

Counting points on hyperelliptic curves with explicit real multiplication in arbitrary genus

Simon Abelard

► **To cite this version:**

Simon Abelard. Counting points on hyperelliptic curves with explicit real multiplication in arbitrary genus. 2018. <hal-01905580>

HAL Id: hal-01905580

<https://hal.inria.fr/hal-01905580>

Submitted on 25 Oct 2018

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.

Counting points on hyperelliptic curves with explicit real multiplication in arbitrary genus

Simon Abelard

Université de Lorraine, CNRS, Inria
Cheriton School of Computer Science, University of Waterloo

October 25, 2018

Abstract

We present a probabilistic Las Vegas algorithm for computing the local zeta function of a genus- g hyperelliptic curve defined over \mathbb{F}_q with explicit real multiplication (RM) by an order $\mathbb{Z}[\eta]$ in a degree- g totally real number field.

It is based on the approaches by Schoof and Pila in a more favorable case where we can split the ℓ -torsion into g kernels of endomorphisms, as introduced by Gaudry, Kohel, and Smith in genus 2. To deal with these kernels in any genus, we adapt a technique that the author, Gaudry, and Spaenlehauer introduced to model the ℓ -torsion by structured polynomial systems. Applying this technique to the kernels, the systems we obtain are much smaller and so is the complexity of solving them.

Our main result is that there exists a constant $c > 0$ such that, for any fixed g , this algorithm has expected time and space complexity $O((\log q)^c)$ as q grows and the characteristic is large enough. We prove that $c \leq 8$ and we also conjecture that the result still holds for $c = 6$.

1 Introduction

Due to its numerous applications in cryptology, number theory, algebraic geometry or even as a primitive used in other algorithms, the problem of counting points on curves and Abelian varieties has been extensively studied over the past three decades. Among the milestones in the history of point-counting, one can mention the first polynomial-time algorithm by Schoof [26] for counting points on elliptic curves, and the subsequent extension to Abelian varieties by Pila [22]. Using similar approaches, we design a probabilistic algorithm for computing the local zeta functions of hyperelliptic curves with explicit real multiplication of arbitrary fixed genus and bound its complexity.

Given an Abelian variety of dimension g over a finite field \mathbb{F}_q , Pila's algorithm computes its local zeta function in time $(\log q)^\Delta$, where Δ is doubly exponential in g . Further contributions were made in [17, 4] so that this exponent Δ is now proven to be polynomial in g in general, and even linear in the hyperelliptic case [2].

In genus 2, a tailor-made extension of Schoof's algorithm due to Gaudry, Harley and Schost [11, 13, 14] allows to count points in time $\tilde{O}(\log^8 q)$. Yet, this remains much larger than the complexity of the Schoof-Elkies-Atkin (SEA) algorithm [27], which is the standard for elliptic point-counting in large characteristic and runs in $\tilde{O}(\log^4 q)$

bit operations. For genus-2 curves with explicit real multiplication (RM), i.e. curves having an additional endomorphism for which an explicit expression is known, a much more efficient point-counting algorithm is introduced in [12] with a bit complexity in $\tilde{O}(\log^5 q)$, thus narrowing the gap between genus 1 and 2.

These algorithms were extended to genus-3 hyperelliptic curves in [3] with an asymptotic complexity in $\tilde{O}(\log^{14} q)$ bit operations that is decreased to $\tilde{O}(\log^6 q)$ bit operations when the curve has explicit RM.

The aim of this paper is to study the asymptotic complexity of point-counting on hyperelliptic curves with explicit RM when g is arbitrary large. In this case, we bound the exponent of $\log q$ by 8 and therefore remove the dependency on g from the exponent of $\log q$.

Another way to avoid such a painful dependency in g in the complexity without restricting to such particular cases is to use the p -adic methods, in the spirit of Satoh's and Kedlaya's algorithms [24, 18] for elliptic and hyperelliptic curves. These methods have also been extended beyond the hyperelliptic case [29, 8] and one can also mention the algorithms of Lauder and Lauder-Wan that also hold for very general varieties [20, 21]. Although these methods are polynomial in g , they are exponential in $\log p$ and therefore cannot be used in large characteristic.

Indeed, the ℓ -adic approaches derived from Schoof's algorithm and the p -adic approaches are complementary when either g or p is small but we still lack a classical algorithm running in time polynomial in both g and $\log q$. However, for counting points on reductions modulo many primes p of the same curve, an algorithm introduced by Harvey in [16] is polynomial in g and polynomial *on average* in $\log p$.

In this paper, we follow the spirit of the Schoof-Pila algorithm and recover the local zeta function by computing the characteristic polynomial χ_π of the action of the Frobenius endomorphism π on the ℓ -torsion subgroups for sufficiently many primes ℓ . The key to our complexity result is that, thanks to the real multiplication, it is sufficient to have π act on much smaller subgroups of the ℓ -torsion, at least for a positive proportion of the primes ℓ .

More precisely, we say that a curve \mathcal{C} has *explicit real multiplication* by $\mathbb{Z}[\eta]$ if the subring $\mathbb{Z}[\eta] \subset \text{End}(\text{Jac}(\mathcal{C}))$ is isomorphic to an order in a totally real degree- g number field, and if we have explicit formulas describing $\eta(P - \infty)$ for some fixed base point ∞ and a generic point P of \mathcal{C} . By explicit formulas, we mean polynomials $(\eta_i^{(u)}(x, y))_{i \in \{0, 1, \dots, g\}}$ and $(\eta_i^{(v)}(x, y))_{i \in \{0, 1, \dots, g\}}$ 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

$$\left\langle \sum_{i=0}^g \eta_i^{(u)}(x, y) X^i, \sum_{i=0}^{g-1} (\eta_i^{(v)}(x, y) / \eta_g^{(v)}(x, y)) X^i \right\rangle,$$

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.

As in [12, 3], we consider primes $\ell \in \mathbb{Z}$ such that $\ell\mathbb{Z}[\eta]$ splits as a product $\mathfrak{p}_1 \cdots \mathfrak{p}_g$ 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 $\text{Jac}(\mathcal{C})[\ell] \subset \text{Ker } \alpha_1 + \cdots + \text{Ker } \alpha_g$. Then, we compute from this representation integers a_0, \dots, a_{g-1} in $\mathbb{Z}/\ell\mathbb{Z}$ such that the sum $\pi + \pi^\vee$ of the Frobenius endomorphism and its dual equals $a_0 + a_1\eta + \cdots + a_{g-1}\eta^{g-1} \pmod{\ell}$. Once enough modular information is known, the values of the a_i 's

such that $\pi + \pi^\vee = \sum_{i=0}^{g-1} a_i \eta_i$ are recovered via the Chinese Remainder Theorem and the coefficients of the characteristic polynomial of the Frobenius can be directly expressed in terms of the a_i 's.

Computing the kernels of the endomorphisms α_i is the dominant step in terms of complexity and thus the cornerstone of our result. We still model these kernels by polynomial systems that we then have to solve, but the resultant-based techniques that were used in [12] and [3] are no longer satisfying when g is arbitrary large. We therefore use the modelling strategy of [2] and apply it to the kernels instead of applying it to the whole ℓ -torsion. The polynomial systems we derive from this approach are in fact very similar to those of [2], except that our kernels are comparable in size to the “ $\ell^{1/g}$ -torsion”, resulting in much smaller degrees and ultimately in a complexity gain by a factor g in the exponent of $\log q$, decreasing it from linear to constant. Using the geometric resolution algorithm just as in [2], we solve these systems in time $K(\log q)^{8+o(1)}$ where K depends on η (and thus on g too) but not on q . It is interesting to note that this result suffers from the pessimistic cubic bounds on the degrees of Cantor’s polynomials established in [2] and that—assuming quadratic bounds as proven in genus 1, 2 and 3—we get a complexity in $K(\log q)^{6+o(1)}$, which is similar to the complexity bound proven in [3] for genus-3 hyperelliptic curves with explicit RM.

For hyperelliptic curves with RM, we have thus been able to eliminate the dependency in g in the exponent of $\log q$, but this does not mean that our algorithm reaches polynomial-time complexity in both g and $\log q$. Indeed, we also discuss the reasons why the “constant” K depends exponentially on g . Among them, we shall see that some can actually be discarded by considering even more particular cases while some appear to be inherent to our geometric-resolution based approach.

Organization. In Section 2, we give an overview of our point-counting algorithm, along with an example of families of hyperelliptic curves of arbitrary high genus with RM by a real subfield of a cyclotomic field. In particular, we prove a bound on the size and number of primes ℓ to consider in our algorithm. Section 3 focuses on the main primitive of our algorithm: the computation of a non-zero element in the kernel of an endomorphism α whose degree is a small multiple of ℓ^2 . This section adapts methods and results of [2, Sec. 4 & 5] to design structured polynomial systems whose solution sets are subsets of $J[\alpha]$. Section 4 concludes on the complexity of solving these systems, and on the overall complexity of our point-counting algorithm. We also present an analysis on the dependency of the final complexity in g , investigating the various places where exponential factors may occur and how to avoid them when it is possible.

Acknowledgements. Most of this work already appears as Chapter VII in the author’s thesis manuscript [1]. As such, the author received helpful feedback from his advisors Pierrick Gaudry and Pierre-Jean Spaenlehauer; and from his thesis referees Christophe Ritzenthaler and Frédérik Vercauteren. The author is also grateful to Benjamin Smith and David Kohel for pointing out references and for fruitful discussions.

2 Overview

The main result of this paper can be summarized by the following theorem, in which we give more precision on the notation $O_\eta(\log^c q)$ for our complexity result, and make the dependency in η explicit. In Section 4, we also bound c by 8 and conjecture that it should be 6. Note that whenever we give a bound with an explicit constant, we can no longer hide the polylogarithmic factor in the exponent, so we use the notation $\tilde{O}_\eta()$ to hide both factors depending only on η and factors that are polylogarithmic in q .

Theorem 1. *For any g and any $\eta \in \overline{\mathbb{Q}}$ such that $\mathbb{Q}(\eta)$ is a totally-real number field of degree g , there exists an explicitly computable $c(\eta) > 0$ such that there is an integer $q_0(g, \eta)$ such that for all prime power $q = p^n$ larger than $q_0(g, \eta)$ with $p \geq (\log q)^{c(\eta)}$ and for all genus- g hyperelliptic curves \mathcal{C} with explicit RM by $\mathbb{Z}[\eta]$ defined over \mathbb{F}_q , the local zeta function of \mathcal{C} can be computed with a probabilistic algorithm in expected time bounded by $(\log q)^{c(\eta)}$.*

2.1 Families of RM curves

We present one-dimensional families of hyperelliptic curves from [28], constructed via cyclotomic covers. They have an affine model $\mathcal{C}_{n,t} : Y^2 = D_n(X) + t$, where t is a parameter and D_n is the n -th Dickson polynomial with parameter 1 defined inductively by $D_0(X) = 2$, $D_1(X) = X$, and

$$D_n(X) = XD_{n-1}(X) - D_{n-2}(X).$$

Since $D_n(X)$ has degree n , setting $n = 2g + 1$ for odd n yields a one-dimensional family $\mathcal{C}_{n,t}$ of genus g hyperelliptic curves given by an odd-degree Weierstrass model. Their Jacobians all have an explicit endomorphism η , and when n is prime, [19, Prop. 2] shows that $\mathbb{Z}[\eta] \cong \mathbb{Z}[\zeta_n + \zeta_n^{-1}]$, where ζ_n is a primitive n -th root of unity over \mathbb{Q} . Note that the construction of an explicit endomorphism is still possible whenever $n = 2g + 1$ is not prime, but then the curves in $\mathcal{C}_{n,t}$ have non-simple Jacobians, which means there are better alternatives than using our algorithm for counting points on them.

Another family based on Artin-Schreier covering is detailed in the same paper but these curves have genus $(p - 1)/2$ where p is the characteristic of the base field, so that our complexity study using the $O_\eta()$ notation would be pointless in that case. Since g becomes much larger than $\log p$ in that case, it would be more efficient to use p -adic algorithms anyway.

Let \mathcal{C} be a (genus- g) hyperelliptic curve in the family $\mathcal{C}_{2g+1,t}$, defined over a finite field \mathbb{F}_q . In [19], Kohel and Smith compute formulas for the Mumford form of $\eta((x, y) - P_\infty)$, where (x, y) is the generic point on \mathcal{C} . These formulas are given explicitly for some examples in genus 2 and 3, and an algorithm [19, Algorithm 5] is presented to compute them for any \mathcal{C} . This algorithm has a time complexity in $O(g^2)$ field operations and requires to store $O(g^3)$ field elements. Thus, given a curve from that family as input, an explicit endomorphism of its Jacobian can be computed once and for all in $\tilde{O}(g^3 \log q)$ time and space complexity, which is negligible compared to the cost of counting points on the curve.

2.2 The characteristic equation

As in [12, 3], let us consider $\psi = \pi + \pi^\vee$ and recall that $\psi \in \mathbb{Q}[\eta]$. We still have $\psi\pi = \pi^2 + q$ and once again, we test this equation to determine ψ instead of the

characteristic equation of π . The link between ψ and π needs to be made explicit, which is the aim of the present section.

Since χ_π is a Weil polynomial, we can write

$$\chi_\pi(X) = \sum_{i=0}^g (-1)^i \sigma_i (X^{2g-i} + q^{g-i} X^i),$$

with $\sigma_0 = 1$ and the convention that σ_g is actually twice smaller than the g -th coefficient of χ_π . Note that we have $q^{-g}(\pi^\vee)^g \chi_\pi(\pi) = 0$ by the Cayley-Hamilton theorem, and using the fact that $\pi\pi^\vee = q$, we rewrite that as

$$\sum_{i=0}^g (-1)^{g-i} \sigma_{g-i} (\pi^i + (\pi^\vee)^i) = 0.$$

Our plan is to compute $\chi_\pi \bmod \ell$ by determining ψ . Let us write $\psi = \sum_{i=0}^{g-1} a_i \eta^i$, the goal of the section is to prove bounds on the coefficients a_i , so that we can estimate the number and maximal size of primes ℓ required to compute ψ without ambiguity. Note that ψ is in the maximal order of $\mathbb{Q}(\eta)$, but not necessarily in $\mathbb{Z}[\eta]$. However, as in [12, 3], $\mathbb{Z}[\eta]$ has finite index Δ in the maximal order and the possible common denominator of the a_i 's has to divide Δ . This denominator entails that additional primes may be required to fully determine ψ , however Δ depends only on η so that it will disappear in the O_η -notation of our complexity estimates. Therefore, we do not detail further this subtlety and assume for simplicity that the a_i 's are integers, which we wish to bound by $O_\eta(\sqrt{q})$.

Let us first express the quantities $\pi^i + (\pi^\vee)^i$ in terms of powers of ψ as a first step towards expressing the σ_i 's as functions of the a_i 's.

Lemma 2. *For any $i \in \{1, \dots, g\}$, there exist integers $(\alpha_{i,j})_{0 \leq j < i}$ such that $\alpha_{i,j} = O(q^{(i-j)/2})$ and*

$$\pi^i + (\pi^\vee)^i = \psi^i + \sum_{j=0}^{i-1} \alpha_{i,j} \psi^j.$$

Proof. The statement holds for $i = 1$ with $\alpha_{1,0} = 0$ by the definition of ψ . For $i = 2$, we have $\psi^2 = \pi^2 + (\pi^\vee)^2 + 2\pi\pi^\vee$, so that we have the result with $\alpha_{2,0} = -2q$ and $\alpha_{2,1} = 0$.

In this proof, we set the convention $\alpha_{i,i} = 1$ to simplify our recurrence relations. Let us now assume the lemma holds for any positive integer no greater than a certain i . We therefore have

$$\psi^{i+1} = (\pi + \pi^\vee)\psi^i = (\pi + \pi^\vee) \left[(\pi^i + (\pi^\vee)^i) - \sum_{j=0}^{i-1} \alpha_{i,j} \psi^j \right].$$

The first term is equal to $\pi^{i+1} + (\pi^\vee)^{i+1} + q(\pi^{i-1} + (\pi^\vee)^{i-1})$ so that we can use the lemma once again for $i - 1$ and get

$$\psi^{i+1} = \pi^{i+1} + (\pi^\vee)^{i+1} - \alpha_{i,i-1} \psi^i + q\alpha_{i-1,0} + \sum_{j=1}^{i-1} (q\alpha_{i-1,j} - \alpha_{i,j-1}) \psi^j.$$

Thus, we have computed the $\alpha_{i+1,j}$ given by

$$\alpha_{i+1,j} = \begin{cases} \alpha_{i,i-1} & \text{if } j = i, \\ -q\alpha_{i-1,0} & \text{if } j = 0, \\ \alpha_{i,j-1} - q\alpha_{i-1,j} & \text{else.} \end{cases}$$

Let us now study the order of magnitude of the $\alpha_{i+1,j}$: from the recurrence hypothesis on both i and $i-1$, $\alpha_{i,i-1} = \alpha_{i+1,i}$ is in $O(\sqrt{q})$, $\alpha_{i-1,0}$ is in $O(q^{(i-1)/2})$ so that $\alpha_{i+1,0}$ is in $O(q^{(i+1)/2})$, and both $q\alpha_{i-1,j}$ and $\alpha_{i,j-1}$ are in $O(q^{(i+1-j)/2})$, which proves the result for any other $\alpha_{i+1,j}$. By induction, the lemma is proven. \square

Note that our O -notation in the previous statement and proof can be a bit misleading as there may not be an absolute constant bounding all the $\alpha_{i,j}/q^{(i-j)/2}$. However, from the recurrence relation between the $a_{i,j}$'s, one sees that each $\alpha_{i,j}$ is equal to $q^{(i-j)/2}$ plus an error term that is in $O_\eta(q^{i-j-1}/2)$ and at worst quadratic in g , hence the error term is negligible compared to $q^{(i-j)/2}$.

Proposition 3. *The polynomial χ_π is uniquely determined by the coefficients a_i 's of ψ in the basis $(1, \eta, \dots, \eta^{g-1})$, and there exists $C_\eta > 0$ depending only on g and η such that for any $i \in \{0, \dots, g-1\}$, we have $|a_i| \leq C_\eta \sqrt{q}$.*

Proof. Recall that σ_i is the i -th coefficient of χ_π , or half this coefficient if $i = g$. Using Lemma 2 for any $i \in \{1, \dots, g\}$ and setting $\alpha_{i,i} = 1$, we have

$$\sum_{i=0}^g (-1)^{g-i} \sigma_{g-i} \sum_{j=0}^i \alpha_{i,j} \psi^j = \sum_{j=0}^g \psi^j \sum_{i=j}^g (-1)^{g-i} \alpha_{i,j} \sigma_{g-i} = 0.$$

Let us define $s_i = \sum_{j=i}^g (-1)^{g-i} \alpha_{i,j} \sigma_{g-i}$ and $\chi_\psi(X) = X^g + s_{g-1}X^{g-1} + \dots + s_0$. Invoking the Weil conjectures for the σ_{g-i} 's and Lemma 2 for the $\alpha_{i,j}$, one concludes that each s_i is in $O(q^{(g-i)/2})$. Furthermore, the expressions of the s_i 's in terms of the σ_i 's form a linear triangular system whose determinant equals 1, so that there is an efficiently computable one-to-one correspondence between χ_ψ and χ_π .

Let us now explicit the link between the coordinates a_i of $\psi = \sum_{i=0}^{g-1} a_i \eta^i$ and the coefficients s_i of χ_ψ . For instance, we have $s_{g-1} = -\text{Tr}(\psi) = -\sum_{i=0}^{g-1} a_i \text{Tr}(\eta^i)$. To get the other relations, let us now order the g conjugates of η (possibly in the Galois-closure of $\mathbb{Q}(\eta)$), numbering them from η_1 to η_g , and proceed to the linear change of variables $\psi_k = \sum_{i=0}^{g-1} a_i \eta_k^i$ for any k in $\{1, \dots, g\}$. The matrix associated to this linear transformation is the Vandermonde matrix of the conjugates η_k 's. This matrix is invertible because η is separable so that the η_i are all distinct reals.

Note that χ_ψ is a degree- g monic polynomial vanishing on ψ , and it is therefore its characteristic polynomial. Since the ψ_k are exactly the real roots (possibly in the Galois-closure of $\mathbb{Q}(\eta)$) of χ_ψ , by Vieta's formula they satisfy the g equations

$$s_{g-i} = (-1)^i S_i(\psi_1, \dots, \psi_g) \text{ for } 1 \leq i \leq g,$$

where the S_i 's are the elementary symmetric polynomials in g variables. Thus, once the a_i 's are known, the values for ψ and its conjugates are known and a unique value for each s_i is deduced. Furthermore, the Fujiwara bounds from [10] imply that for any $k \in \{1, \dots, g\}$ we have

$$|\psi_k| \leq 2 \max_{0 \leq k \leq g} (|s_{g-k}|^{1/k}).$$

We already know that $|s_{g-k}| = O(\sqrt{q}^k)$, so we deduce that the $|\psi_k|$ are in $O(\sqrt{q})$. Then, inverting the linear change of variable, we prove that the a_i are also in $O_\eta(\sqrt{q})$ since the matrix norm of the inverse of the Vandermonde matrix only depends on η . \square

Our algorithm is based on determining the a_i 's modulo ℓ for sufficiently many ℓ until they are known without ambiguity and we can deduce χ_π . While the Weil bounds on the σ_i 's are enough for our purpose, we have proven that the a_i 's are in $O_\eta(\sqrt{q})$ as in genus 2 and 3 [12, 3]. The next section details the process of recovering such modular information on the a_i 's.

2.3 Overview of our algorithm

The general RM point counting algorithm is Algorithm 1. As mentioned above, we want to compute the coefficients a_0, \dots, a_{g-1} of the endomorphism ψ . More precisely, we compute their values modulo sufficiently many totally-split primes ℓ until we can deduce their values from the bounds of Prop 3 and the Chinese Remainder Theorem. Then, the coefficients of χ_π are deduced from the a_i 's.

input : q an odd prime power, and $f \in \mathbb{F}_q[X]$ a monic squarefree polynomial of degree $2g + 1$ such that the hyperelliptic curve $Y^2 = f(X)$ has explicit RM by $\mathbb{Z}[\eta]$.

output: The characteristic polynomial $\chi_\pi \in \mathbb{Z}[T]$ of the Frobenius endomorphism on the Jacobian J of the curve.

$w \leftarrow 1$;
 Define C_g as in Prop. 3;
while $w \leq 2 \Delta C_g \sqrt{q} + 1$ **do**
 Pick the next prime ℓ that satisfies conditions (C1) to (C4);
 Compute the ideal decomposition $\ell \mathbb{Z}[\eta] = \mathfrak{p}_1 \cdots \mathfrak{p}_g$, corresponding to the eigenvalues $\lambda_1, \dots, \lambda_g$ of η in $J[\ell]$;
 for $i \leftarrow 1$ **to** g **do**
 Compute a small element α_i of \mathfrak{p}_i as in Lemma 4;
 Compute a non-zero element D_i of order ℓ in $J[\alpha_i]$;
 Find the unique $k_i \in \mathbb{Z}/\ell\mathbb{Z}$ such that $k_i\pi(D_i) = \pi^2(D_i) + qD_i$;
 end
 Find the unique tuple (a_0, \dots, a_{g-1}) in $(\mathbb{Z}/\ell\mathbb{Z})^g$ such that $\sum_{j=0}^{g-1} a_j \lambda_i^j = k_i$,
 for i in $\{1, \dots, g\}$;
 $w \leftarrow w \cdot \ell$;
end
 Reconstruct (a_0, \dots, a_{g-1}) using the Chinese Remainder Theorem ;
 Deduce χ_π from ψ .

Algorithm 1: Overview of our RM point-counting algorithm

We now explain how the algorithm works for a given split ℓ . First its decomposition as a product of prime ideals $\ell \mathbb{Z}[\eta] = \mathfrak{p}_1 \cdots \mathfrak{p}_g$ is computed, and for each prime ideal \mathfrak{p}_i , a non-zero element α_i in \mathfrak{p}_i is found with a small representation as in Lemma 4 below. 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 viewed as an endomorphism of $J[\ell] \cong (\mathbb{Z}/\ell\mathbb{Z})^{2g}$.

On $G_i \subset J[\alpha_i]$, the endomorphism η acts as the multiplication by λ_i . Therefore, the endomorphism $\psi = \sum_{i=0}^{g-1} a_i \eta^i$ also acts as a scalar multiplication on this 2-dimensional space, and we write $k_i \in \mathbb{Z}/\ell\mathbb{Z}$ 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

$$k_i \pi(D_i) = \pi^2(D_i) + qD_i. \quad (1)$$

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 the a_j 's modulo ℓ :

$$\sum_{j=0}^{g-1} a_j \lambda_i^j \equiv k_i \pmod{\ell}.$$

Therefore we have a linear system of g equations in g 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 a modelling similar to that of [2]. Likewise, we use the geometric resolution algorithm to find the solutions of this system, perhaps in a finite extension of the base field, 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 non-trivial one.

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 4 below with norm non-divisible by ℓ^3 for $i \in \{1, \dots, g\}$;
- (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 . The condition (C3) implies that there is a unique subgroup G_i of order ℓ^2 in $J[\alpha_i]$.

Given a genus- g curve \mathcal{C} with RM by $\mathbb{Z}[\eta]$, by Chebotarev's density theorem, the proportion of primes ℓ satisfying the last condition is at least $1/\#\text{Gal}(\mathbb{Q}(\eta)/\mathbb{Q})$, which is bounded below by $1/(g!)$. To count points on \mathcal{C} , we need to find L a set of primes satisfying all the above conditions and such that $\prod_{\ell \in L} \ell > 2\Delta C_\eta \sqrt{q}$. By the prime number theorem, both the number and size of the primes contained in L are in $O((g!) \log(C_g q))$. In some particular cases, the proportion of "nice" primes may be much larger: for instance when the RM field is the totally real subfield of a cyclotomic field. In the field

$Q(\zeta_n + \zeta_n^{-1})$, a prime ℓ totally splits if and only if $\ell \equiv \pm 1 \pmod n$, and therefore condition (C4) is satisfied by a proportion of primes equal to $2/(n-1) = 1/g$. In that case, the number and size of primes in the set L can be reduced to $O(g \log(C_g q))$.

Lemma 4. *For any prime ℓ that splits completely in $\mathbb{Z}[\eta]$, each prime ideal \mathfrak{p} above ℓ contains a non-zero element α of the form $\alpha = \sum_{i=0}^{g-1} \alpha_i \eta^i$, where the $|\alpha_i|$ are integers smaller than $\Delta^{1/g} \ell^{1/g}$, where we recall that Δ is the index $[\mathcal{O}_{\mathbb{Q}(\eta)} : \mathbb{Z}[\eta]]$.*

Proof. The coefficients of the elements of the ideal \mathfrak{p} represented by polynomials in η form a lattice L of dimension g . In $\mathbb{Z}[\eta]$, its volume is the norm of \mathfrak{p} , i.e. ℓ . Thus, its actual volume in \mathbb{R}^g is $\ell \Delta$. Let us consider $C = \{x \in \mathbb{R}^g \mid \|x\|_\infty \leq \Delta^{1/g} \ell^{1/g}\}$. The volume of the convex C is $2^g \Delta \ell$. Since g is the dimension of L and $\Delta \ell$ its volume, Minkowski's theorem guarantees the existence of a non-zero element v of L belonging to C . By definition, $v = \sum_{i=0}^{g-1} v_i \eta^i$ is an element of \mathfrak{p} whose coordinates v_i 's are integer of absolute values bounded by $\Delta^{1/g} \ell^{1/g}$, which concludes the proof. \square

Since we know it exists, given one of the ideals \mathfrak{p}_i , we can find α_i a small element of \mathfrak{p}_i as in Lemma 4 by exhaustive search in at most $2^g \Delta \ell$ operations in $\mathbb{Z}[\eta]$. Note that there is an extensive literature on finding short vectors in a lattice of dimension d , motivated for instance by cryptographic applications. An example is the quantum algorithm of [9] which computes a $2^{\tilde{O}(\sqrt{d})}$ -approximation of the shortest non-zero vector in time polynomial in d . Restricting to classical algorithms, the best option in general is the BKZ algorithm [25] that computes a $2^{\tilde{O}(d^\alpha)}$ -approximation in time $2^{\tilde{O}(d^{1-\alpha})}$, for any $\alpha \in [0, 1]$. In our case however, the existence of a very short vector is already known and, more importantly, the factor 2^g due to the dimension is acceptable since it vanishes in the O_η -notation.

3 Modelling kernels of endomorphisms

Let α be an explicit endomorphism of degree $O(\ell^2)$ on the Jacobian of \mathcal{C} , which satisfies the properties of Lemma 4. We want to compute a polynomial system that describes the kernel $J[\alpha]$ of α , and then solve it. The resultant-based approach of [3] cannot be used as the degrees are squared each time we eliminate a variable, causing an exponential dependency in g in the exponent of ℓ . Instead, we use the modelling techniques from [2], where the endomorphism α replaces the multiplication by ℓ . This time, the g variables of large degrees have degrees in $O_\eta(\ell^{3/g})$ instead of $O_\eta(\ell^3)$ so that the final complexity bound for computing the kernel α is in $O_\eta(\ell^c)$, with c an absolute constant.

The main change between this section and [2, Sec. 4 & 5] is that the d_i and e_i no longer denote ℓ -division but α -division polynomials, and the polynomials u_j and v_j intervening in the Mumford representation of the candidate kernel element are modified accordingly. The structure of our modelling is very similar but require some adaptations at various places, which is the reason why we repeat the analysis in the generic case. In the non-generic case, we restate the main results of [2, Sec. 5] but only detail the parts requiring adjustments.

3.1 The generic case

Let us first recall the definition of Cantor's ℓ -division polynomials, the coefficients of the polynomials $\delta_\ell(X)$ and $\varepsilon_\ell(X)$ such that, for (x, y) a generic point of the curve and $\ell > g$, we have

$$\ell((x, y) - P_\infty) = \left\langle \delta_\ell \left(\frac{x - X}{4y^2} \right), \varepsilon_\ell \left(\frac{x - X}{4y^2} \right) \right\rangle.$$

An important step towards our complexity bounds is to bound the degrees of these polynomials, so that we can later on deduce degree-bounds for the polynomial systems modelling the $J[\alpha_i]$. To this end, we use the following result proven in [2, Sec. 6].

Theorem 5. [2, Lemma 10] *For any integer $\ell > g$, the polynomial $\delta_\ell(X)$ of degree g in X has coefficients in $\mathbb{F}_q[x]$ whose degrees in x are bounded by $g\ell^3/3 + O_g(\ell^2)$; the polynomial $\varepsilon_\ell(X)/y$ has coefficients in $\mathbb{F}_q(x)$ such that the degrees of the numerators and the denominators have degrees bounded by $2g\ell^3/3 + O_g(\ell^2)$. Furthermore, the roots of the denominators are roots of the leading coefficient of $\delta_\ell(X)$.*

These polynomials describe the multiplication by ℓ , but for our purpose we need to define the α -division polynomials d_i and e_i such that, denoting by $P = (x, y)$ the generic point of \mathcal{C} , the non-normalized Mumford form of $\alpha(P - P_\infty)$ is equal to

$$\left\langle \sum_{i=0}^g d_i(x) X^i, y \sum_{i=0}^{g-1} \frac{e_i(x)}{e_g(x)} X^i \right\rangle.$$

By Lemma 4, we know that $\alpha = \sum_{i=0}^{g-1} \alpha_i \eta^i$ with $|\alpha_i| = O_\eta(\ell^{1/g})$. Since the degrees of the $\eta^i(P - P_\infty)$ do not depend on ℓ , by Theorem 5 applied to Cantor's α_i -division polynomials we prove that the degrees of the d_i 's and e_i 's are in $O_\eta(\ell^{3/g})$.

Definition 6. *In what follows, we will say that an element of J is α -generic if it has weight g and the corresponding reduced divisor $\sum_{i=1}^g (P_i - P_\infty)$ satisfies the following two properties:*

- *For any i , the u -coordinate of the divisor $\alpha(P_i - P_\infty)$ in Mumford form has degree g ;*
- *For any $i \neq j$, the u -coordinates of the divisors $\alpha(P_i - P_\infty)$ and $\alpha(P_j - P_\infty)$ are coprime.*

This implies that if an affine point P occurs in the support of $\alpha(P_i - P_\infty)$ then neither P nor $-P$ appears in the support of another $\alpha(P_j - P_\infty)$.

Let $D = \sum_{i=1}^g (P_i - P_\infty)$ be an α -generic divisor in J . We shall consider a system equivalent to $\alpha(D) = 0$ but let us first introduce some notation. For each point $P_i = (x_i, y_i)$ in the support of D , we denote $\langle u_i, v_i \rangle$ the Mumford form of $\alpha(P_i - P_\infty)$ and $(a_{ij}, b_{ij})_{1 \leq j \leq g}$ the coordinates of the g points in its support counted with multiplicities, which means that for any i the g roots of u_i are exactly the a_{ij} , and that for any j , $b_{ij} = v_i(a_{ij})$.

Proposition 7. *We can model the set of generic α -division elements as the solution set of a bihomogeneous polynomial system consisting of $O(g^2)$ equations in $\mathbb{F}_q[X_1, \dots, X_g, Y_1, \dots, Y_{n_y}]$ such that $n_y = O(g^2)$ and the degrees d_x and d_y in the X_i 's and Y_j 's are respectively in $O_\eta(\ell^{3/g})$ and $O_\eta(1)$.*

Proof. Following the modelling of [2, Sec. 4], we have $\alpha(D) = 0$ if and only if the sum of the divisors $\sum_{i=1}^g \alpha(P_i - P_\infty)$ is a principal divisor. The only pole is at infinity, so this is equivalent to the existence of a non-zero function $\varphi \in \mathbb{F}_q(\mathcal{C})$ of the form $P(X) + YQ(X)$ with P and Q two polynomials such that the g^2 points (a_{ij}, b_{ij}) are the zeros of φ , with multiplicities. Since we want φ to have g^2 affine points of intersection with the curve \mathcal{C} (once again, counted with multiplicities), the polynomial $\text{Res}_Y(Y^2 - f, P + YQ) = P^2 - fQ^2$ must have degree g^2 which yields $2 \deg(P) \leq g^2$ and $2 \deg(Q) \leq g^2 - 2g - 1$. Exactly one of those two bounds is even (it depends on the parity of g), and for this particular bound, the inequality must be an equality, otherwise the degree of the resultant would not be g^2 . Since the function φ is defined up to a multiplicative constant, we can normalize it so that the polynomial $P^2 + fQ^2$ is monic, which is equivalent to enforce that either P or Q is monic depending on the parity of g .

For a fixed $i \in [1, g]$, requiring the (a_{ij}, b_{ij}) to be zeros of φ amounts to asking for the a_{ij} to be roots of $P(X) + Q(X)v_i(X)$, with multiplicities. Since the a_{ij} are by definition the roots of the u_i , $\alpha(D) = 0$ is equivalent to g congruence relations $P + Qv_i \equiv 0 \pmod{u_i}$. Thus, for any α -generic divisor, $\alpha(D) = 0$ is equivalent to the existence of P and Q satisfying the above g congruence relations.

The variables are the coefficients of P and Q , as well as the x_i and y_i . With the degree conditions and the normalization, we have $g^2 - g$ variables coming from P and Q . Adding the $2g$ variables x_i and y_i , we get a total of $g^2 + g$ variables. Each one of the g congruence relations amounts to g equations providing a total of g^2 conditions on the coefficients of P and Q . The fact that the (x_i, y_i) are points of the curve yields the g additional equations $y_i^2 = f(x_i)$. Finally, we have to enforce the α -genericity of the solutions, which can be done by requiring that $\prod_i d_g(x_i) e_g(x_i) \prod_{i < j} \text{Res}(u_i, u_j) \neq 0$. Note that we do not extend Theorem 5 to the α -division polynomials but instead add the non-vanishing condition for the denominator e_g of the v -coordinate of $\alpha(D)$. Still, we get a polynomial system with $g^2 + g$ equations in $g^2 + g$ variables, together with an inequality.

We now estimate the degrees to which the variables occur in the equations. Each congruence relation is obtained by reducing $P + Qv_i$, which is a polynomial of degree $O(g^2)$ in X , by u_i which is of degree g . We can do it by repeatedly replacing X^g by $-\sum_{j < g} (d_j(x_i)/d_g(x_i))X^j$, which we will have to do at most $O(g^2)$ times. Since the d_j have degree in $O_\eta(\ell^{3/g})$ in x_i , the fully reduced polynomial will have coefficients that are fractions for which the degrees of the numerators and of the denominators are at most $O_\eta(\ell^{3/g})$ in the x_i variables. In these equations, the degree in the y_i variables and in the variables for the coefficients of P and Q is 1. The degrees in x_i and y_i in the curve equations are $2g + 1$ and 2 respectively.

It remains to study the degree of the inequality. Each resultant is the determinant of a $2g \times 2g$ Sylvester matrix whose coefficients are the d_i , which have degrees bounded by $O_\eta(\ell^{3/g})$. Since for any i there are exactly g resultants involving x_i in the product, the degree of this inequality in any x_i is in $O_\eta(\ell^{3/g})$, and it does not involve the other variables. In order to be able to use Proposition [2, Prop. 3] that we recall in Section 4, we must model this inequality by an equation, which is done classically by introducing a new variable T and by using the equation $T \cdot \prod_i d_g(x_i) e_g(x_i) \prod_{i < j} \text{Res}(u_i, u_j) = 1$.

To conclude, we have a polynomial system with two blocks of variables: the g variables x_i on the one hand and the $g^2 - g$ variables coming from the coefficients of P and Q , along with the g variables y_i on the other hand. The degrees of the equations

in the first block of variables grows cubically in $\ell^{1/g}$, while the degrees in the other block of variables depends only on η . \square

3.2 Non-generic kernel elements

As in [2, Sec. 4], apart from the neutral element, we expect to capture the whole kernel of the endomorphism α by using the modelling of Section 3.1. Contrary to [2], Algorithm 1 does not require us to find a basis of $J[\alpha]$ because the determination of the k_i 's does only require a single non-zero element in each $J[\alpha_i]$. Thus, a study of non-generic elements in $J[\alpha]$ is necessary only if there is no α -generic element in $J[\alpha]$. Such a case happens if and only if the polynomial $\prod_{i=1}^g d_g(x_i)e_g(x_i) \prod_{i \neq j} \text{Res}(u_i, u_j)$ in the variables x_1, \dots, x_g vanishes on $J[\alpha]$. It seems very unlikely that the whole set $J[\alpha]$ would live in such a hypersurface, and if it happens, one can discard the ℓ for which we fail to find an α -generic element. Although it seems even more unlikely that this situation could happen for sufficiently many ℓ so as to threaten the validity of our complexity bound, we are far from a proven statement and do not exclude it might be possible to design a highly non-generic curve providing a counterexample.

Therefore, we follow the non-genericity analysis of [2, Sec. 5] except that we consider u_i and v_i defined as the Mumford form of $\alpha(P_i - P_\infty)$ instead of $\ell(P_i - P_\infty)$. Let us briefly review the non-generic situations that one can encounter, following [2, Sec. 5.1] and keeping the same numbering.

Case 1: Modelling a kernel element of weight $w < g$. We write $D = \sum_{i=1}^w (P_i - P_\infty)$ and look for a $\varphi = P(X) + YQ(X)$ vanishing at each point of each reduced divisor $\alpha(P_i - P_\infty)$. This is similar to the Case 1 of [2, Sec. 5.1].

Case 2: Modelling a kernel element with multiple points. It may happen that the element we are looking for is $D = \sum_{i=1}^w (P_i - P_\infty)$ but not all the P_i 's are distinct. In that case, we rewrite it $D = \sum_{j=1}^s \lambda_j (P_j - P_\infty)$ such that the P_j 's are distinct and look for a $\varphi = P(X) + YQ(X)$ vanishing at each point of each reduced divisor $\lambda_j \alpha(P_j - P_\infty)$. Apart from the modification of u_i and v_i , the modelling is identical to that of [2].

Case 4: Modelling a kernel element after reduction. Even if all the $\alpha(P_i - P_\infty)$ had full weight, there may still be less than g^2 points in the union of their supports due to possible cancellations of points appearing in the supports of several $\alpha(P_i - P_\infty)$ with different signs. Exactly as in [2, Sec. 5.1], if P appears within $\alpha(P_i - P_\infty)$ and $\alpha(P_j - P_\infty)$ with respective multiplicities ν_i and ν_j of opposite signs, this is modelled by ensuring that the corresponding u_i, u_j , and $v_i + v_j$ share a common factor $(X - \xi)^\nu$ where $\nu = \max(|\nu_i|, |\nu_j|)$. In that case, we look for $\varphi(X, Y) = (X - \xi)^\nu (\tilde{P}(X) + Y\tilde{Q}(X))$, with \tilde{P} coprime to \tilde{Q} . Once modified the values of the u_i and v_i , nothing changes from [2].

Case 5: Modelling a kernel element with multiplicity. Conversely, $\alpha(P_i - P_\infty)$ and $\alpha(P_j - P_\infty)$ can also share the same point with multiplicities of identical sign, leading to multiplicities in the reduced divisor $\alpha(D)$. Similarly to what was done in the Case 5 of [2, Sec. 5.1], we can group the corresponding u_i, u_j, v_i and v_j in polynomials U and V such that $U|V^2 - f$ and $\deg V < \deg U$, and then look for

$\varphi = P(X) + YQ(X)$ such that $P + QV \equiv 0 \pmod{U}$. Once again, nothing changes apart from the definition of the u_i 's and v_i 's.

Case 3: Low weight after applying α . We kept this case for the end because it is not a straightforward extension of the Case 3 appearing in [2, Sec. 5.1]. Until now, we assumed that all the P_i 's in the support of D were such that $\alpha(P_i - P_\infty)$ had weight g , i.e. $d_g(x_i) \neq 0$. We now want to model the case where $D = \sum_{i=1}^w (P_i - P_\infty)$ such that each $\alpha(P_i - P_\infty)$ has weight w_i . In [2], this was done using a result from [7] giving a necessary and sufficient condition for $\ell(P_i - P_\infty)$ to be of weight w_i . When α is an endomorphism other than scalar multiplication, no such result holds a priori. In what follows, we address this issue by designing non-generic α -division polynomials $\Gamma_{\alpha,t}$ and $\Delta_{\alpha,t}$ such that $\alpha((x,y) - P_\infty)$ has weight w if and only if $\Delta_{\alpha,w}(x) = 0$ and $\Gamma_{\alpha,w-1}(x) \neq 0$.

Combining all degeneracies. As in [2, Sec. 5.2], we have to consider situations in which several of the previous cases occur simultaneously. Note that while we wanted to compute the whole ℓ -torsion in [2], we now only need one kernel element per endomorphism α_i to determine $\chi_\pi \pmod{\ell}$. Therefore, after finding a non-zero solution to any of the subsequent systems, one need not consider the others. Once again, we will not perform a complete analysis as in [2, Sec. 5.2] but rather detail when modifying the values of u_i and v_i is not sufficient. We also update the analysis on the numbers and degrees of equations and variables. The aim of the Section is to prove the following proposition.

Proposition 8. *We can model the set of non-generic elements of $J[\alpha]$ as the solution set of $O_\eta(1)$ bihomogeneous polynomial systems each consisting of $O(g^2)$ equations in $\mathbb{F}_q[X_1, \dots, X_g, Y_1, \dots, Y_{n_y}]$ such that $n_y = O(g^2)$ and the degrees d_x and d_y in the X_i 's and Y_j 's are respectively in $O_\eta(\ell^{3/g})$ and $O_\eta(1)$.*

Proof. As in [2], we encode each possible non-generic situation by a non-genericity tuple $(w, \lambda, \tau, \varepsilon, M)$ in the sense of Definition 9 below, and derive an associated polynomial system whose solution set corresponds to elements $D \in J[\alpha]$ such that:

- the reduced divisor D of weight w has the form $\sum_{i=1}^k \lambda_i P_i$ with distinct P_i 's,
- each $\lambda_i \alpha(P_i - P_\infty)$ has weight τ_i ,
- each ε_i is in $\{0, 1\}$ and such that $\varepsilon_i = 1$ if and only if $\tau_i = \lambda_i = 1$.
- the $k \times s$ matrix M represents the points shared by the $\lambda_i \alpha(P_i - P_\infty)$ as in [2, Sec. 5.2], with $s \leq gk$.

Definition 9. [2, Def. 13] *A normalized non-genericity tuple is a tuple $(w, \lambda, t, \epsilon, M)$, where $1 \leq w \leq g$ is an integer, $\lambda = (\lambda_1, \dots, \lambda_k)$ is a partition of w , t and ϵ are vectors $t = (t_1, \dots, t_k)$ and $\epsilon = (\epsilon_1, \dots, \epsilon_k)$ of the same length as λ with $1 \leq t_i \leq g$ and $\epsilon_i \in \{0, 1\}$, where ϵ_i can be 1 only if $t_i = 1$ and $\lambda_i = 1$, and finally M is a matrix with k rows and s columns, where $0 \leq s \leq gk$, and its entries are integers such that:*

- For all $1 \leq i \leq k$, the sum of the absolute values of the entries on the row i is equal to t_i ;

- The columns are sorted in lexicographical order;
- The sum of the rows of the matrix is a vector whose coordinates are nonnegative.

We can follow the analysis of [2, Sec. 5.2] to describe more explicitly the equations and their degrees / number of variables, and remark that the only part that does not generalize readily is the definition of non-generic α -division polynomials, as in the Case 3 above. Let us first fix this issue.

When the weight t_i of $\lambda_i \alpha(P_i - P_\infty)$ is strictly smaller than g , the usual coordinate system given by the Mumford form is no longer available, due to the vanishing of the denominator $e_g(x_i)$. We define an adequate coordinate system to describe non-generic elements of weight t . Let us consider the variety

$$V_{\alpha,t} = \{(x, y) \in \mathcal{C} \mid \alpha((x, y) - P_\infty) \text{ has weight } t\}.$$

We want to define polynomials $\Delta_{\alpha,t}$ and $\Gamma_{\alpha,t}$ such that a point is in $V_{\alpha,w}$ if and only if $\Delta_{\alpha,w}(x) = 0$ and $\Gamma_{\alpha,w-1}(x) \neq 0$ iteratively. First, $\Delta_{\alpha,g-1} = \text{GCD}(d_g, e_g)$, so that the points (x, y) of $V_{\alpha,g-1}$ satisfy $\Delta_{\alpha,g-1}(x, y) = 0$. Assuming that for $k < g$ we have already constructed a squarefree polynomial $\Delta_{\alpha,k}$ vanishing on the abscissae of points in $V_{\alpha,k}$, then one can compute $\alpha((x, y) - P_\infty)$ over $\mathbb{F}_p[x, y]/(\Delta_{\alpha,k}(x), y^2 - f(x))$. By our recurrence hypothesis, the Mumford form of the result is $\langle u, v \rangle$, with u of degree k and v of degree $k-1$. Let $\Gamma_{\alpha,k-1}$ be the product of $\text{LC}(u)$ with the denominator of $\text{LC}(v)$, then $V_{\alpha,k}$ is the set of points (x, y) such that $\Delta_{\alpha,k}(x) = 0$ and $\Gamma_{\alpha,k-1}(x) \neq 0$. Furthermore, $\Delta_{\alpha,k-1} = \text{GCD}(\Delta_{\alpha,k}, \Gamma_{\alpha,k-1})$ vanishes on the points of $V_{\alpha,k-1}$.

To avoid multiplicities, we replace $\Delta_{\alpha,t}(x)$ by the square-free polynomial whose roots are exactly the roots of $\Delta_{\alpha,t}(x)$ that are not roots of $\Gamma_{\alpha,t-1}(x)$ when it is necessary. Note that the degrees of the Δ and Γ are by construction bounded by $\deg \Delta_{\alpha,g-1} \leq \deg d_g$ with $\deg d_g$ itself bounded by $O_\eta(\ell^{1/g})$. This way, we state an analogue of [2, Def. 14] for non-generic α -division polynomials:

Definition 10. *The non-generic α -division polynomials $\mathbf{u}_{\alpha,t}$ and $\mathbf{v}_{\alpha,t}$ are the polynomials in X with coefficients in $\mathbb{F}_p[x, y]/(\Delta_{\alpha,t}(x), y^2 - f(x))$ such that*

$$\alpha((x, y) - \infty) = \langle \mathbf{u}_{\alpha,t}(X), \mathbf{v}_{\alpha,t}(X) \rangle,$$

in weight- t Mumford representation: $\mathbf{u}_{\alpha,t}(X)$ is monic of degree t , $\mathbf{v}_{\alpha,t}(X)$ is of degree at most $t-1$ and they satisfy $\mathbf{u}_{\alpha,t} \mid \mathbf{v}_{\alpha,t}^2 - f$.

All the equations associated to a non-genericity tuple $(w, \lambda, t, \epsilon, M)$ are merely identical to those of [2, Sec. 5.2] except that the d_i, e_i and have different definitions and that $\Delta_{\alpha,t}$ replaces $\hat{\Delta}_{\ell,t}$ so that Equation 3 of [2] now reads

$$\begin{cases} \Delta_{\lambda_i \alpha, t_i}(x_i) = 0, \\ \Gamma_{\lambda_i \alpha, t_i - 1}(x_i) \neq 0, \end{cases} \quad \text{for all } i \text{ in } [1, k] \text{ such that } t_i < g, \quad (\text{Sys.3b})$$

While turning the systems describing $J[\ell]$ into systems describing $J[\alpha]$, we did not add any new variable, so that the study of [2, Sec. 5.2] summed up in [2, Tab. 1] is still valid and we just recall that the total number of variables is bounded by $4g^2 + g$.

As for the number of equations and their respective degrees, the only difference with [2, Tab. 2] comes from the fact that the coefficients of the u_i and v_i have degrees

in the x_i bounded by $O_\eta(\ell^{3/g})$ instead of $O_\eta(\ell^3)$. In particular, there are at most $O(g^4)$ equations involving at most $O(g^2)$ variables, and apart from the x_i 's, the variables have degrees bounded by $O(g^3)$. This shows that any system corresponding to a non-genericity tuple satisfies the degree conditions of Proposition 8. As in [2], the number of such tuples is bounded by $g^{O(g^3)}$ and Proposition 8 is proved. \square

4 Complexity analysis

Now that we have modelled subsets of $J[\alpha]$ by polynomial systems whose sizes in terms of equations, variables and degrees have been carefully bounded, we apply the geometric resolution algorithm and bound its complexity.

4.1 Solving the polynomial systems modelling $J[\alpha]$

Just as in [2], we use geometric resolutions to describe 0-dimensional (i.e. finite) sets $V \subset \overline{\mathbb{F}_q}^n$ where V is defined over \mathbb{F}_q . The terminology here is borrowed from [6], see also [15].

Definition 11 (Geometric resolution). *An \mathbb{F}_{q^e} -geometric resolution of V is a tuple $((\ell_1, \dots, \ell_n), Q, (Q_1, \dots, Q_n))$ where:*

- The vector $(\ell_1, \dots, \ell_n) \in \mathbb{F}_{q^e}^n$ is such that the linear form

$$\begin{aligned} \ell : \quad \overline{\mathbb{F}_q}^n &\rightarrow \overline{\mathbb{F}_q} \\ (x_1, \dots, x_n) &\mapsto \sum_{i=1}^n \ell_i x_i \end{aligned}$$

takes distinct values at all points in V . The linear form ℓ is called the primitive element of the geometric resolution;

- The polynomial $Q \in \mathbb{F}_{q^e}[T]$ equals $\prod_{\mathbf{x} \in V} (T - \ell(\mathbf{x}))$;
- The polynomials $Q_1, \dots, Q_n \in \mathbb{F}_{q^e}[T]$ parametrize V by the roots of the polynomial Q , i.e.

$$V = \{(Q_1(t), \dots, Q_n(t)) \mid t \in \overline{\mathbb{F}_q}, Q(t) = 0\}.$$

We will need to bound the complexity of computing geometric resolutions of bihomogeneous polynomial systems. We do so by using a variant of [2, Prop. 3], which is restated here.

Proposition. [2, Prop. 3] *There exists a probabilistic Turing machine \mathbf{T} which takes as input polynomial systems with coefficients in a finite field \mathbb{F}_q and which satisfies the following property. For any function $h : \mathbb{Z}_{>0} \rightarrow \mathbb{Z}_{>0}$, for any positive number $C > 0$ and for any $\varepsilon > 0$, there exists a function $\nu : \mathbb{Z}_{>0} \rightarrow \mathbb{Z}_{>0}$ and a positive number $D > 0$ such that for all positive integers $g, \ell, n_x, n_y, d_x, d_y, m > 0$ such that $n_x < Cg$, $n_y < h(g)$, $d_x < h(g)\ell^C$, $d_y < h(g)$, $m < h(g)$, for any prime power q such that the prime number p dividing q satisfies $2^{n_x+n_y} d_x^{n_x} d_y^{n_y} < p$, and for any polynomial system $f_1, \dots, f_m \in \mathbb{F}_q[X_1, \dots, X_{n_x}, Y_1, \dots, Y_{n_y}]$ such that*

- for all $i \in [1, m]$, $\deg_x(f_i) \leq d_x$ and $\deg_y(f_i) \leq d_y$,
- the ideal $I = \langle f_1, \dots, f_m \rangle$ has dimension 0 and is radical,

the Turing machine \mathbf{T} with input f_1, \dots, f_m returns an $\mathbb{F}_{q^{\lceil \nu(g) \log \ell \rceil}}$ -geometric resolution of the variety $\{\mathbf{x} \in \overline{\mathbb{F}_q} \mid f_1(\mathbf{x}) = \dots = f_m(\mathbf{x}) = 0\}$ with probability at least $5/6$, using space and time bounded above by $\nu(g) \ell^{Dg} (\log q)^{2+\varepsilon}$.

Proof. This is done in [2, Sec. 3]. □

Remark Our systems are very similar to those presented in [2], which is the reason why we will be using the above proposition. In this paper, however, we bound d_x by some $h(g) \ell^{C/g}$ instead of $h(g) \ell^C$. Following the proof provided in [2, Sec. 3], the above proposition can be readily adapted and the factor $1/g$ in the exponent propagates which yields a final complexity bound bounded by $\nu(g) \ell^D (\log q)^{2+\varepsilon}$ (the exponent of ℓ is now a constant). This remark is essential to prove the following proposition.

Proposition 12. *For any $\varepsilon > 0$, there is a constant D such that for any endomorphism $\alpha \in \mathbb{Z}[\eta]$ of norm a multiple of $\ell > g$ coprime to the base field characteristic, there is a Monte Carlo algorithm which computes an \mathbb{F}_{q^e} -geometric resolution of the sub-variety of $J[\alpha]$ consisting of α -generic α -torsion elements, where $e = O_\eta(\log \ell)$. The time and space complexities of this algorithm are bounded by $O_\eta(\ell^D (\log q)^{2+\varepsilon})$ and it returns the correct result with probability at least $5/6$.*

Proof. Let us consider the sub-variety $S \subset J[\alpha]$ consisting of α -generic elements, and I the corresponding ideal. More precisely, we see I as the ideal of a sub-scheme of the scheme $J[\alpha]$, itself subscheme of $J[\deg \alpha]$, which is the kernel of a finite and étale map because $\deg \alpha$ is a small multiple of ℓ and is hence coprime to the characteristic p thanks to our assumptions on the size of p in the statement of Theorem 1.

Therefore, I is 0-dimensional and radical. Since all the elements in S have the same weight g we can use the Mumford coordinates $\langle u(X), v(X) \rangle$ with $\deg u = g$ and $\deg v < g - 1$ as a local system of coordinates to represent them. But the polynomial system that we have built is with the (x_i, y_i) coordinates, that is, it generates the ideal I^{unsym} obtained by adjoining to the equations defining I the $2g$ equations coming from $u(X) = \prod (X - x_i)$ and $y_i = v(x_i)$. Then we have $\deg I^{\text{unsym}} = g! \deg I$. By the α -genericity condition, all the fibers in the variety have exactly $g!$ distinct points corresponding to permuting the (x_i, y_i) which are all distinct. Therefore the radicality of I implies the radicality of I^{unsym} and we can apply the modified version of [2, Prop. 3] to our polynomial system.

Indeed, by Proposition 7 we now have a function h and a constant C such that $d_x \leq h(g) \ell^{C/g}$ instead of $h(g) \ell^C$. As we remarked, we can propagate this factor $1/g$ and compute an \mathbb{F}_{q^e} -geometric resolution of S in time and space bounded by $O_\eta(\ell^D (\log q)^{2+\varepsilon})$, with $e = O_\eta(\log \ell)$. □

Following the same proof but invoking Proposition 8 instead of Proposition 7, the same complexity bound holds for solving the polynomial system associated to any non-genericity tuple. Even if a non-zero α -torsion element is only found after solving all the systems associated to non-genericity tuples, the cost for computing $\psi \bmod \ell$ is only multiplied by a factor in $O_\eta(1)$.

4.2 An explicit bound for the exponent of $\log q$

We have proven that there exists a constant c such that for any prime ℓ satisfying conditions (C1) to (C4), computing $\psi \bmod \ell$ is achieved within $O_\eta(\ell^c)$ field operations.

Taking into account the size of the largest ℓ to consider and the cost of field operations, the overall complexity of our point-counting algorithm is in $O_\eta((\log q)^{c+2})$. The bottleneck is the computation of geometric resolutions of polynomial systems which is quadratic in the maximum of the degrees of the intermediate ideals $\langle f_1, \dots, f_i \rangle$, up to a factor in $O_\eta(1)$ (see for instance [15] for a detailed complexity analysis).

This quantity being hard to assess, we bound it by $2^{g+n_y} d_x^g d_y^{n_y}$ using [2, Prop. 8] (itself derived from [23, Prop. I.1]). Neglecting factors in $O_\eta(1)$, this bound boils down to $O_\eta(d_x^g)$. Keeping in mind that $d_x = O_\eta(\ell^{3/g})$, we can set $c = 6$ and deduce an overall complexity bound in $\tilde{O}_\eta(\log^8 q)$.

Note that our bound on d_x is pessimistic because we used the proven cubic bound for the degrees of Cantor's division polynomials while we expect them to be actually quadratic (see the final remark of [2, Sec. 6] for detailed experiments and conjectures). Under this assumption, d_x is reduced to $O_\eta(\ell^{2/g})$ and the overall complexity would therefore be in $\tilde{O}_\eta(\log^6 q)$ for any g . Apart from the hidden factor within the \tilde{O}_η notation, this conjectural result is identical to what was proven in genus 3 in [3]. In the next section, we push the analysis further by investigating the dependency in g of that factor.

4.3 Dependency in g of the complexity

The goal of this section is to assess the potential of our algorithm to achieve a polynomial-time complexity both in g and $\log q$ on some family of curves. To this end, we review our complexity analysis with additional attention given to the factors that previously vanished in the O_η .

Dependency in g of the largest ℓ Let us first come back to the constant C_η of Section 2.2. We have seen that the only non-polynomial dependency in g came from the matrix norm when inverting the linear change of variables $\psi_k = \sum_{i=0}^{g-1} a_i \eta_k^i$, which is described by the Vandermonde matrix of the g conjugates of η , denoted by η_k for $k \in \{1, \dots, g\}$. Let B be the inverse of this matrix, then we have

$$B_{ij} = \frac{\sum_{\substack{1 \leq k_1 < \dots < k_{g-j} < g \\ k_1, \dots, k_{g-j} \neq i}} (-1)^{j-1} \eta_{k_1} \cdots \eta_{k_{g-j}}}{\eta_i \prod_{k \neq i} (\eta_k - \eta_i)}.$$

Let $E = \max_k (|\eta_1|, \dots, |\eta_k|)$, $e = 1/\min_k (|\eta_1|, \dots, |\eta_k|)$, and $D = \max_{i \neq j} (|\eta_i - \eta_j|^{-1})$, then we can bound the absolute value of any entry of B very roughly either by $ge(2ED)^g$ or by ge if $2ED \leq 1$, and the matrix-norm of B is bounded by g times this previous bound. Note that the factor Δ is also a nuisance but it is bounded by the discriminant of $\mathbb{Z}[\eta]$. This discriminant is in turn bounded by $\max_{i \neq j} (|\eta_i - \eta_j|)^{2g}$. Thus, the constant C_η can be bounded by $g^2 c^g$, where c has a polynomial dependency in η and its conjugates.

By the prime number theorem, the set L of primes such that $\prod_{\ell \in L} \ell > 2C_\eta \sqrt{q}$ is such that the number and size of primes in L is in $\tilde{O}(g \log q)$. As we already mentioned, the primes to consider must satisfy the conditions (C1) to (C4) and that may cause them to be larger by a factor depending exponentially on g *a priori*. Since the complexity of computing $\chi_\pi \bmod \ell$ is polynomial in ℓ , this implies that the overall complexity depends exponentially in g in general.

However, a curve in the family $\mathcal{C}_{n,t}$ introduced in Section 2.1 has RM by the real subfield of $\mathbb{Q}(\zeta_n)$, for which we know that the proportion of split primes is $2/(n-1) = 1/g$. Therefore, this first obstacle due to the size of primes to consider can be overcome provided that we further strengthen the assumptions on the RM-curves we consider.

Finding small elements in lattices This time, the exhaustive search is no longer sufficient for our needs because of the exponential factor 2^g in the size of the ball $\{v \mid \|v\|_\infty \leq \Delta^{1/g} \ell^{1/g}\}$. Unfortunately, the currently known algorithms for finding short vectors in time subexponential in the dimension of the lattice have a drawback that makes them unusable in our point-counting algorithm. Indeed, although they run faster than the naive approach, they do not necessarily output the shortest non-zero vector on the lattice, but an approximation that may be greater by a factor which is also subexponential in the dimension. The size of the short vector plays a prominent role in the complexity analysis of our point-counting algorithm as it gives a bound on the degrees of the equations modelling $J[\alpha]$. Even if we find an α whose coordinates are in $C\ell^{1/g}$ instead of $C^{1/g}\ell^{1/g}$, the constant factor C will cause a factor C^g in the bound $2^{g+n_y} d_x^g d_y^{n_y}$, and hence in the final complexity of solving the polynomial systems.

Although finding short generators of ideal in number fields is believed to be hard in general, we may still expect to further restrict the RM curves we consider so as to fall in a case for which the complexity of such task becomes affordable. Examples are given in [5], where a classical algorithm is shown to compute short generators of principal ideals in particular number fields called multiquadratics, i.e. fields of the form $\mathbb{Q}(\sqrt{d_1}, \dots, \sqrt{d_n})$, in time quasipolynomial in the degree (which is g in our context). While we acknowledge that it is quite speculative to hope for families of curves of arbitrary high genus with RM by a $\mathbb{Z}[\eta]$ satisfying all the previous hypotheses, we do not linger on this because the next point is much more of a concern anyway.

Solving polynomial systems Using the strategy of Section 3, the complexity is quadratic in the bound $2^{g+n_y} d_x^g d_y^{n_y}$ of [2, Prop. 8], which includes a factor $g^{O(g^2)}$. Indeed, although the ideals of α -torsion have degree ℓ^2 independent of g , this is not true for the number of variables involved in our modelling, which is at least g^2 in the generic case.

However, even if none of the current complexity bounds for solving polynomial systems is sufficient to derive a polynomial-time algorithm both in g and $\log q$, there are still reasons to hope. Indeed, while the analysis made in [1] pointed out the fact that the systems themselves could have exponential size in g , these fears were based on very rough estimates of their size as straight-line programs. In fact, the cost of evaluating our equations of the form $P + Qv_i = 0 \pmod{u_i}$ can be split into two parts : first computing u_i and v_i , which amounts to computing $\alpha((x_i, y_i) - P_\infty)$ in $\mathbb{F}_q[x_i, y_i]/(y_i^2 - f(x_i))$. This is done within $O(\|\alpha\|_\infty/g \log \ell + g^2)$ operations on polynomials whose sizes are bounded by $O(g\|\alpha\|_\infty \ell^{3/g})$ field elements. Then, one has to finally reduce the degree- g^2 polynomial $P + Qv_i$ modulo the degree- g polynomial u_i , which can be done naively by replacing powers of X larger than g , for at most g^4 operations on polynomials of degrees $\leq g^2$ with coefficients in $\text{Frac}(\mathbb{F}_q[x_i, y_i]/(y_i^2 - f(x_i)))$ whose sizes are bounded by $2g^3 \ell^{3/g}$ field elements.

Thus, our systems have polynomial sizes in both g and $\log q$, which still fosters the hope that it could still be possible to solve them in time also polynomial in these pa-

rameters, although we recognize that improving on the estimate given by the multi-homogeneous Bézout bound would be a significant progress. Other possible workarounds to avoid an exponential dependency in g could be looking for easier instances in which we could model the α -torsion by even smaller polynomial systems, or cases for which there are simpler ways of obtaining a generic α -torsion divisor than the one we used.

5 Future work

Based on the facts that the genus-3 RM point-counting algorithm of [3] is practical and that we extended it to arbitrary genus with a similar complexity (at least conjecturally), one could hope to use it for practical computations in genus larger than 3. Using the RM families that we presented in Section 2.1, the smallest case to consider would be $g = 5$ and $\ell = 23$. Given that the case $\ell = 29$ is still a challenge in genus 3 (using Magma, it currently requires more than 1.5 TB of RAM), we are quite pessimistic on the feasibility of such an attempt in a close future.

An interesting problem that could have both practical and theoretical impact in terms of complexity is to find new (families of) curves with explicit real multiplication. While RM by multiquaratics is theoretically interesting, practical experiments could be made much easier if we could find curves with RM by an order in which a small prime (say $\ell \leq 11$) happens to be totally split.

Lastly, even if we were to find a way of solving the polynomial systems within a polynomial (or at least subexponential) complexity, the number of non-generic systems is still exponential in g . Heuristically, non-genericity should never be a problem, but in order to reach a proven subexponential complexity, one also needs to find another way of dealing with non-genericity.

References

- [1] Simon Abelard. *Counting points on hyperelliptic curves in large characteristic: algorithms and complexity*. PhD thesis, Université de Lorraine, 2018.
- [2] Simon Abelard, Pierrick Gaudry, and Pierre-Jean Spaenlehauer. Improved complexity bounds for counting points on hyperelliptic curves, 2017. To appear in *Foundations of Computational Mathematics*, ArXiv preprint 1710.03448.
- [3] Simon Abelard, Pierrick Gaudry, and Pierre-Jean Spaenlehauer. Counting points on genus-3 hyperelliptic curves with explicit real multiplication, 2018. To appear in the Proceedings of the ANTS-XIII Conference (Thirteenth Algorithmic Number Theory Symposium), ArXiv preprint 1806.05834.
- [4] Leonard M. Adleman and Ming-Deh Huang. Counting points on curves and Abelian varieties over finite fields. *Journal of Symbolic Computation*, 32(3):171–189, 2001.
- [5] Jens Bauch, Daniel J. Bernstein, Henry de Valence, Tanja Lange, and Christine Van Vredendaal. Short generators without quantum computers: the case of multiquadratics. In *EUROCRYPT 2017*, volume 10210 of *LNCS*, pages 27–59. Springer, 2017.

- [6] Antonio Cafure and Guillermo Matera. Fast computation of a rational point of a variety over a finite field. *Mathematics of Computation*, 75(256):2049–2085, 2006.
- [7] David G. Cantor. On the analogue of the division polynomials for hyperelliptic curves. *Journal für die reine und angewandte Mathematik*, 447:91–146, 1994.
- [8] Edgar Costa, David Harvey, and Kiran S. Kedlaya. Zeta functions of nondegenerate hypersurfaces in toric varieties via controlled reduction in p -adic cohomology. *To appear in the Proceedings of the ANTS-XIII Conference (Thirteenth Algorithmic Number Theory Symposium), ArXiv preprint 1806.00368*, 2018.
- [9] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT 2016*, volume 9666 of *LNCS*, pages 559–585. Springer, 2016.
- [10] Matsusaburô Fujiwara. Über die obere schranke des absoluten betrages der wurzeln einer algebraischen gleichung. *Tohoku Mathematical Journal, First Series*, 10:167–171, 1916.
- [11] Pierrick Gaudry and Robert Harley. Counting points on hyperelliptic curves over finite fields. In *ANTS 2000*, volume 1838 of *LNCS*, pages 313–332. Springer, 2000.
- [12] Pierrick Gaudry, David R. Kohel, and Benjamin A. Smith. Counting points on genus 2 curves with real multiplication. In *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 504–519. Springer, 2011.
- [13] Pierrick Gaudry and Éric Schost. Construction of secure random curves of genus 2 over prime fields. In *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 239–256. Springer, 2004.
- [14] Pierrick Gaudry and Éric Schost. Genus 2 point counting over prime fields. *Journal of Symbolic Computation*, 47(4):368–400, 2012.
- [15] Marc Giusti, Grégoire Lecerf, and Bruno Salvy. A Gröbner free alternative for polynomial system solving. *Journal of complexity*, 17(1):154–211, 2001.
- [16] David Harvey. Counting points on hyperelliptic curves in average polynomial time. *Annals of Mathematics*, 179(2):783–803, 2014.
- [17] Ming-Deh Huang and Doug Ierardi. Counting points on curves over finite fields. *Journal of Symbolic Computation*, 25(1):1–21, 1998.
- [18] Kiran S. Kedlaya. Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology. *Journal of the Ramanujan Mathematical Society*, 16(4):323–338, 2001.
- [19] David R. Kohel and Benjamin A. Smith. Efficiently computable endomorphisms for hyperelliptic curves. In *ANTS VII*, volume 4076 of *LNCS*, pages 495–509. Springer Verlag, 2006.
- [20] Alan G. B. Lauder. Deformation theory and the computation of zeta functions. *Proceedings of the London Mathematical Society*, 88(3):565–602, 2004.

- [21] Alan G. B. Lauder and Daqing Wan. Counting points on varieties over finite fields of small characteristic. In J. P. Buhler and P. Stevenhagen, editors, *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*, Mathematical Sciences Research Institute Publications, pages 579–612. Cambridge University Press, 2008.
- [22] Jonathan Pila. Frobenius maps of Abelian varieties and finding roots of unity in finite fields. *Mathematics of Computation*, 55(192):745–763, 1990.
- [23] Mohab Safey El Din and Éric Schost. A nearly optimal algorithm for deciding connectivity queries in smooth and bounded real algebraic sets. *Journal of the ACM*, 63(6):1–48, 2017.
- [24] Takakazu Satoh. The canonical lift of an ordinary elliptic curve over a finite field and its point counting. *Journal of the Ramanujan Mathematical Society*, 15(4):247–270, 2000.
- [25] Claus-Peter Schnorr and Martin Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66(1-3):181–199, 1994.
- [26] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p . *Mathematics of Computation*, 44(170):483–494, 1985.
- [27] René Schoof. Counting points on elliptic curves over finite fields. *J. Théor. Nombres Bordeaux*, 7(1):219–254, 1995.
- [28] Walter Tautz, Jaap Top, and Alain Verberkmoes. Explicit hyperelliptic curves with real multiplication and permutation polynomials. *Canad. J. Math*, 43(5):1055–1064, 1991.
- [29] Jan Tuitman. Counting points on curves using a map to \mathbb{P}^1 , II. *Finite Fields and Their Applications*, 45:301–322, 2017.