [19; 15; 16] allow us to count the points of a curve defined over the rationals modulo many primes in average polynomial time, which is especially relevant when experimenting with the Sato–Tate conjecture.

The aim of this paper is to show — both with theoretical proofs and practical experiments — that the complexity of ℓ-adic methods for genus-3 hyperelliptic curves can be dramatically decreased as soon as an explicitly computable noninteger endomorphism $\eta \in \mathrm{End}(\mathrm{Jac}(\mathcal{C}))$ is known. More precisely, we say that a curve $\mathcal{C}$ has *explicit real multiplication* by $\mathbb{Z}[\eta]$ if the subring $\mathbb{Z}[\eta] \subset \mathrm{End}(\mathrm{Jac}(\mathcal{C}))$ is isomorphic to an order in a totally real cubic number field, and if we have explicit formulas describing

MSC2010: primary 11G20; secondary 11M38, 11Y16.
Keywords: point-counting, hyperelliptic curves, Schoof's algorithm, real multiplication.

$\eta(P-\infty)$ for some fixed base point ∞ and a generic point P of C. By explicit formulas, we mean polynomials $(\eta_i^{(u)}(x,y))_{i\in\{0,1,2,3\}}$ and $(\eta_i^{(v)}(x,y))_{i\in\{0,1,2,3\}}$ in $\mathbb{F}_q[x,y]$, such that, when $\mathcal{C}$ is given in odd-degree Weierstrass form, the Mumford coordinates of $\eta((x,y)-\infty)$ are

$$\left\langle \sum_{i=0}^{3} \eta_i^{(u)}(x,y)X^i, \sum_{i=0}^{2} \left(\eta_i^{(v)}(x,y)/\eta_3^{(v)}(x,y)\right)X^i \right\rangle,$$

where (x,y) is the generic point of the curve. In cases where $\mathcal{C}$ does not have an odd-degree Weierstrass model, we can work in an extension of degree at most 8 of the base field in order to ensure the existence of a rational Weierstrass point.

The influence of real multiplication (RM) on the complexity of point counting was investigated for genus-2 curves in [11], where the complexity was lowered from $\widetilde{O}((\log q)^8)$ [13] to $\widetilde{O}((\log q)^5)$. For genus-2 curves, another related active line of research is to mimic the improvement of Elkies and Atkin by using modular polynomials [3]. However, the main difficulty of this method is to precompute the modular polynomials, which are much larger than their genus-1 counterparts.

Our main result is the following theorem.

Theorem 1. *Let $\mathcal{C}$ be a genus-3 hyperelliptic curve defined over a finite field $\mathbb{F}_q$ having explicit real multiplication by $\mathbb{Z}[\eta]$, where $\eta \in \mathrm{End}(\mathrm{Jac}(C))$. We assume that $\mathcal{C}$ is given by an odd-degree Weierstrass equation $Y^2 = f(X)$. The characteristic polynomial of the Frobenius endomorphism on the Jacobian of $\mathcal{C}$ can be computed with a Las Vegas probabilistic algorithm in expected time bounded by $c(\log q)^6(\log\log q)^k$, where k is an absolute constant and c depends only on the degrees of the polynomials $\eta_i^{(u)}$ and $\eta_i^{(v)}$ and on the ring $\mathbb{Z}[\eta]$.*

In this paper, we use the notation $\widetilde{O}(\,)$ as a shorthand for complexity statements hiding polylogarithmic terms: the complexity in the theorem would be abbreviated $\widetilde{O}((\log q)^6)$. We insist on the fact that all the $O(\,)$ and the $\widetilde{O}(\,)$ notation used throughout the paper should be understood up to a multiplicative constant which may depend on the ring $\mathbb{Z}[\eta]$ and on the degrees of the polynomials $\eta_i^{(u)}$ and $\eta_i^{(v)}$. There are natural families of curves for which these degrees are bounded by an absolute constant and for which $\mathbb{Z}[\eta]$ is fixed: reductions at primes (of good reduction) of a hyperelliptic curve with explicit RM defined over a number field.

As in Schoof's algorithm and its generalizations in [23; 17; 2], the ℓ-adic approach consists in computing the characteristic polynomial of the Frobenius endomorphism by computing its action on the ℓ-torsion of the Jacobian of the curve for sufficiently many ℓ. In order to prove the claimed complexity bound, we consider primes $\ell \in \mathbb{Z}$ such that $\ell\mathbb{Z}[\eta]$ splits as a product $\mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$ of prime ideals. Computing the kernels of endomorphisms α_i in each $\mathfrak{p}_i$ provides us with an algebraic representation of the ℓ-torsion $\mathrm{Jac}(\mathcal{C})[\ell] \subset \mathrm{Ker}\,\alpha_1 + \mathrm{Ker}\,\alpha_2 + \mathrm{Ker}\,\alpha_3$. Then, we compute from this representation integers $a, b, c \in \mathbb{Z}/\ell\mathbb{Z}$ such that the sum $\pi + \pi^\vee$ of the Frobenius endomorphism and its dual equals $a + b\eta + c\eta^2 \bmod \ell$. Once enough modular information is known, the values of a, b, c such that $\pi + \pi^\vee = a + b\eta + c\eta^2$ are recovered via the Chinese remainder theorem and the coefficients of the characteristic polynomial of the Frobenius can be directly expressed in terms of a, b and c. In fact, in practice we do not have to restrict to split

primes: any partial factorization of $\ell\mathbb{Z}[\eta]$ provides some modular information on $a, b, c \bmod \ell$. We give an example with a ramified prime in Section 7.1, but on the theoretical side, considering nonsplit primes does not improve the asymptotic complexity.

The cornerstone of the complexity analysis is the cost of the computation of the kernels of endomorphisms, which is done by solving a polynomial system. Using resultant-based elimination techniques and degree bounds on Cantor's polynomials, we prove that we can solve these equations in time quadratic in the number of solutions, which leads to the claimed complexity bound. For practical computations, we replace the resultants by Gröbner bases and we retrieve modular information only for small ℓ to speed up an exponential collision search which can be massively run in parallel. Although using Gröbner bases seems to be more efficient in practice, we do not see any hope of proving with rigorous arguments that it is asymptotically competitive.

As a proof-of-concept, we have implemented our algorithm and we provide experimental results. In particular, we were able to compute the zeta function of a genus-3 hyperelliptic curve with explicit RM defined over $\mathbb{F}_p$ with $p = 2^{64} - 59$. To our knowledge, the largest genus-3 computation that had been achieved previously was the computation of the zeta function of a hyperelliptic curve defined over $\mathbb{F}_p$ with $p = 2^{61} - 1$, done by Sutherland [27] using generic group methods.

Examples of curves with RM are given by modular curves. For instance, the genus-3 curve $y^2 = x^7 + 3x^6 + 2x^5 - x^4 - 2x^3 - 2x^2 - x - 1$ is a quotient of $X_0(284)$ and therefore has real multiplication by an element of $\mathbb{Q}[x]/(x^3 - 3x - 1)$. This follows from the properties of the Hecke operators as explained in [26, Chapter 7]. Based on this theory, algorithms for constructing such curves are explained in [10]; however the explicit expression for the real endomorphism is not given. We expect that tracking the Hecke correspondences along their construction, and using techniques like those in [29] to reconstruct the rational fractions describing the real endomorphism, could solve this question. In any case, these are only isolated points in the moduli space. Larger families are obtained from cyclotomic covering. This line of research has produced several families of hyperelliptic genus-3 curves having explicit RM by $\mathbb{Z}[2\cos(2\pi/7)]$. In particular, explicit such families are given in [22] and [28], and explicit formulas for their RM endomorphisms are obtained in [20]. We use the 1-dimensional family of curves from [28, Theorem 1 with $p = 7$] for our experiments. Other families of genus-3 curves (but not necessarily hyperelliptic) with RM have been made explicit in [5, Chapter 2], following [9]. We would like to point out that within the moduli space of complex polarized abelian varieties of dimension 3, those with RM by a fixed order in a cubic field form a moduli space of codimension 3 [21, Section 9.2]. Since Jacobians of hyperelliptic curves form a codimension-1 space, we would expect the moduli space of hyperelliptic curves of genus 3 with RM by a given cubic order to have dimension 2.

We finally briefly mention how our algorithm and analysis could be extended in several directions. First, the complexity analysis leads, with small modifications, to a point-counting algorithm for general genus-3 hyperelliptic curves (i.e., without RM) with complexity $\widetilde{O}((\log q)^{14})$. Second, if the curve is not hyperelliptic, the main difficulty is to define analogues of Cantor's division polynomials and get bounds on their degrees. Without them, it is still possible to use an explicit group law to derive a polynomial

system for the kernel of an endomorphism, but getting a proof for its degree would require a different approach to the one we took. Still, the complexities with or without RM are expected to remain the same for plane quartics as for genus-3 hyperelliptic curves. Third, if we go to higher genus hyperelliptic curves with RM, the main difficulty in extending our approach is in the complexity estimate of the polynomial system solving, because resultant-based approaches are not competitive when the number of variables grows, and a tedious analysis like that in [1] seems to be necessary.

The article is organized as follows. Section 2 gives a bird's-eye view of our algorithm, along with a complexity analysis relying on the technical results detailed in Sections 3 to 6. Practical experiments are presented in Section 7.

2. Overview of the algorithm

Let $\mathcal{C}$ be a genus-3 hyperelliptic curve over a finite field $\mathbb{F}_q$ with explicit RM, and let η be the given explicit endomorphism. We denote by μ_0, μ_1, μ_2 the coefficients of the minimal polynomial $T^3+\mu_2T^2+\mu_1T+\mu_0$ of η over $\mathbb{Q}$.

2.1. ***Bounds.*** The characteristic polynomial of the Frobenius endomorphism π is of the form $\chi_\pi(T) = T^6 - \sigma_1T^5 + \sigma_2T^4 - \sigma_3T^3 + q\sigma_2T^2 - q^2\sigma_1T + q^3$, and Weil's bounds give

$$|\sigma_1| \leq 6\sqrt{q}, \quad |\sigma_2| \leq 15q, \quad |\sigma_3| \leq 20q^{3/2}.$$

In order to take advantage of the explicit RM, we consider the endomorphism $\psi = \pi + \pi^\vee$, for which we can derive the real Weil's polynomial $\chi_\psi(T) = T^3 - \sigma_1T^2 + (\sigma_2 - 3q)T - (\sigma_3 - 2q\sigma_1)$, which corresponds to the characteristic polynomial of ψ viewed as an element of the real subfield of $\mathrm{End}(\mathrm{Jac}(C)) \otimes \mathbb{Q}$. The endomorphism ψ belongs to the ring of integers of $\mathbb{Q}(\eta)$. The ring $\mathbb{Z}[\eta]$ might be a proper suborder of the ring of integers, so let us call Δ its index, so that ψ can be written $\psi = a + b\eta + c\eta^2$, where a, b, c are rationals with a denominator that divides Δ. By computing formally the characteristic polynomial of $a + b\eta + c\eta^2$ in $\mathbb{Q}(\eta)$ and by equating it with the expression for the real Weil's polynomial $\chi_\psi(T)$, we obtain a direct way to compute σ_1, σ_2 and σ_3 in terms of a, b, c:

$$\begin{aligned}
\sigma_1 &= 3a - b\mu_2 - 2c\mu_1 + c\mu_2^2, \\
\sigma_2 - 3q &= 3a^2 - 2ab\mu_2 + 2ac(\mu_2^2 - 2\mu_1) + b^2\mu_1 + 3bc\mu_0 - bc\mu_1\mu_2 - c^2(2\mu_0\mu_2 + \mu_1^2), \\
\sigma_3 - 2q\sigma_1 &= a^3 - a^2b\mu_2 + a^2c(\mu_2^2 - 2\mu_1) + ab^2\mu_1 + abc(3\mu_0 - \mu_1\mu_2) \\
&\qquad + ac^2(\mu_1^2 - 2\mu_0\mu_2) - b^3\mu_0 + b^2c\mu_0\mu_2 - bc^2\mu_0\mu_1 + c^3\mu_0^2.
\end{aligned} \tag{1}$$

In Section 4, it is shown that the coefficients a, b and c can be bounded in $O(\sqrt{q})$. More precisely, we denote by C_{abc} a constant that depends only on η such that their absolute values are bounded by $C_{abc}\sqrt{q}$. Since these bounds are much smaller than the bounds for σ_1, σ_2, σ_3, it makes sense to design an algorithm that reconstructs these coefficients of ψ instead of the coefficients of χ_π as in the classical Schoof algorithm, and this is what we are going to do later on.

Another important bound that we need concerns the size of small elements that can be found in ideals of $\mathbb{Z}[\eta]$. Let ℓ be a prime that splits completely in $\mathbb{Z}[\eta]$, so that we can write $\ell = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$, where the $\mathfrak{p}_i$'s are distinct prime ideals of norm ℓ. In Section 5, it is shown that each $\mathfrak{p}_i$ contains a nonzero element $\alpha_i = a_i + b_i\eta + c_i\eta^2$, where a_i, b_i and c_i are integers and are bounded in absolute value by $O(\ell^{1/3})$.

2.2. ***Algorithms.*** The general RM point counting algorithm is Algorithm 1. We give a description of it, allowing some black-box primitives that will be detailed in dedicated sections. As mentioned above, we will work with the a, b, c coefficients of the ψ endomorphism. More precisely, we compute their values modulo sufficiently many completely split primes ℓ until we can deduce their values from the bounds of Lemma 5 by the Chinese remainder theorem, taking into account their potential denominator Δ. Then the coefficients of χ_π are deduced by (1).

We now explain how the algorithm works for a given split ℓ. First its decomposition as a product of prime ideals $\ell\,\mathbb{Z}[\eta] = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$ is computed, and for each prime ideal $\mathfrak{p}_i$, a nonzero element α_i of $\mathfrak{p}_i$ is found with a small representation $\alpha_i = a_i + b_i\eta + c_i\eta^2$ as in Lemma 6. In fact, $\mathfrak{p}_i$ is not necessarily principal and α_i need not generate $\mathfrak{p}_i$. The kernel of α_i is denoted by $J[\alpha_i]$ and it contains a subgroup G_i isomorphic to $\mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z}$, since the norm of α_i is a multiple of ℓ. The two-element representation $(\ell, \eta - \lambda_i)$ of the ideal $\mathfrak{p}_i$ implies that λ_i is an eigenvalue of η regarded as an endomorphism of $J[\ell] \cong (\mathbb{Z}/\ell\mathbb{Z})^6$.

On $G_i \subset J[\alpha_i]$, the endomorphism η acts as the multiplication by λ_i. Therefore, $\psi = a + b\eta + c\eta^2$ also acts as a scalar multiplication on this 2-dimensional space, and we write $k_i \in \mathbb{Z}/\ell\mathbb{Z}$ as the corresponding eigenvalue: for any D_i in G_i, we have $\psi(D_i) = k_i D_i$. On the other hand, from the definition of ψ, it follows that $\psi\pi = \pi^2 + q$. Therefore, if such a D_i is known, we can test which value of $k_i \in \mathbb{Z}/\ell\mathbb{Z}$ satisfies

$$k_i\pi(D_i) = \pi^2(D_i) + q\,D_i. \tag{2}$$

Since ℓ is a prime and D_i is of order exactly ℓ, this is also the case for $\pi(D_i)$. Finding k_i can then be seen as a discrete logarithm problem in the subgroup of order ℓ generated by $\pi(D_i)$; hence the solution is unique. Equating the two expressions for ψ, we get explicit relations between a, b, c modulo ℓ:

$$a + b\lambda_i + c\lambda_i^2 \equiv k_i \bmod \ell.$$

Therefore we have a linear system of three equations in three unknowns, the determinant of which is the Vandermonde determinant of the λ_i, which are distinct by hypothesis. Hence the system can be solved and it has a unique solution modulo ℓ.

It remains to show how to construct a divisor D_i in G_i, i.e., an element of order ℓ in the kernel $J[\alpha_i]$. Since an explicit expression of η as an endomorphism of the Jacobian of $\mathcal{C}$ is known, an explicit expression can be deduced for α_i, using the explicit group law. The coordinates of the elements of this kernel are solutions of a polynomial system that can be directly derived from this expression of α_i. Using standard techniques, it is possible to find the solutions of this system in a finite extension of the base field (of degree bounded by the degree of the ideal generated by the system, i.e., in $O(\ell^2)$), from which divisors in $J[\alpha_i]$ can be constructed. Multiplying by the appropriate cofactor, we can reach all the elements of G_i, but we stop as soon as we get a nontrivial one.

Algorithm 1: Overview of our RM point-counting algorithm

Data: q an odd prime power, and $f \in \mathbb{F}_q[X]$ a monic squarefree polynomial of degree 7 such that the curve $Y^2 = f(X)$ has explicit RM by $\mathbb{Z}[\eta]$.
Result: the characteristic polynomial $\chi_\pi \in \mathbb{Z}[T]$ of the Frobenius endomorphism on the Jacobian J of the curve.

$R \leftarrow 1$;
while $R \leq 2\,\Delta\, C_{abc}\,\sqrt{q} + 1$ **do**
  pick the next prime ℓ that satisfies conditions (C1) to (C4);
  compute the ideal decomposition $\ell\,\mathbb{Z}[\eta] = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$, corresponding to the eigenvalues $\lambda_1,\ \lambda_2,\ \lambda_3$ of η in $J[\ell]$;
  for $i \leftarrow 1$ **to** 3 **do**
    compute a small element α_i of $\mathfrak{p}_i$ as in Lemma 6;
    compute a nonzero element D_i of order ℓ in $J[\alpha_i]$;
    find the unique $k_i \in \mathbb{Z}/\ell\mathbb{Z}$ such that $k_i\pi(D_i) = \pi^2(D_i) + q\,D_i$;
  find the unique triple (a, b, c) in $(\mathbb{Z}/\ell\mathbb{Z})^3$ such that $a + b\lambda_i + c\lambda_i^2 = k_i$, for i in $\{1, 2, 3\}$;
  $R \leftarrow R \cdot \ell$;
reconstruct (a, b, c) using the Chinese remainder theorem;
deduce χ_π from (1);

We summarize the conditions that must be satisfied by the primes ℓ that we work with:

(C1) ℓ must be different from the characteristic of the base field.

(C2) ℓ must be coprime to the discriminant of the minimal polynomial of η.

(C3) There must exist $\alpha_i \in \mathfrak{p}_i$ as in Lemma 6 with norm nondivisible by ℓ^3 for $i \in \{1, 2, 3\}$.

(C4) The ideal $\ell\mathbb{Z}[\eta]$ must split completely.

The first 3 conditions eliminate only a finite number of ℓ's that depends only on η, while the last one eliminates a constant proportion. The condition (C3) implies that there is a unique subgroup G_i of order ℓ^2 in $J[\alpha_i]$ (our description of the algorithm could actually be adapted to handle the cases where this is not true).

Algorithm 1 is a very natural extension of the one described in [11] for genus-2 curves with RM. In [11], the action of the real endomorphism $\psi = \pi + \pi^\vee$ is studied on subspaces $J[\mathfrak{p}_i]$ of the ℓ-torsion, and the corresponding eigenvalues are collected and used to reconstruct information modulo ℓ. In genus 3, we have three such 2-dimensional subspaces and eigenvalues to compute and recombine instead of two in genus 2. The main differences between the present work and [11] are the way the ℓ-torsion elements are constructed with polynomial systems and the bounds on the coefficients of ψ. In both cases, going from dimension 2 to 3 is not immediate.

2.3. ***Complexity analysis.*** The field $\mathbb{Q}(\eta)$ is of degree 3, so its Galois group has order at most 6 and by Chebotarev's density theorem the density of primes that split completely is at least $1/6$. Therefore the

main loop is done $O(\log q/\log\log q)$ times, with primes ℓ that are in $O(\log q)$. All the steps that take place in the number field take a negligible time. For instance, a small generator like the one in Lemma 6 can be found by exhaustive search: only $O(\ell)$ trials are needed since we are searching over all elements of the form $a+b\eta+c\eta^2$, with $|a|, |b|, |c|$ in $O(\ell^{1/3})$.

The bottleneck of the algorithm is the computation of a nonzero element of order ℓ in the kernel $J[\alpha_i]$ of α_i. This part will be treated in detail in Section 3, where it is shown to be feasible in $\widetilde{O}(\ell^4)$ operations in $\mathbb{F}_q$. The output is a divisor D_i of order ℓ in $J[\alpha_i]$ that is defined over an extension field $\mathbb{F}_{q^\delta}$, where δ is in $O(\ell^2)$.

In order to check (2), we first need to compute $\pi(D_i)$ and $\pi^2(D_i)$ which amounts to raising the coordinates to the q-th power. The cost is $\widetilde{O}(\ell^2\log q)$ operations in $\mathbb{F}_q$. Then, each Jacobian operation in the group generated by $\pi(D_i)$ costs $\widetilde{O}(\ell^2)$ operations in the base field, and we need $O(\sqrt{\ell})$ of them to solve the discrete logarithm problem given by (2). The overall cost of finding k_i, once D_i is known is therefore $\widetilde{O}(\ell^2(\sqrt{\ell}+\log q))$ operations in $\mathbb{F}_q$.

Finally, the amount of work performed for each ℓ is $\widetilde{O}(\ell^2(\ell^2+\log q))$ operations in the base field $\mathbb{F}_q$. Summing up for all the primes, and taking into account the cost of the operations in $\mathbb{F}_q$, we obtain a global bit-complexity of $\widetilde{O}((\log q)^6)$.

3. Computing kernels of endomorphisms

3.1. ***Modeling the kernel computation by a polynomial system.*** Let α be an explicit endomorphism of degree $O(\ell^2)$ on the Jacobian of $\mathcal{C}$, which satisfies the properties of Lemma 6. In particular, α vanishes on a subspace of $J[\ell]$. We want to compute a triangular polynomial system that describes the kernel $J[\alpha]$ of α. This will provide us with a nice description of a subgroup of the ℓ-torsion on which we will be able to test the action of $\psi=\pi+\pi^\vee$ and deduce a, b, c such that $\psi = a+b\eta+c\eta^2 \bmod \ell$.

We first model $J[\alpha]$ by a system of polynomial equations that we will then put in triangular form. To do so, we consider a generic divisor $D=P_1+P_2+P_3-3\infty$, where P_i is an affine point of $\mathcal{C}$ of coordinates (x_i, y_i). We then write $\alpha(D)=0$, i.e, $\alpha(P_1-\infty)+\alpha(P_2-\infty)=-\alpha(P_3-\infty)$. Generically, we expect each $\alpha(P_i-\infty)$ to be of weight 3, and we write $\langle u_i, v_i\rangle$ for its Mumford form. We derive our equations by computing the Mumford form $\langle u_{12}, v_{12}\rangle$ of $\alpha(P_1-\infty)+\alpha(P_2-\infty)$ and then writing coefficient-wise the conditions $u_{12}=u_3$ and $v_{12}=-v_3$. The case where the genericity conditions are not satisfied is discussed at the end of the section.

Similarly to the Schoof–Pila algorithm, we define polynomials — which are equivalent to Cantor's division polynomials — by the formulas

$$u_{12}(X)=X^3+\sum_{i=0}^{2}\frac{\tilde{d}_i(x_1,x_2,y_1,y_2)}{\tilde{d}_3(x_1,x_2)}X^i, \qquad v_{12}(X)=\sum_{i=0}^{2}\frac{\tilde{e}_i(x_1,x_2,y_1,y_2)}{\tilde{e}_3(x_1,x_2)}X^i,$$

$$u_3(X)=X^3+\sum_{i=0}^{2}\frac{d_i(x_3)}{d_3(x_3)}X^i, \qquad v_3(X)=y_3\sum_{i=0}^{2}\frac{e_i(x_3)}{e_3(x_3)}X^i.$$

Lemma 2. *For any $i \in \{1, 2, 3\}$, the degrees of $\tilde{d}_i$, $\tilde{e}_i$, d_i and e_i are in $O(\ell^{2/3})$.*

Proof. Let us first remark that the $\tilde{d}_i$'s and $\tilde{e}_i$'s are obtained after adding two divisors $\langle u_1, v_1\rangle$ and $\langle u_2, v_2\rangle$ such that the coefficients of the u_i and v_i are, respectively, the d_j/d_3 and $y_i e_j/e_3$ evaluated at x_i. Thus, since this application of the group law involves a number of operations that is bounded independently of ℓ and q, the degree stays within a constant multiplicative factor, which is captured by the $O(\,)$. Therefore it is enough to prove the result for the d_i's and e_i's.

Since the endomorphism α satisfies the properties of Lemma 6, it is a linear combination of 1, η and η^2 with coefficients of size $O(\ell^{1/3})$. Using the same argument about the group law, we can further reduce our proof to the case where $\alpha = n\eta^k$, with $k \in \{0, 1, 2\}$ and n an integer in $O(\ell^{1/3})$. But once again, η^k does not depend on ℓ so that, provided we can prove that Cantor's n-division polynomials have degrees in $O(n^2)$, we have proven that $n\eta^k(P - \infty) = \eta^k(n(P - \infty))$ have coefficients whose degrees are in $O(n^2)$, and then so does $\alpha(P - \infty)$. This quadratic bound on the degrees of Cantor's division polynomials is proven in Lemma 8 of Section 6 and the result follows. □

3.2. ***Solving the system with resultants.*** Typical tools for solving a polynomial system are the F4 algorithm, methods based on geometric resolution, or homotopy techniques. To obtain reasonable complexity bounds, they all require some knowledge of the properties of the system, and this might be hard to provide. Since we have a system in essentially three variables (in fact, there are six variables $x_1, x_2, x_3, y_1, y_2, y_3$, but the y_i variables can be directly eliminated by using the equation defining the curve), we prefer to stick to an approach based on resultants. It ends up having a complexity that is quasiquadratic in the degree of the ideal, which is the best that can be hoped for anyway for all of the advanced techniques, and the complexity analysis requires only elementary tools. A complication that can occur with resultants is that $\mathrm{Res}_x(f, g)$ is identically zero when f and g have a nonconstant GCD. This is not a problem in our case since we can divide polynomials f and g by their GCD, by factoring them at the cost of $O(\max(\deg(f), \deg(g))^\omega)$ field operations — where $\omega \leq 3$ is the exponent of linear algebra — using the bivariate recombination methods in [4] (the trivariate case can be reduced to the bivariate case by using the techniques in [31, Section 21.2]). In what follows, the complexities of computing the resultants are larger than $O(\max(\deg(f), \deg(g))^\omega)$, so we can forget about this complication. We also note that since the system is symmetric with respect to x_1 and x_2, it may be possible to decrease the degrees by rewriting the system in terms of elementary symmetric polynomials in x_1 and x_2; however, we do not consider this symmetrization process in the analysis since it may only gain a constant factor in the complexity.

Following our modeling, the equality of the u-coordinates gives three equations

$$\tilde{d}_i(x_1, x_2, y_1, y_2)d_3(x_3) = \tilde{d}_3(x_1, x_2)d_i(x_3), \quad \text{for } i \in \{0, 1, 2\}, \tag{3}$$

of degree $O(\ell^{2/3})$ in the x_i's. By computing resultants with the equations $y_i^2 = f(x_i)$, we derive three equations $E_i(x_1, x_2, x_3) = 0$ whose degrees are still in $O(\ell^{2/3})$.

We then eliminate x_1 by computing three trivariate resultants R_i (between the two equations E_j with $j \neq i$). We get three equations $R_i(x_2, x_3) = 0$ of degrees $O(\ell^{4/3})$ within a complexity of $\widetilde{O}(\ell^{10/3})$ field operations, as proven in Lemma 4 below.

Then, we compute bivariate resultants S_i (between the two equations R_j with $j \neq i$) to eliminate x_2. From Lemma 3, we get three univariate equations $S_i(x_3) = 0$ of degree bounded by $O(\ell^{8/3})$ for a complexity of $\widetilde{O}(\ell^4)$ field operations. And we compute the polynomial $S(x_3)$ as the GCD of the $S_i(x_3)$, which belongs to the ideal defined by our original system.

The bound on the degree of S is much larger than $\ell^2 - 1$, the expected degree of the kernel. Although we can expect the actual degree to be in $O(\ell^2)$, we need to add the constraints coming from the v-coordinates to be able to prove it.

The polynomial system coming from $v_{12} = -v_3$ has the same characteristics as the one coming from the u-coordinates. Therefore, we can proceed in a similar way and deduce, at a cost of $\widetilde{O}(\ell^4)$ operations another univariate polynomial $\tilde{S}(x_3)$ belonging to the ideal. Now, since all the original equations have been taken into account, all common roots of S and $\tilde{S}$ will correspond to a solution of the original system for which we know that there are $O(\ell^2)$ solutions. Therefore taking the squarefree part of the GCD of S and $\tilde{S}$ yields a polynomial of degree $O(\ell^2)$.

This univariate polynomial can be factored at a cost of $\widetilde{O}(\ell^4)$ operations in $\mathbb{F}_q$ with standard algorithms [30] (there exist asymptotically faster algorithms, but we already fit within our target complexity). We then deal with each irreducible factor in turn, until one is found that leads to a genuine solution of the original system. Let δ be the degree of such an irreducible factor $\phi(x_3)$. In the field extension $\mathbb{F}_{q^\delta} = \mathbb{F}_q[x_3]/\phi(x_3)$, we have by construction a root x_3 of ϕ. We then solve again the original polynomial system where x_3 is instantiated with this root. This system is bivariate in x_1 and x_2 and there are $O(1)$ solutions, that possibly live in another finite extension $\mathbb{F}_{q^{\delta'}}$ of $\mathbb{F}_{q^\delta}$. Since the degrees of the bivariate polynomials are in $O(\ell^{2/3})$, by Lemma 3, solving this system costs $\widetilde{O}(\ell^2)$ operations in $\mathbb{F}_{q^\delta}$.

A solution obtained in this way must be checked, because it could come from a vanishing denominator that has been cleared when constructing the system or from nongeneric situations. But given a set of candidate coordinates for a D_i element of $J[\alpha_i]$, it is cheap to check that this is indeed an element of the Jacobian and that it is killed by α_i. Also, if α_i is not a generator of $\mathfrak{p}_i$, it is necessary to check the order of D_i: if this is a multiple of ℓ, then multiplying D_i by the cofactor gives an order-ℓ element. But it is also possible to get an unlucky element of small order coprime to ℓ, and then we have to take another solution of the system.

Since an operation in $\mathbb{F}_{q^\delta}$ requires a number of operations in $\mathbb{F}_q$ that is quasilinear in δ, and since the sum of all the degrees δ of the irreducible factors of $\mathrm{GCD}(S, \tilde{S})$ is in $O(\ell^2)$, the amortized cost is $\widetilde{O}(\ell^4)$ operations in $\mathbb{F}_q$ to deduce a divisor D_i in $J[\alpha_i]$.

3.3. ***Complexity of bi- and tri-variate resultants.*** In this section, the algorithms work by evaluation and interpolation, and require that there be enough elements in the base field. If there were not, we would simply take a field extension $\mathbb{F}_{q^\delta}$ of $\mathbb{F}_q$, which would add a factor $\widetilde{O}(\delta)$ to the complexity. The complexity of the algorithms will be polynomial in the number of evaluation points; therefore, the final complexity will be logarithmic in δ, so the cost of taking a field extension would be negligible in the $\widetilde{O}(\,)$ notation. We will therefore not mention this potential complication further.

Another difficulty is that an evaluation / interpolation strategy assumes that the points of evaluation are generic enough, so that all the degrees after evaluation are generic. This is again guaranteed by taking a large enough base field. Still, the algorithm remains a Monte–Carlo one. However, the ultimate goal is to construct kernel elements, which is an easily verified property. Turning this into a Las Vegas algorithm can therefore be done with standard techniques.

Lemma 3 [30, Theorem 6.22 and Corollary 11.21]. *Let $P(x, y)$ and $Q(x, y)$ be two polynomials whose degrees in x and y are bounded by d_x and d_y respectively. Then, $R(y) = \mathrm{Res}_x(P, Q)$ can be computed in $\widetilde{O}(d_x^2 d_y)$ field operations, and the degree of R is bounded by $2d_x d_y$.*

Lemma 4. *Let $P(x, y, z)$ and $Q(x, y, z)$ be two polynomials whose degrees in each variable are bounded by d. Then, $R(y, z) = \mathrm{Res}_x(P, Q)$ can be computed in $\widetilde{O}(d^5)$ field operations, and the degree of R in each variable is bounded by $2d^2$.*

Proof. The Sylvester matrix has at most $2d$ columns and its entries are bivariate polynomials whose degrees in y and z are bounded by d. Thus, its determinant is a polynomial whose degrees in y and z are bounded by $2d^2$.

We first perform a Kronecker substitution by considering $\tilde{P}(x, y) = P(x, y, y^{2d^2+1})$ and $\tilde{Q}(x, y) = Q(x, y, y^{2d^2+1})$, which are polynomials of degrees $\leq d$ in x and $\leq 2d^3 + d$ in y. Note that the choice to replace z by y^{2d^2+1} is made to be able to invert the Kronecker substitution after the resultant computation.

Next, we compute $\tilde{R}(y) = \mathrm{Res}_x(\tilde{P}(x, y), \tilde{Q}(x, y))$. By Lemma 3, it is a univariate polynomial of degree at most $4d^4 + 2d^2$ and can be computed in $\widetilde{O}(d^5)$ operations. We can then invert the Kronecker substitution to get $R(y, z)$, which can be done in time linear in the number of monomials, that is, in $O(d^4)$. □

3.4. ***Nongeneric situations.*** Our analysis assumes in the first place that the ℓ-torsion elements are generic in a rather strong sense, see, e.g., [1, Definition 11] for details. This is expected to be the case with overwhelming probability, when the base field is large enough and the curve is taken at random in a large family. However, to obtain a proven complexity we must also consider the cases where there exist ℓ-torsion elements that are nongeneric. We follow the strategy of [1] where another polynomial system is designed and solved for each nongeneric situation; for instance, the fact that an ℓ-torsion divisor is of weight less than 3, or that some points involved in the modeling are not distinct while they generically are. We do not give all the details, but the number of polynomial systems to consider is bounded by a constant, and each of these polynomial systems describes a situation that is smaller than the generic one in the sense that it has either fewer variables or a lower degree, so the complexity bound is maintained.

4. Bounds on the coefficients of ψ

The system of equations (1) giving σ_1, σ_2 and σ_3 in terms of a, b, c is homogeneous if we put weight $1/2$ to a, b, c and σ_1, weight 1 to q and σ_2, weight $3/2$ to σ_3, and weight 0 to μ_0, μ_1, and μ_2 so any polynomial in a reduced Gröbner basis of the corresponding ideal will have the same property. Computing such a Gröbner basis with the lexicographical ordering $a > b > c > \sigma_1 > \sigma_2 > \sigma_3 > \mu_0 > \mu_1 > \mu_2 > q$ (we did

this computation with the Magma v2.23-4 software), we get a polynomial Ψ_c of degree 6 in c that does not involve a or b, and which has the form

$$\Psi_c(q, c, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2) = D(\mu_0, \mu_1, \mu_2)^3\, c^6 + \sum_{i=0}^{5} \psi_c^{(i)}(q, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2)\, c^i,$$

where $D(\mu_0, \mu_1, \mu_2) = -27\,\mu_0^2 + 18\,\mu_0\mu_1\mu_2 - 4\,\mu_0\mu_2^3 - 4\,\mu_1^3 + \mu_1^2\mu_2^2$ is the discriminant of the polynomial $T^3 + \mu_2\,T^2 + \mu_1\,T + \mu_0$.

By computing Gröbner bases for other lexicographical orderings (with $a > c > b > \sigma_1 > \sigma_2 > \sigma_3 > \mu_0 > \mu_1 > \mu_2 > q$ and $b > c > a > \sigma_1 > \sigma_2 > \sigma_3 > \mu_0 > \mu_1 > \mu_2 > q$, respectively), we get that polynomials of the following form also belong to the ideal generated by the polynomials in the system (1):

$$\Psi_b(q, b, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2) = D(\mu_0, \mu_1, \mu_2)^3\, b^6 + \sum_{i=0}^{5} \psi_b^{(i)}(q, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2)\, b^i,$$

$$\Psi_a(q, a, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2) = D(\mu_0, \mu_1, \mu_2)^3\, a^6 + \sum_{i=0}^{5} \psi_a^{(i)}(q, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2)\, a^i.$$

The polynomials $\psi_a^{(i)}$, $\psi_b^{(i)}$ and $\psi_c^{(i)}$ are homogeneous of weighted degree $3 - i/2$ with respect to the grading given above.

Lemma 5. *The absolute values of the coefficients a, b, c of $\psi = a + b\eta + c\eta^2$ are bounded above by $O(q^{1/2})$.*

Proof. First, we consider the equation $\Psi_c = 0$. We write $c = \tilde{c}\,q^{1/2}$, $\sigma_1 = \tilde{\sigma}_1\,q^{1/2}$, $\sigma_2 = \tilde{\sigma}_2\,q$, $\sigma_3 = \tilde{\sigma}_3\,q^{3/2}$. Since $\psi_c^{(i)}$ is homogeneous and has weighted degree $3 - i/2$, a polynomial $\theta_c^{(i)}(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{\sigma}_3, \mu_0, \mu_1, \mu_2)$ exists such that

$$\psi_c^{(i)}(q, \sigma_1, \sigma_2, \sigma_3, \mu_0, \mu_1, \mu_2) \cdot c^i = q^3\,\tilde{c}^{\,i}\,\theta_c^{(i)}(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{\sigma}_3, \mu_0, \mu_1, \mu_2). \tag{4}$$

Weil's bounds imply that $|\tilde{\sigma}_i| = O(1)$ for $i \in \{1, 2, 3\}$. Therefore, for all $i \in \{0, \ldots, 5\}$, we obtain that $|\theta_c^{(i)}(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{\sigma}_3, \mu_0, \mu_1, \mu_2)| = O(1)$. For fixed $\mu_0, \mu_1, \mu_2 \in \mathbb{Q}$ such that $\mu_0 + \mu_1 T + \mu_2 T^2 + T^3$ is the minimal polynomial of a totally real algebraic number, the discriminant $D(\mu_0, \mu_1, \mu_2)$ must be nonzero. Equations $\Psi_c = 0$ and (4) imply the following inequality:

$$|\tilde{c}|^6 - \sum_{i=0}^{5} \frac{|\theta_c^{(i)}(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{\sigma}_3, \mu_0, \mu_1, \mu_2)|}{|D(\mu_0, \mu_1, \mu_2)|^3} |\tilde{c}|^i \leq 0.$$

Then $|\tilde{c}|$ must be smaller than or equal to the largest root of this polynomial inequality, which can itself be bounded, for instance, with Cauchy's bound

$$|\tilde{c}| \leq 1 + \max_{0 \leq i \leq 5} \left\{ \frac{|\theta_c^{(i)}(\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{\sigma}_3, \mu_0, \mu_1, \mu_2)|}{|D(\mu_0, \mu_1, \mu_2)|^3} \right\},$$

which shows that $|\tilde{c}| = O(1)$, and hence $|c| = O(q^{1/2})$. The proofs for the bounds on $|a|$ and $|b|$ are similar, using the equations $\Psi_a = 0$ and $\Psi_b = 0$. □

5. Small elements in ideals of $\mathbb{Z}[\eta]$

We first recall that we consider only primes ℓ that do not divide the discriminant of the minimal polynomial of η (condition (C2)). Hence, if $\mathbb{Z}[\eta]$ is not the maximal order of $\mathbb{Q}(\eta)$, this has no consequence on the factorization properties of ℓ.

Lemma 6. *For any prime ℓ that splits completely in $\mathbb{Z}[\eta]$, each prime ideal $\mathfrak{p}_i$ above ℓ contains a nonzero element α_i of the form $\alpha_i = a_i + b_i\eta + c_i\eta^2$, where $|a_i|$, $|b_i|$ and $|c_i|$ are integers in $O(\ell^{1/3})$, and the norm of α_i is in $O(\ell)$.*

Proof. The coefficients of the elements of the ideal $\mathfrak{p}_i$ represented by polynomials in η form a lattice. Applying Minkowski's bound to this lattice, we obtain the existence of a nonzero element $\alpha_i = a_i + b_i\eta + c_i\eta^2$ in $\mathfrak{p}_i$ for which the L_2-norm of (a_i, b_i, c_i) is in $O(\ell^{1/3})$. From this bound on the L_2-norm, we derive a bound on the L_∞-norm, and finally on the norm of α_i as an algebraic number. At each step, the constant hidden in the $O(\,)$ gets worse but still depends only on $\mathbb{Z}[\eta]$. □

For any given η, it is not difficult to make the constants in the $O(\,)$ fully explicit. We do it in the particular case of $\mathbb{Z}[\eta_7]$, with $\eta_7 = 2\cos(2\pi/7)$, which is the RM used in our practical experiments. Since $\mathbb{Z}[\eta_7]$ is a principal ring, a more direct approach leads to bounds for a generator that are tighter than what would be obtained by a naive application of Lemma 6.

Lemma 7. *Every ideal $\mathfrak{p}_i$ of norm ℓ in $\mathbb{Z}[\eta_7]$ has a generator α_i of the form $a_i + b_i\eta_7 + c_i\eta_7^2$, where $a_i, b_i, c_i \in \mathbb{Z}$ satisfy*

$$|a_i| < 2.415 \cdot \ell^{1/3}, \quad |b_i| < 1.850 \cdot \ell^{1/3}, \quad |c_i| < 1.764 \cdot \ell^{1/3}.$$

Proof. By abuse of notation, we identify $\mathbb{Q}(\eta_7)$ with the algebraic number field $\mathbb{Q}[X]/(X^3 + X^2 - 2X - 1)$ and we let σ_1, σ_2, σ_3 be the three real embeddings of $\mathbb{Q}(\eta_7)$ in $\mathbb{R}$. Let $\epsilon_1 = 1 - \eta_7^2$ and $\epsilon_2 = 1 + \eta_7$ be a pair of fundamental units, and let μ_i be a generator of $\mathfrak{p}_i$. The logarithmic embedding

$$\varphi : x \mapsto (\log|\sigma_1(x)|, \log|\sigma_2(x)|, \log|\sigma_3(x)|)$$

sends the set of generators of $\mathfrak{p}_i$ to the lattice generated by $\varphi(\epsilon_1)$ and $\varphi(\epsilon_2)$ translated by $\varphi(\mu_i)$. Solving a closest vector problem for the projection of $\varphi(\mu_i)$ on the plane where the three coordinates sum-up to 0, we deduce a unit ξ_i such that $\alpha_i = \xi_i\mu_i$ is a generator whose real embeddings are bounded by

$$|\sigma_1(\alpha_i)| \le 2.247 \cdot \ell^{1/3}, \quad |\sigma_2(\alpha_i)| \le 1.803 \cdot \ell^{1/3}, \quad |\sigma_3(\alpha_i)| \le 2.247 \cdot \ell^{1/3}.$$

Writing $\alpha_i = a_i + b_i\eta_7 + c_i\eta_7^2$, the real embeddings can also be expressed as $(\sigma_1(\alpha_i), \sigma_2(\alpha_i), \sigma_3(\alpha_i))^T = V \cdot (a_i, b_i, c_i)^T$, where V is the Vandermonde matrix of $(\sigma_1(\eta_7), \sigma_2(\eta_7), \sigma_3(\eta_7))$. A numerical evaluation of its inverse allows the translation of the bounds on $\sigma_1(\alpha_i)$, $\sigma_2(\alpha_i)$, $\sigma_3(\alpha_i)$ into the claimed bounds on a_i, b_i, c_i. □

6. Bounding the degrees of Cantor's division polynomials in genus 3

The purpose of this section is to prove the following lemma on the Cantor's division polynomials, which are explicit formulas for the endomorphism corresponding to scalar multiplication [6].

Lemma 8. *In genus* 3, *the degrees of Cantor's ℓ-division polynomials are bounded by* $O(\ell^2)$.

In [6], there are exact formulas for the degrees of the leading and the constant coefficients d_3 and d_0. However, there is no formula or bound for the degrees of the other coefficients of the ℓ-division polynomials. Still, our proof strongly relies on [6] and we do not try to make it standalone: we assume that the reader is familiar with that article, and all references to expressions, propositions or definitions in this proof are taken from that paper.

For a polynomial P whose coefficients are themselves univariate polynomials, we denote by $\operatorname{maxdeg}(P)$ the maximum of the degrees of its coefficients.

We first prove a bound on the degrees of the coefficients of the quantities α_r and γ_r defined in [6], from which the wanted bounds will follow. The key tools are the recurrence formulas (8.31) and (8.33) that relate quantities at index r to quantities at index around $r/2$, in a similar fashion as for the division polynomials of elliptic curves. More precisely, the following lemma shows that when the index r is (roughly) doubled, $\operatorname{maxdeg}\alpha_r$ and $\operatorname{maxdeg}\gamma_r$ are roughly multiplied by 4, which leads to the expected quadratic growth.

Lemma 9. *Let* $\ell \geq 12$, *and assume that for all* $i \leq (\ell+9)/2$ *the degrees* $\operatorname{maxdeg}\alpha_i$ *and* $\operatorname{maxdeg}\gamma_i$ *are bounded by* C, *then* $\operatorname{maxdeg}\alpha_\ell$ *and* $\operatorname{maxdeg}\gamma_\ell$ *are bounded by* $4C + 36\ell + 108$.

Proof. We first deal with the bound on $\operatorname{maxdeg}\gamma_\ell$. Let us consider r and s around $\ell/2$ such that $\ell = r + s - 5$: we take either $r = s - 3 = \ell/2 + 1$ if ℓ is even, or $r = s - 4 = (\ell+1)/2$ otherwise.

From (8.30) and (8.31), the degree of $\gamma_\ell[h]\psi_{s-r}\psi_{r-2}\psi_{s-2}\psi_{r-1}\psi_{s-1}$ is that of the determinant of the matrix $\mathcal{E}_{rs}[h]$, defined

$$\mathcal{E}_{rs}[h] = \begin{pmatrix} \alpha_{r-3}\alpha_s[0] & \alpha_{r-3}\alpha_s[1] & \psi_{r-3}\psi_s & \gamma_{r-3}\gamma_s[h] \\ \alpha_{r-2}\alpha_{s-1}[0] & \alpha_{r-2}\alpha_{s-1}[1] & \psi_{r-2}\psi_{s-1} & \gamma_{r-2}\gamma_{s-1}[h] \\ \alpha_{r-1}\alpha_{s-2}[0] & \alpha_{r-1}\alpha_{s-2}[1] & \psi_{r-1}\psi_{s-2} & \gamma_{r-1}\gamma_{s-2}[h] \\ \alpha_r\alpha_{s-3}[0] & \alpha_r\alpha_{s-3}[1] & \psi_r\psi_{s-3} & \gamma_r\gamma_{s-3}[h] \end{pmatrix}.$$

Therefore we have an expression for the degrees of the coefficients of γ_ℓ in terms of objects at index around r and s:

$$\deg\gamma_\ell[h] \leq \deg\det\mathcal{E}_{rs}[h] - \deg(\psi_{r-2}\psi_{s-2}\psi_{r-1}\psi_{s-1}).$$

In this last formula, the factor ψ_{s-r} has been omitted, because $s-r$ is either 3 or 4, and by (8.17) this has nonnegative degree in any case. Thus, we simply bounded it below by 0 in the previous inequality. Before entering a more detailed analysis, we use (8.8) to rewrite the first column with expressions for which we have exact formulas for the degree:

$$\mathcal{E}_{rs}[h] = \begin{pmatrix} \psi_{r-4}\psi_{s-1} & \alpha_{r-3}\alpha_s[1] & \psi_{r-3}\psi_s & \gamma_{r-3}\gamma_s[h] \\ \psi_{r-3}\psi_{s-2} & \alpha_{r-2}\alpha_{s-1}[1] & \psi_{r-2}\psi_{s-1} & \gamma_{r-2}\gamma_{s-1}[h] \\ \psi_{r-2}\psi_{s-3} & \alpha_{r-1}\alpha_{s-2}[1] & \psi_{r-1}\psi_{s-2} & \gamma_{r-1}\gamma_{s-2}[h] \\ \psi_{r-1}\psi_{s-4} & \alpha_r\alpha_{s-3}[1] & \psi_r\psi_{s-3} & \gamma_r\gamma_{s-3}[h] \end{pmatrix}.$$

The determinant of $\mathcal{E}_{rs}[h]$ is the sum of products of four ψ factors and four α or γ factors. The degrees of the former are explicitly known, while by hypothesis we have upper bounds on the latter,

since all the indices are at most $(\ell+9)/2$. We can then deduce an upper bound on the degree of this determinant. All the ψ_i have indices with i in the range $[r-4, s]$ (remember that $r \leq s$), and since their degrees increase with the indices, we can upper bound the degree of the products of the four ψ factors by $4 \deg \psi_s$. Therefore we have

$$\deg \det \mathcal{E}_{rs}[h] \leq 4(\deg \psi_s + C).$$

In order to deduce an upper bound on $\operatorname{maxdeg} \gamma_\ell$, it remains to get a lower bound on the degree of the $\deg(\psi_{r-2}\psi_{s-2}\psi_{r-1}\psi_{s-1})$ term, and again by monotonicity of the degree in the index, we lower bound it by $4 \deg \psi_{r-2}$. So finally, we get

$$\operatorname{maxdeg} \gamma_\ell \leq 4C + (\deg \psi_s^4 - \deg \psi_{r-2}^4).$$

Using (8.16) and (8.17), we deduce that for all k, we have $\deg(\psi_k^2) = 3(k^2-9)$ and substituting this value and the expression of $r-2$ and s in terms of ℓ, we obtain

$$\deg \psi_s^4 - \deg \psi_{r-2}^4 = \begin{cases} 30\ell + 90 & \text{if } \ell \text{ is even,} \\ 36\ell + 108 & \text{if } \ell \text{ is odd,} \end{cases}$$

and the result follows for $\operatorname{maxdeg} \gamma_\ell$.

The proof for $\operatorname{maxdeg} \alpha_\ell$ follows the same line. Using the matrix $\mathcal{F}_{rs}[h]$ defined in (8.32) in a similar way to which we used the matrix $\mathcal{E}_{rs}[h]$, and with the help of the formula (8.33), we end up with the bounds

$$\operatorname{maxdeg} \alpha_\ell \leq \begin{cases} 4C + 30\ell - 30 & \text{if } \ell \text{ is even,} \\ 4C + 36\ell - 36 & \text{if } \ell \text{ is odd,} \end{cases}$$

which are stricter than our target.

Finally, the bound $\ell \geq 12$ is necessary to ensure that the quantities r and s are at least 5, as required in [6] to apply the formulas (8.31) and (8.33). □

We can now finish the proof of Lemma 8. We define two sequences $(\ell_i)_{i\geq 0}$ and $(C_i)_{i\geq 0}$ as follows: let $\ell_0 = 12$ and let C_0 be a bound on the degrees of the coefficients of all the α_i and γ_i for $i \leq \ell_0$. Then for all $i \geq 1$, we define the sequences inductively by

$$\begin{cases} \ell_{i+1} = 2\ell_i - 9, \\ C_{i+1} = 4C_i + 36\ell_{i+1} + 108. \end{cases}$$

By Lemma 9, for all i and all $\ell \leq \ell_i$, the degrees $\operatorname{maxdeg} \alpha_\ell$ and $\operatorname{maxdeg} \gamma_\ell$ are bounded by C_i. The expression $\ell_i = (\ell_0 - 9)2^i + 9 = 3 \cdot 2^i + 9$ can be derived directly from the definition and substituted in the recurrence formula of C_{i+1} to get $C_{i+1} = 4C_i + 216 \cdot 2^i + 432$. This recurrence can be solved by setting $\Gamma_i = C_i + 108 \cdot 2^i + 144$, so that $\Gamma_{i+1} = 4\,\Gamma_i$, and we obtain $C_i = (C_0 + 252)\,4^i - 108 \cdot 2^i - 144$. Finally, for any ℓ, we select the smallest i such that $\ell \leq \ell_i$. This value of i is $\lceil \log_2((\ell-9)/3) \rceil$. The corresponding bound for $\operatorname{maxdeg} \alpha_\ell$ and $\operatorname{maxdeg} \gamma_\ell$ is then C_i, which grows like $O(\ell^2)$ (and we remark that the effect of the ceiling can make the constant hidden in the $O(\,)$ expression grow by a factor of at most 3).

Using the expression (8.10), we have $\operatorname{maxdeg} \delta_\ell \leq \operatorname{maxdeg} \alpha_\ell + \operatorname{maxdeg} \gamma_\ell$, and therefore the bound $O(\ell^2)$ also applies to the degrees of the coefficients of δ_ℓ. Using the formula (8.13), the same holds for the coefficients of ϵ_ℓ / y.

This concludes the proof of Lemma 8.

7. Experimental results

In order to evaluate the practicality of our algorithm, we have tested it on one of the families of genus-3 hyperelliptic curves having explicit RM given in [28, Theorem 1]. Formulas for their RM endomorphisms are described in [20]: for $t \neq \pm 2$, the curve $\mathcal{C}_t$ with equation

$$y^2 = x^7 - 7x^5 + 14x^3 - 7x + t,$$

admits an endomorphism given in Mumford representation by

$$\eta_7(x, y) = \langle X^2 + 11\, x X/2 + x^2 - 16/9,\ y \rangle.$$

The fact that this expression has degree 2 while one would generically expect a degree 3 is no accident: it comes from the construction in [28] of the endomorphism as a sum of two automorphisms on a double cover of the curve. We have $\eta_7^3 + \eta_7^2 - 2\eta_7 - 1 = 0$, so the ring $\mathbb{Z}[\eta_7]$ is isomorphic to the ring of integers $\mathbb{Z}[2\cos(2\pi/7)]$ of the real subfield of the cyclotomic field $\mathbb{Q}(e^{2i\pi/7})$. All the numerical data in this section have been obtained for the parameter $t = 42$, on the prime field $\mathbb{F}_p$ with $p = 2^{64} - 59$.

In our practical computations, the main differences to the theoretical description are the following: we use Gröbner basis algorithms instead of resultants, we consider also small nonsplit primes ℓ and small powers, and we finish the computation with a parallel collision search. The source code for our experiments is available here: https://members.loria.fr/SAbelard/RMg3.tgz.

7.1. ***Computing modular information with Gröbner basis.*** Although the polynomial system resolution using resultants has a complexity of $\widetilde{O}(\ell^4)$, the real cost for small values of ℓ is already pretty large. In the resolution method described in Section 3.2, each bivariate resultant is computed by evaluation / interpolation and hence requires the computation of many univariate resultants. We illustrate this by counting the number of univariate resultants to perform and their degrees for the main step of the resolution (the part that reaches the peak complexity). In the following table, we measure the cost of such resultant computations using the NTL 10.5.0 and FLINT 2.5.2 libraries, both linked against GMP 6, when the base field is $\mathbb{F}_{2^{64}-59}$. These costs do not include the evaluation / interpolation steps which might also be problematic for large instances, because they are hard to parallelize.

ℓ	# res	deg	cost (NTL)	cost (FLINT)
13	5.25×10^8	16,000	1,850 days	735 days
29	1.28×10^{10}	80,000	310,000 days	190,000 days

We were more successful with the direct approach using Gröbner bases that we now describe. For computing the kernel of a given endomorphism, we computed a Gröbner basis of (3) with some small

modifications. First, we observe that the only occurrences of y_1 and y_2 are within the monomial y_1y_2. Consequently, we can remove one variable by replacing each occurrence of y_1y_2 by a fresh variable y. Next, we need to make the system 0-dimensional by encoding the fact that $d_3(x_3)$ and $\tilde{d}_3(x_1, x_2)$ are nonzero. This is done by introducing another fresh variable t and by adding the polynomial $S(x_1, x_2, x_3)t - 1$ to the system, where $S(x_1, x_2, x_3)$ is the squarefree part of $d_3(x_3)\tilde{d}_3(x_1, x_2)$. Finally, since each polynomial is symmetric with respect to the transposition of the variables x_1 and x_2, we can rewrite the equations using the symmetric polynomials $s_1 = x_1 + x_2$ and $s_2 = x_1 x_2$. This divides by 2 the degree in x_1 and x_2 of the equations. We end up with a system in five variables.

The whole construction can be slightly modified to compute the preimage of a given divisor by the endomorphism: to model $\alpha(D) = Q$, we write $D = P_1 + P_2 + P_3 - 3\infty$ and solve for $\alpha(P_1 - \infty) + \alpha(P_2 - \infty) = Q - \alpha(P_3 - \infty)$. In that case, the variable y_3 gets involved in all the equations, so that we get a system in six variables.

For $\ell = 2$, the 2-torsion elements are easily deduced from the factorization of f, and by computing a preimage of a 2-torsion divisor, we get a point in $J[4]$ from which we could deduce $a, b, c \bmod 4$. Dividing again by 2 was too costly, due to the fact that the 4-torsion point was in an extension of degree 4. For $\ell = 3$, which is an inert prime, we ran the kernel computation for the multiplication-by-3 endomorphism, without using the RM property. The norm being 27, this is the largest modular computation that we performed (and the most costly in terms of time and memory). The prime $\ell = 7$ ramifies in $\mathbb{Z}[\eta_7]$ as the cube of the ideal generated by $\alpha_7 = -2 - \eta_7 + \eta_7^2$. The kernel of α_7 can be computed but it yields only one linear relation in $a, b, c \bmod 7$. Dividing the kernel elements by α_7 would give more information, but this computation did not finish due to the field extension in which the divisors are defined. The first split prime is $\ell = 13$. We use the following small generators: $(13) = (2 - \eta_7 - 2\eta_7^2)(-2 + 2\eta_7 + \eta_7^2)(3 + \eta_7 - \eta_7^2)$, which seem to produce the polynomial systems with the smallest degrees. For instance, the apparently smaller element $1 + \eta_7^2$ of norm 13 yields equations of much higher degrees; 7, 71, 72, 73, 72. The next split prime is 29, which would maybe have been feasible, but was not necessary for our setting. In the following table, we summarize the data for these systems, which were obtained with Magma v2.23-4 on a Xeon E7-4850v3 at 2.20GHz, with 1.5 TB RAM.[1]

mod ℓ^k	#var	degree of each eq.	time	memory	$a, b, c \bmod \ell^k$
2	—	—	—	—	0, 0, 0
4 (inert^2)	6	7, 7, 14, 15, 15, 10	1 min	negl.	2, 2, 2
3 (inert)	5	7, 53, 54, 55, 26	14 days	140 GB	1, 2, 1
$7 = \mathfrak{p}_1^3$	5	7, 35, 36, 37, 36	3.5h	6.6 GB	$a + 2b + 4c \equiv 2$
$13 = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$	5	7, 44, 45, 46, 52	3×3 days	41 GB	12, 10, 9
$29 = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$	5	7, 92, 93, 94, 100	$>3\times2$ weeks	>0.8 TB	—

[1]The F4 algorithm can be highly sensitive to the modeling of the problem and we refer to the source code. In particular, thanks to serendipity, we saved a factor greater than 12 in the runtime for $\ell = 7, 13$ by forgetting to take the squarefree part of the saturation polynomial. We have no explanation for this phenomenon.

7.2. ***Parallel collision search for RM curves.*** The classical square-root-complexity search in genus 3 requires $O(q)$ group operations [8]. For RM curves, this can be improved by searching for the coefficients a, b, c of $\psi = \pi + \pi^\vee$ in $\mathbb{Z}[\eta]$. This readily yields a complexity of $O(q^{3/4})$, using the equation $aD + b\eta(D) + c\eta^2(D) = (q+1)D$, which must be satisfied for any rational divisor D. While a baby-step, giant-step approach is straightforward to design, it needs $O(q^{3/4})$ space and this is the bottleneck. A low-memory, parallel version of this search can be obtained with the algorithm of [12]. There the details are given only for a 2-dimensional problem, while here we have a 3-dimensional problem, but we did not encounter any unexpected issues when adapting the parameters to our case. Also, just like in [12], including some anterior modular knowledge is straightforward: if a, b, c are known modulo m, the expected time is $O(q^{3/4}/m^{3/2})$.

We wrote a dedicated C implementation with a few lines of assembly to speed up the addition and multiplication in $\mathbb{F}_p$, taking advantage of the special form of p. This implementation performs 10.7 million operations in the Jacobian per second using 32 (hyperthreaded) threads of a 16-core bi-Xeon E5-2650 at 2 GHz. We used the knowledge of ψ modulo 156 but not of the known relation modulo 7 for simplicity (there is no obstruction to using it and saving an additional $7^{1/2}$ factor).

After computing about 190,000 chains of average length 32,000,000, we got a collision, from which we deduced

$$\psi = 2551309006 + 2431319810\ \eta_7 - 847267802\ \eta_7^2,$$

and the coefficients of the characteristic polynomial χ_π of the Frobenius are then

$$\sigma_1 = 986268198, \quad \sigma_2 = 35389772484832465583, \quad \sigma_3 = 109560528621042368187702122 44.$$

The number of group operations that were done is slightly less than $43\,(p^{3/4}/156^{3/2})$. This factor 43 is close to the average that we observed in our numerous experiments with smaller sizes. Scaled on a single (physical) core, we can estimate the cost of this collision search to be 105 core-days.

Acknowledgements.

We are grateful to Benjamin Smith for fruitful discussions and to Allan Steel for his help with memory issues with Magma. We also wish to thank anonymous referees for their comments which helped improve the paper.

References

[1] Simon Abelard, Pierrick Gaudry, and Pierre-Jean Spaenlehauer, *Improved complexity bounds for counting points on hyperelliptic curves*, Found. Comput. Math. (2018).

[2] Leonard M. Adleman and Ming-Deh Huang, *Counting points on curves and abelian varieties over finite fields*, J. Symbolic Comput. **32** (2001), no. 3, 171–189. MR 1851164

[3] Sean Ballentine, Aurore Guillevic, Elisa Lorenzo García, Chloe Martindale, Maike Massierer, Benjamin Smith, and Jaap Top, *Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication*, Algebraic geometry for coding theory and cryptography, Assoc. Women Math. Ser., no. 9, Springer, 2017, pp. 63–94. MR 3775425

[4] A. Bostan, G. Lecerf, B. Salvy, É. Schost, and B. Wiebelt, *Complexity issues in bivariate polynomial factorization*, ISSAC 2004, ACM, New York, 2004, pp. 42–49. MR 2126923

[5] Ivan Boyer, *Variétés abéliennes et jacobiennes de courbes hyperelliptiques, en particulier à multiplication réelle ou complexe*, Ph.D. thesis, Université Paris Diderot, 2014.

[6] David G. Cantor, *On the analogue of the division polynomials for hyperelliptic curves*, J. Reine Angew. Math. **447** (1994), 91–145. MR 1263171

[7] Robert Carls and David Lubicz, *A p-adic quasi-quadratic time point counting algorithm*, Int. Math. Res. Not. **2009** (2009), no. 4, 698–735. MR 2480098

[8] Noam D. Elkies, *Elliptic and modular curves over finite fields and related computational issues*, Computational perspectives on number theory (Chicago, 1995), AMS/IP Stud. Adv. Math., no. 7, Amer. Math. Soc., Providence, RI, 1998, pp. 21–76. MR 1486831

[9] Jordan S. Ellenberg, *Endomorphism algebras of Jacobians*, Adv. Math. **162** (2001), no. 2, 243–271. MR 1859248

[10] Gerhard Frey and Michael Müller, *Arithmetic of modular curves and applications*, Algorithmic algebra and number theory (Heidelberg, 1997), Springer, 1999, pp. 11–48. MR 1672093

[11] Pierrick Gaudry, David Kohel, and Benjamin Smith, *Counting points on genus 2 curves with real multiplication*, Advances in cryptology—ASIACRYPT 2011, Lecture Notes in Comput. Sci., no. 7073, Springer, 2011, pp. 504–519. MR 2935020

[12] Pierrick Gaudry and Éric Schost, *A low-memory parallel version of Matsuo, Chao, and Tsujii's algorithm*, Algorithmic number theory—ANTS VI, Lecture Notes in Comput. Sci., no. 3076, Springer, 2004, pp. 208–222. MR 2137355

[13] ———, *Genus 2 point counting over prime fields*, J. Symbolic Comput. **47** (2012), no. 4, 368–400. MR 2890878

[14] Michael C. Harrison, *An extension of Kedlaya's algorithm for hyperelliptic curves*, J. Symbolic Comput. **47** (2012), no. 1, 89–101. MR 2854849

[15] David Harvey, *Counting points on hyperelliptic curves in average polynomial time*, Ann. of Math. (2) **179** (2014), no. 2, 783–803. MR 3152945

[16] David Harvey and Andrew V. Sutherland, *Computing Hasse–Witt matrices of hyperelliptic curves in average polynomial time, II*, Frobenius distributions: Lang–Trotter and Sato–Tate conjectures, Contemp. Math., no. 663, Amer. Math. Soc., Providence, RI, 2016, pp. 127–147. MR 3502941

[17] Ming-Deh Huang and Doug Ierardi, *Counting points on curves over finite fields*, J. Symbolic Comput. **25** (1998), no. 1, 1–21. MR 1600606

[18] Kiran S. Kedlaya, *Counting points on hyperelliptic curves using Monsky–Washnitzer cohomology*, J. Ramanujan Math. Soc. **16** (2001), no. 4, 323–338, Correction in **18**:4 (2003), 417–418. MR 1877805 arXiv math/0105031

[19] Kiran S. Kedlaya and Andrew V. Sutherland, *Computing L-series of hyperelliptic curves*, Algorithmic number theory—ANTS VIII, Lecture Notes in Comput. Sci., no. 5011, Springer, 2008, pp. 312–326. MR 2467855

[20] David R. Kohel and Benjamin A. Smith, *Efficiently computable endomorphisms for hyperelliptic curves*, Algorithmic number theory—ANTS VII, Lecture Notes in Comput. Sci., no. 4076, Springer, 2006, pp. 495–509. MR 2282945

[21] Herbert Lange and Christina Birkenhake, *Complex abelian varieties*, Grundlehren der Math. Wissenschaften, no. 302, Springer, 1992. MR 1217487

[22] J.-F. Mestre, *Familles de courbes hyperelliptiques à multiplications réelles*, Arithmetic algebraic geometry (Texel, 1989), Progr. Math., no. 89, Birkhäuser, Boston, 1991, pp. 193–208. MR 1085260

[23] J. Pila, *Frobenius maps of abelian varieties and finding roots of unity in finite fields*, Math. Comp. **55** (1990), no. 192, 745–763. MR 1035941

[24] Takakazu Satoh, *The canonical lift of an ordinary elliptic curve over a finite field and its point counting*, J. Ramanujan Math. Soc. **15** (2000), no. 4, 247–270. MR 1801221

[25] René Schoof, *Elliptic curves over finite fields and the computation of square roots mod p*, Math. Comp. **44** (1985), no. 170, 483–494. MR 777280

[26] Goro Shimura, *Introduction to the arithmetic theory of automorphic functions*, Publ. Math. Soc. Japan, no. 11, Iwanami Shoten, Tokyo, 1971. MR 0314766

[27] Andrew V. Sutherland, *A generic approach to searching for Jacobians*, Math. Comp. **78** (2009), no. 265, 485–507. MR 2448717

[28] Walter Tautz, Jaap Top, and Alain Verberkmoes, *Explicit hyperelliptic curves with real multiplication and permutation polynomials*, Canad. J. Math. **43** (1991), no. 5, 1055–1064. MR 1138583

[29] Paul van Wamelen, *Proving that a genus* 2 *curve has complex multiplication*, Math. Comp. **68** (1999), no. 228, 1663–1677. MR 1648415

[30] Joachim von zur Gathen and Jürgen Gerhard, *Modern computer algebra*, 3rd ed., Cambridge Univ. Press, 2013. MR 3087522

[31] Richard Zippel, *Effective polynomial computation*, Springer Int. Series Engrg. Comput. Sci., no. 241, Springer, 1993.

Received 27 Feb 2018. Revised 18 Sep 2018.

SIMON ABELARD: simon.abelard@loria.fr
Université de Lorraine, CNRS, Inria, Nancy, France

PIERRICK GAUDRY: pierrick.gaudry@loria.fr
Université de Lorraine, CNRS, Inria, Nancy, France

PIERRE-JEAN SPAENLEHAUER: pierre-jean.spaenlehauer@loria.fr
Université de Lorraine, CNRS, Inria, Nancy, France

VOLUME EDITORS

Renate Scheidler
University of Calgary
Calgary, AB T2N 1N4
Canada

Jonathan Sorenson
Butler University
Indianapolis, IN 46208
United States

ISSN: 2329-9061 (print), 2329-907X (electronic)

ISBN: 978-1-935107-02-6 (print), 978-1-935107-03-3 (electronic)

First published 2019.

MATHEMATICAL SCIENCES PUBLISHERS
798 Evans Hall #3840, c/o University of California, Berkeley CA 94720-3840
contact@msp.org **http://msp.org**

THE OPEN BOOK SERIES 2

Thirteenth Algorithmic Number Theory Symposium

The Algorithmic Number Theory Symposium (ANTS), held biennially since 1994, is the premier international forum for research in computational number theory. ANTS is devoted to algorithmic aspects of number theory, including elementary, algebraic, and analytic number theory, the geometry of numbers, arithmetic algebraic geometry, the theory of finite fields, and cryptography.

This volume is the proceedings of the thirteenth ANTS meeting, held July 16-20, 2018, at the University of Wisconsin-Madison. It includes revised and edited versions of 28 refereed papers presented at the conference.

Edited by Renate Scheidler and Jonathan Sorenson

CONTRIBUTORS

Simon Abelard
Sonny Arora
Vishal Arul
Angelica Babei
Jens-Dietrich Bauch
Alex J. Best
Jean-François Biasse
Alin Bostan
Reinier Bröker
Nils Bruin
Xavier Caruso
Stephanie Chan
Qi Cheng
Gilles Christol
Owen Colman
Edgar Costa
Philippe Dumas
Kirsten Eisenträger
Claus Fieker
Shuhong Gao
Pierrick Gaudry
Alexandre Gélin
Alexandru Ghitza
Laurent Grémy
Jeroen Hanselman
David Harvey
Tommy Hofmann
Everett W. Howe
David Hubbard
Kiran S. Kedlaya
Thorsten Kleinjung
David Kohel
Wanlin Li
Richard Magner
Anna Medvedovsky
Michael Musty
Ha Thanh Nguyen Tran
Christophe Ritzenthaler
David Roe
J. Maurice Rojas
Nathan C. Ryan
Renate Scheidler
Sam Schiavone
Andrew Shallue
Jeroen Sijsling
Carlo Sircana
Jonathan Sorenson
Pierre-Jean Spaenlehauer
Andrew V. Sutherland
Nicholas Triantafillou
Joris van der Hoeven
Christine Van Vredendaal
John Voight
Daqing Wan
Lawrence C. Washington
Jonathan Webster
Benjamin Wesolowski
Yinan Zhang
Alexandre Zotine