Counting Points on Genus 2 Curves with Real Multiplication

Pierrick Gaudry, David Kohel, Benjamin Smith

Counting Points on Genus 2 Curves with Real Multiplication

P. Gaudry, D. Kohel, and B. Smith

Abstract. We present an accelerated Schoof-type point-counting algorithm for curves of genus 2 equipped with an efficiently computable real multiplication endomorphism. Our new algorithm reduces the complexity of genus 2 point counting over a finite field $\mathbb{F}_q$ of large characteristic from $\widetilde{O}(\log^8 q)$ to $\widetilde{O}(\log^5 q)$. Using our algorithm we compute a 256-bit prime-order Jacobian, suitable for cryptographic applications, and also the order of a 1024-bit Jacobian.

1 Introduction

Cryptosystems based on curves of genus 2 offer per-bit security and efficiency comparable with elliptic curve cryptosystems. However, many of the computational problems related to creating secure instances of genus 2 cryptosystems are considerably more difficult than their elliptic curve analogues. Point counting—or, from a cryptographic point of view, computing the cardinality of a cryptographic group—offers a good example of this disparity, at least for curves defined over large prime fields. Indeed, while computing the order of a cryptographic-sized elliptic curve with the Schoof–Elkies–Atkin algorithm is now routine, computing the order of a comparable genus 2 Jacobian requires a significant computational effort [7, 9].

In this article we describe a number of improvements to the classical Schoof–Pila algorithm for genus 2 curves with explicit and efficient real multiplication (RM). For explicit RM curves over $\mathbb{F}_p$, we reduce the complexity of Schoof–Pila from $\widetilde{O}(\log^8 p)$ to $\widetilde{O}(\log^5 p)$. We applied a first implementation of our algorithms to find prime-order Jacobians over 128-bit fields (comparable to prime-order elliptic curves over 256-bit fields, and therefore suitable for contemporary cryptographic applications). Going further, we were able to compute the order of an RM Jacobian defined over a 512-bit prime field, far beyond the cryptographic range. (For comparison, the previous record computation in genus 2 was over a 128-bit field.)

While these RM curves are special, they are not "too special": Every ordinary genus 2 Jacobian over a finite field has RM; our special requirement is that this RM be known in advance and be efficiently computable. The moduli of curves with RM by a fixed ring form 2-dimensional subvarieties (Humbert surfaces) in the 3-dimensional moduli space of all genus 2 curves. We can generate random curves with the specified RM by choosing random points on an explicit model of the corresponding Humbert surface [10]. In comparison with elliptic curves, for which the moduli space is one-dimensional, this still gives an additional degree of freedom in the random curve selection. To generate random curves with efficiently computable RM, we choose random curves from some known one and two-parameter families (see §4).

Curves with efficiently computable RM have an additional benefit in cryptography: the efficient endomorphism can be used to accelerate scalar multiplication on the Jacobian, yielding faster encryption and decryption [11, 15, 18]. The RM formulæ are also compatible with fast arithmetic based on theta functions [6].

2 Conventional Point Counting for Genus 2 Curves

Let $\mathcal{C}$ be a curve of genus 2 over a finite field $\mathbb{F}_q$, of odd characteristic, defined by an affine model $y^2 = f(x)$, where f is a squarefree polynomial of degree 5 or 6 over $\mathbb{F}_q$. Let $J_\mathcal{C}$ be the Jacobian of $\mathcal{C}$; we assume $J_\mathcal{C}$ is ordinary and absolutely simple. Points on $J_\mathcal{C}$ correspond to degree-0 divisor classes on $\mathcal{C}$; we use the Mumford representation for divisor classes together with the usual Cantor-style composition and reduction algorithms for divisor class arithmetic [5, 2]. Multiplication by ℓ on $J_\mathcal{C}$ is denoted by $[\ell]$, and its kernel by $J_\mathcal{C}[\ell]$. More generally, if ϕ is an endomorphism of $J_\mathcal{C}$ then $J_\mathcal{C}[\phi] = \ker(\phi)$, and if S is a set of endomorphisms then $J_\mathcal{C}[S]$ denotes the intersection of $\ker(\phi)$ for ϕ in S.

2.1 The Characteristic Polynomial of Frobenius

We let π denote the Frobenius endomorphism of $J_\mathcal{C}$, with Rosati dual $\pi^\dagger$ (so $\pi\pi^\dagger = [q]$). The characteristic polynomial of π has the form

$$\chi(T) = T^4 - s_1 T^3 + (s_2 + 2q)\, T^2 - q s_1 T + q^2, \tag{1}$$

where s_1 and s_2 are integers, and s_2 is a translation of the standard definition. The polynomial $\chi(T)$ determines the cardinality of $J_\mathcal{C}(\mathbb{F}_{q^k})$ for all k: in particular, $\#J_\mathcal{C}(\mathbb{F}_q) = \chi(1)$. We refer to the determination of $\chi(T)$ as the *point counting problem*.

The polynomial $\chi(T)$ is a *Weil polynomial*: all of its complex roots lie on the circle $|z| = \sqrt{q}$. This implies the Weil bounds

$$|s_1| \le 4\sqrt{q} \quad \text{and} \quad |s_2| \le 4q. \tag{2}$$

However, the possible values of (s_1, s_2) do not fill the whole rectangle specified by the Weil bounds. Rück [17, Theorem 1.1] shows that in fact s_1 and s_2 satisfy

$$s_1^2 - 4s_2 \ge 0 \text{ and } s_2 + 4q \ge 2|s_1|,$$

so the possible values of (s_1, s_2) are in the following domain:

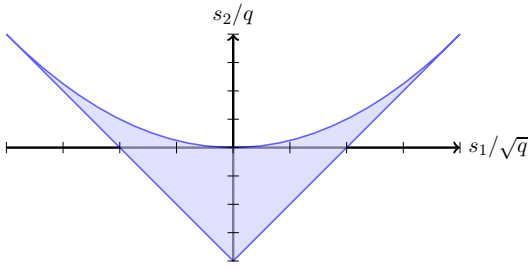

2.2 The Classical Schoof–Pila Algorithm for Genus 2 Curves

The objective of point counting is to compute $\chi(T)$, or equivalently the tuple of integers (s_1, s_2). When the characteristic of $\mathbb{F}_q$ is large, the conventional approach is to apply the Schoof–Pila algorithm as far as is practical, before passing to a baby-step giant-step algorithm if necessary (see §2.5). The strategy of Schoof's algorithm and its generalizations is to compute the polynomials $\chi_\ell(T) = \chi(T) \bmod (\ell)$ for sufficiently many primes (or prime powers) ℓ to reconstruct $\chi(T)$ using the Chinese Remainder Theorem.

Since $\chi_\ell(T)$ is the characteristic polynomial of π restricted to $J_{\mathcal{C}}[\ell]$ (see [16, Proposition 2.1]), we have

$$\chi_\ell(\pi)(D) = 0 \text{ for all } D \text{ in } J_{\mathcal{C}}[\ell].$$

Conversely, to compute $\chi_\ell(T)$ we let D be a generic element of $J_{\mathcal{C}}[\ell]$ (as in §2.3 below), compute the three points

$$(\pi^2 + [\bar{q}])^2(D), \quad (\pi^2 + [\bar{q}])\pi(D), \text{ and } \quad \pi^2(D),$$

and then search for the coefficients $(\bar{s}_1, \bar{s}_2)$ of $\chi_\ell(T)$ in $(\mathbb{Z}/\ell\mathbb{Z})^2$, for which the linear relation

$$(\pi^2 + [\bar{q}])^2(D) - [\bar{s}_1]\,(\pi^2 + [\bar{q}])\pi(D) + [\bar{s}_2]\,\pi^2(D) = 0 \tag{3}$$

holds. If the minimal polynomial of π on $J_{\mathcal{C}}[\ell]$ is a proper divisor of $\chi_\ell(T)$—which occurs for at most a finite number of ℓ dividing $\mathrm{disc}(\chi)$—then the polynomial so determined is not unique, but $\chi_\ell(T)$ can be determined by deducing the correct multiplicities of its factors.

Once we have computed $\chi_\ell(T)$ for sufficiently many ℓ, we reconstruct $\chi(T)$ using the Chinese Remainder Theorem. The Weil and Rück bounds together with a weak version of the prime number theorem tell us how many ℓ are required: Pila notes in [16, §1] that the set of $O(\log q)$ primes $\ell < 21 \log q$ will suffice. We analyse the complexity of the classical Schoof–Pila algorithm in §2.4.

2.3 Endomorphisms and Generic Kernel Elements

We now recall how to contruct an effective version of a generic ℓ-torsion element. We present it in a slightly more general setting, so that we can use this ingredient in the subsequent RM-specific algorithm. Therefore, we show how to compute with a generic element of the kernel of some endomorphism ϕ of $J_{\mathcal{C}}$, whereas ϕ is just $[\ell]$ in the classical algorithm.

Definition 1. *Fix an embedding $P \mapsto D_P$ of $\mathcal{C}$ in $J_{\mathcal{C}}$. We say that an endomorphism ϕ of $J_{\mathcal{C}}$ is* explicit *if we can effectively compute polynomials d_0, d_1, d_2, e_0, e_1, and e_2 such that if $P = (x_P, y_P)$ is a generic point of $\mathcal{C}$, then the Mumford representation of $\phi(D_P)$ is given by*

$$\phi(D_P) = \left(x^2 + \frac{d_1(x_P)}{d_2(x_P)}x + \frac{d_0(x_P)}{d_2(x_P)}, y - y_P\left(\frac{e_1(x_P)}{e_2(x_P)}x + \frac{e_0(x_P)}{e_2(x_P)}\right)\right). \tag{4}$$

The d_0, d_1, d_2, e_0, e_1, and e_2 are called the ϕ-division polynomials.

If ϕ is an explicit endomorphism, then we can use (4) (extending $\mathbb{Z}$-linearly) to evaluate $\phi(D)$ for general divisor classes D in $J_{\mathcal{C}}$. In the case $\phi = [\ell]$, the $[\ell]$-division polynomials are the ℓ-division polynomials of Cantor [3]. The ϕ-division polynomials depend on the choice of embedding $P \mapsto D_P$; we will make this choice explicit when computing the ϕ-division polynomials for each of our families in §4.

To compute a generic element of $J_{\mathcal{C}}[\phi]$, we generalize the approach of [7] (which computes generic elements of $J_{\mathcal{C}}[\ell]$). The resulting algorithm is essentially the same as in [7, §3] (except for the parasite computation step, which we omit) with ϕ-division polynomials replacing ℓ-division polynomials, so we will only briefly sketch it here.

Let $D = (x^2 + a_1x + a_0, y - (b_1x + b_0))$ be (the Mumford representation of) a generic point of $J_{\mathcal{C}}$. We want to compute a triangular ideal I_ϕ in $\mathbb{F}_q[a_1, a_0, b_1, b_0]$ vanishing on the nonzero elements of $J_{\mathcal{C}}[\phi]$. The element D equals $D_{(x_1,y_1)} + D_{(x_2,y_2)}$, where (x_1, y_1) and (x_2, y_2)

are generic points of $\mathcal{C}$. To find a triangular system of relations on the a_i and b_i such that D is in $J_{\mathcal{C}}[\phi]$ we solve for x_1, y_1, x_2, and y_2 in

$$\phi(D_{(x_1,y_1)}) = -\phi(D_{(x_2,y_2)}),$$

applying (4) and using resultants computed with the evaluation–interpolation technique of [7, §3.1]. We then resymmetrize as in [7, §3.2] to express the result in terms of the a_i and b_i. We can now compute with a "generic" element $(x^2 + a_1x + a_0, y - (b_1x + b_0))$ of $J_{\mathcal{C}}[\phi]$ by reducing the coefficients modulo I_ϕ after each operation.

Following the complexity analysis of [7, §3.5], we can compute a triangular representation for I_ϕ in $O(\delta^2\mathsf{M}(\delta)\log\delta + \mathsf{M}(\delta^2)\log\delta)$ field operations, where δ is the maximum among the degrees of the ϕ-division polynomials, and $\mathsf{M}(d)$ is the number of operations required to multiply polynomials of degree d over $\mathbb{F}_q$. Using asymptotically fast multiplication algorithms, we can therefore compute I_ϕ in $\widetilde{O}(\delta^3)$ field operations. The degree of I_ϕ is in $O(\delta^2)$; with this triangular representation, each multiplication modulo I_ϕ costs $\widetilde{O}(\delta^2)$ field operations.

2.4 Complexity of Classical Schoof–Pila Point Counting

Proposition 1. *The complexity of the classical Schoof–Pila algorithm for a curve of genus* 2 *over* $\mathbb{F}_q$ *is in* $\widetilde{O}((\log q)^8)$.

Proof. To determine $\chi(T)$, we need to compute $\chi_\ell(T)$ for $O(\log q)$ primes ℓ in $O(\log q)$. To compute $\chi_\ell(T)$, we must first compute the ℓ-division polynomials, which have degrees in $O(\ell^2)$. We then compute the kernel ideal I_ℓ; according to the previous subsection, the total cost is in $\widetilde{O}(\ell^6)$ field operations. The cost of checking (3) against a generic element of $J_{\mathcal{C}}[\ell]$ decomposes into the cost of computing Frobenius images of the generic element in $\widetilde{O}(\ell^4 \log q)$ and of finding the matching pair $(\bar{s}_1, \bar{s}_2)$ in $\widetilde{O}(\ell^5)$ field operations. So the total complexity for computing $\chi_\ell(T)$ is in $\widetilde{O}(\ell^4(\ell^2 + \log q))$ field operations. In terms of bit operations, for each ℓ bounded by $O(\log q)$, we compute $\chi_\ell(T)$ in time $\widetilde{O}((\log q)^7)$, and the result follows from the addition of these costs for all the different ℓ's. □

2.5 Baby-Step Giant-Step Algorithms

In practice, computing $\chi_\ell(T)$ with classical Schoof–Pila becomes impractical for large values of ℓ. The usual approach is to carry out the Schoof–Pila algorithm to the extent possible, obtaining congruences for s_1 and s_2 modulo some integer M, before completing the calculation using a generic group algorithm such as baby-step giant-step (BSGS). Our BSGS algorithm of choice is the low-memory parallelized variant of the Matsuo–Chao–Tsuji algorithm [8, 12].

The bounds in (2) imply that the search space of candidates for (s_1, s_2) is in $O(q^{3/2})$, and a pure BSGS approach finds (s_1, s_2) in time and space $\widetilde{O}(q^{3/4})$. However, when we apply BSGS after a partial Schoof–Pila computation, we obtain a congruence for (s_1, s_2) modulo M. If $M < 8q$, then the size of the search space is reduced to $O(q^{3/2}/M^2)$, and the complexity for finding (s_1, s_2) is reduced to $\widetilde{O}(q^{3/4}/M)$. For larger M, the value of s_1 is fully determined, and the problem is reduced to a one-dimensional search space of size $O(q/M)$ for which the complexity becomes $\widetilde{O}(\sqrt{q/M})$.

3 Point Counting in Genus 2 with Real Multiplication

By assumption, $J_{\mathcal{C}}$ is ordinary and simple, so $\chi(T)$ is an irreducible polynomial defining a quartic CM-field with real quadratic subfield $\mathbb{Q}(\sqrt{\Delta})$. We say that $J_{\mathcal{C}}$ (and $\mathcal{C}$) has *real multiplication* (RM) by $\mathbb{Q}(\sqrt{\Delta})$. For a randomly selected curve, Δ is in $O(q)$; but in the sequel we consider families of curves with RM by $\mathbb{Q}(\sqrt{\Delta})$ for small Δ (= 5 or 8), admitting an explicit (in the sense of Definition 1) endomorphism ϕ such that

$$\mathbb{Z}[\phi] = \mathbb{Q}(\sqrt{\Delta}) \cap \mathrm{End}(J_{\mathcal{C}}) \tag{5}$$

(that is, $\mathbb{Z}[\phi]$ is the full real subring of $\mathrm{End}(J_{\mathcal{C}})$), and

$$\mathrm{disc}\,(\mathbb{Z}[\phi]) = \Delta.$$

We presume that the trace $\mathrm{Tr}(\phi)$ and norm $\mathrm{N}(\phi)$, such that $\phi^2 - \mathrm{Tr}(\phi)\phi + \mathrm{N}(\phi) = 0$, are known. We also suppose that ϕ is *efficient*, in the following sense:

Definition 2. *We say that an explicit endomorphism ϕ is* efficiently computable *if the cost of evaluating ϕ at points of $J_{\mathcal{C}}(\mathbb{F}_q)$ requires only $O(1)$ field operations (comparable to a few group operations in $J_{\mathcal{C}}$). In practice, this means that the ϕ-division polynomials have small degree.*

The existence of an efficiently computable ϕ and knowledge of Δ allows us to make significant improvements to each stage of the Schoof–Pila algorithm. Briefly: in §3.2 we use ϕ to simplify the testing procedure for each ℓ; in §3.3 we show that when ℓ splits in $\mathbb{Z}[\phi]$, we can use ϕ to obtain a radical reduction in complexity for computing $\chi_\ell(T)$; and in §3.4 we show that knowing an effective ϕ allows us to use many fewer primes ℓ.

3.1 The RM Characteristic Polynomial

Let $\psi = \pi + \pi^\dagger$; we consider $\mathbb{Z}[\psi]$, a subring of the real quadratic subring of $\mathrm{End}(J_{\mathcal{C}})$. The characteristic polynomial of ψ is the *real Weil polynomial*

$$\xi(T) = T^2 - s_1 T + s_2; \tag{6}$$

the discriminant of $\mathbb{Z}[\psi]$ is $\Delta_0 = s_1^2 - 4s_2$. The analogue for (s_1, Δ_0) of Rück's bounds is

$$(|s_1| - 4\sqrt{q})^2 \geq \Delta_0 = s_1^2 - 4s_2 \geq 0. \tag{7}$$

Equation (5) implies that $\mathbb{Z}[\psi]$ is contained in $\mathbb{Z}[\phi]$, so there exist integers m and n such that

$$\psi = m + n\phi. \tag{8}$$

Both s_1 and s_2 are determined by m and n: we have

$$s_1 = \mathrm{Tr}(\psi) = 2m + n\mathrm{Tr}(\phi) \quad \text{and} \quad s_2 = \mathrm{N}(\psi) = (s_1^2 - n^2\Delta)/4. \tag{9}$$

In fact n is the conductor of $\mathbb{Z}[\psi]$ in $\mathbb{Z}[\phi]$ up to sign: $|n| = [\mathbb{Z}[\phi] : \mathbb{Z}[\psi]]$, and hence

$$\Delta_0 = \mathrm{disc}(\mathbb{Z}[\psi]) = s_1^2 - 4s_2 = n^2\Delta.$$

The square root of the bounds in (7) gives bounds on s_1 and n:

$$4\sqrt{q} - |s_1| \geq \sqrt{\Delta_0} = |n|\sqrt{\Delta} \geq 0;$$

In particular, $|s_1| \leq 4\sqrt{q}$ and $|n| \leq 4\sqrt{q/\Delta}$. Applying the relation in (9), we have the bounds

$$|m| \leq 2(|\mathrm{Tr}(\phi)| + \sqrt{\Delta})\sqrt{q/\Delta} \quad \text{and} \quad |n| \leq 4\sqrt{q/\Delta}. \tag{10}$$

Both $|m|$ and $|n|$ are in $O(\sqrt{q})$.

3.2 An Efficiently Computable RM Relation

We can use our efficiently computable endomorphism ϕ to replace the relation of (3) with a more efficiently computable alternative. Multiplying (8) through by π, we have

$$\psi\pi = \pi^2 + [q] = m\pi + n\phi\pi.$$

We can therefore compute $\bar{m} = m \bmod \ell$ and $\bar{n} = n \bmod \ell$ by letting D be a generic ℓ-torsion point, computing the three points

$$(\pi^2 + [\bar{q}])(D), \quad \pi(D), \quad \text{and} \quad \phi\pi(D),$$

and then searching for the $\bar{m}$ and $\bar{n}$ in $\mathbb{Z}/\ell\mathbb{Z}$ such that

$$(\pi^2 + [\bar{q}])(D) - [\bar{m}]\pi(D) - [\bar{n}]\phi\pi(D) = 0 \tag{11}$$

holds; we can find such an $\bar{m}$ and $\bar{n}$ in $O(\ell)$ group operations.

Solving (11) rather than (3) has several advantages. First, computing $(\pi^2 + [\bar{q}])(D)$, $\pi(D)$, and $\phi\pi(D)$ requires only two applications of Frobenius, instead of the four required to compute $(\pi^2 + [\bar{q}])^2(D)$, $(\pi^2 + [\bar{q}])\pi(D)$, and $\pi^2(D)$ (and Frobenius applications are costly in practice). Moreover, either s_2 needs to be determined in $O(q)$, or else the value of n in (3) leaves a sign ambiguity for each prime ℓ, because only $n^2 \bmod \ell$ can be deduced from $(\bar{s}_1, \bar{s}_2)$. In contrast, (11) determines n directly.

3.3 Exploiting Split Primes in $\mathbb{Q}(\sqrt{\Delta})$

Let $\mathbb{Z}[\phi] \subset \mathrm{End}(J_{\mathcal{C}})$ be an RM order in $\mathbb{Q}(\phi) \cong \mathbb{Q}(\sqrt{\Delta})$. Asymptotically, half of all primes ℓ split: $(\ell) = \mathfrak{p}_1\mathfrak{p}_2$ in $\mathbb{Z}[\phi]$, where $\mathfrak{p}_1 + \mathfrak{p}_2 = (1)$ (and this carries over to prime powers ℓ). This factorization gives a decomposition of the ℓ-torsion

$$J_{\mathcal{C}}[\ell] = J_{\mathcal{C}}[\mathfrak{p}_1] \oplus J_{\mathcal{C}}[\mathfrak{p}_2].$$

In particular, any ℓ-torsion point D can be uniquely expressed as a sum $D = D_1 + D_2$ where D_i is in $J_{\mathcal{C}}[\mathfrak{p}_i]$.

According to the Cohen–Lenstra heuristics [4], more than 75% of RM fields have class number 1; in each of the explicit RM families in §4, the order $\mathbb{Z}[\phi]$ has class number 1. All ideals are principal in such an order, so we may find a generator for each of the ideals $\mathfrak{p}_i$. Furthermore, the following lemma shows that we can find a generator which is not too large.

Lemma 1. *If $\mathfrak{p}$ is a principal ideal of norm ℓ in a real quadratic order $\mathbb{Z}[\phi]$, then there exists an effectively computable generator of $\mathfrak{p}$ with coefficients in $O(\sqrt{\ell})$.*

Proof. Let α be a generator of $\mathfrak{p}$, and ε a fundamental unit of $\mathbb{Z}[\phi]$. Let $\gamma \mapsto \gamma_1$ and $\gamma \mapsto \gamma_2$ be the two embeddings of $\mathbb{Z}[\phi]$ in $\mathbb{R}$, indexed so that $|\alpha_1| \geq |\alpha_2|$ and $|\varepsilon_1| > 1$ (replacing ε with ε^{-1} if necessary). Then $R = \log(|\varepsilon_1|)$ is the regulator of $\mathbb{Z}[\phi]$. Set $\beta = \varepsilon^{-k}\alpha$, where $k = [\log(|\alpha_1/\sqrt{\ell}|)/R]$; then $\beta = a + b\phi$ is a new generator for $\mathfrak{p}$ such that

$$-\frac{1}{2} \leq \frac{\log(|\beta_i/\sqrt{\ell}|)}{R} \leq \frac{1}{2}.$$

From the preceding bounds, $|\beta_1 + \beta_2| = |2a + b\mathrm{Tr}(\phi)|$ and $|\beta_1 - \beta_2| = |b\sqrt{\Delta}|$ are bounded by $2e^{R/2}\sqrt{\ell}$. Since $\mathrm{Tr}(\phi)$, Δ and R are fixed constants, $|a|$ and $|b|$ are in $O(\sqrt{\ell})$. The "effective" part of the result follows from classical algorithms for quadratic fields. □

Lemma 2. *Let $J_{\mathcal{C}}$ be the Jacobian of a genus 2 curve over a finite field $\mathbb{F}_q$ with an efficiently computable RM endomorphism ϕ. There exists an algorithm which, given a principal ideal $\mathfrak{p}$ of norm ℓ in $\mathbb{Z}[\phi]$, computes an explicit generator α of $\mathfrak{p}$ and the α-division polynomials in $O(\ell)$ field operations.*

Proof. By Lemma 1, we can compute a generator $\alpha = [a] + [b]\phi$ with a and b in $O(\sqrt{\ell})$. The $[a]$- and $[b]$-division polynomials have degrees in $O(\ell)$, and can be determined in $O(\ell)$ field operations. The division polynomials for the sum $\alpha = [a] + [b]\phi$ require one sum and one application of ϕ; and since ϕ is efficiently computable, this increases the division polynomial degrees and computing time by at most a constant factor. □

We can now state the main theorem for RM point counting.

Theorem 1. *There exists an algorithm for the point counting problem in a family of genus 2 curves with efficiently computable RM of class number 1, whose complexity is in $\widetilde{O}((\log q)^5)$.*

Proof. Let $J_{\mathcal{C}}$ be a Jacobian in a family with efficiently computable RM by $\mathbb{Z}[\phi]$. Suppose that ℓ is prime, $(\ell) = \mathfrak{p}_1\mathfrak{p}_2$ in $\mathbb{Z}[\phi]$, and that the $\mathfrak{p}_i$ are principal. By Lemma 2 we can compute representative α-division polynomials for $J_{\mathcal{C}}$ for each $\mathfrak{p}$ in $\{\mathfrak{p}_1, \mathfrak{p}_2\}$ in time $\widetilde{O}(\ell)$, hence generic points D_i in $J_{\mathcal{C}}[\mathfrak{p}_i]$.

We recall that (11) is the homomorphic image under π of the equation

$$\psi(D) - [\bar{m}](D) - [\bar{n}]\phi(D) = 0.$$

When applied to D_i in $J_{\mathcal{C}}[\mathfrak{p}_i]$, both ψ and ϕ act as elements of $\mathbb{Z}[\phi]/\mathfrak{p}_i \cong \mathbb{Z}/\ell\mathbb{Z}$. Moreover $\bar{x}_i = \phi \bmod \mathfrak{p}_i$ is known, and it remains to determine $\bar{y}_i = \psi \bmod \mathfrak{p}_i$ by means of the discrete logarithm

$$\psi(D_i) = [\bar{y}_i](D_i) = [\bar{m} + \bar{n}\bar{x}_i](D_i)$$

in the cyclic group $\langle D_i \rangle \cong \mathbb{Z}/\ell\mathbb{Z}$. The application of π transports this discrete logarithm problem to that of solving for $\bar{y}_i$ in

$$D_i'' = [\bar{y}_i]D',$$

where $D_i' = \pi(D_i)$ and $D_i'' = (\pi^2 + [\bar{q}])(D_i)$. By the CRT, from $(\bar{y}_1, \bar{y}_2)$ in $(\mathbb{Z}/\ell\mathbb{Z})^2$ we recover $\bar{y}$ in $\mathbb{Z}[\phi]/(\ell)$, from which we solve for $(\bar{m}, \bar{n})$ in $(\mathbb{Z}/\ell\mathbb{Z})^2$ such that

$$\bar{y} = \bar{m} + \bar{n}\phi \in \mathbb{Z}[\phi]/(\ell).$$

The values of $(\bar{s}_1, \bar{s}_2)$ are then recovered from (9).

The ring $\mathbb{Z}[\phi]$ is fixed, so as $\log q$ goes to infinity we find that 50% of all primes ℓ split in $\mathbb{Z}[\phi]$ by the Chebotarev density theorem. It therefore suffices to consider split primes in $O(\log q)$. In comparison with the conventional algorithm presented in §2.2, we reduce from computation modulo the ideal for $J_{\mathcal{C}}[\ell]$ of degree in $O(\ell^4)$, to computation modulo the ideals for $J_{\mathcal{C}}[\mathfrak{p}_i]$ of degree in $O(\ell^2)$. This means a reduction from $\widetilde{O}(\ell^4(\ell^2 + \log q))$ to $\widetilde{O}(\ell^2(\ell + \log q))$ field operations for the determination of each $\chi_\ell(T)$, giving the stated reduction in total complexity from $\widetilde{O}((\log q)^8)$ to $\widetilde{O}((\log q)^5)$. □

Remark 1. Computing (m, n) instead of (s_1, s_2) allows us to reduce the number of primes ℓ to be considered by about a half, since by (10) their product needs to be in $O(\sqrt{q})$ instead of $O(q)$. While this changes only the constant in the asymptotic complexity of the algorithm, it yields a significant improvement in practice.

Remark 2. If $\mathbb{Z}[\phi]$ does not have class number 1, and if $(\ell) = \mathfrak{p}_1\mathfrak{p}_2$ where the $\mathfrak{p}_i$ are not principal, then we may use a small complementary ideal $(c) = \mathfrak{c}_1\mathfrak{c}_2$ such that $\mathfrak{c}_i\mathfrak{p}_i$ are principal in order to apply Lemma 2 to a larger proportion of small ideals. Moreover, if $(\bar{m}, \bar{n})$ is known modulo c, this can be used to reduce the discrete log problem modulo ℓ. Again, since a fixed positive density $1/2h$ of primes are both split and principal, where h is the class number of $\mathbb{Z}[\phi]$, this does not affect the asymptotic complexity. Moreover, the first occurrence of a nontrivial class group is for $\Delta = 65$, beyond the current range for which an explicit RM construction is currently known.

3.4 Shrinking the BSGS Search Space

In the context of the conventional Schoof-Pila algorithm, we need to find s_1 in $O(\sqrt{q})$ and s_2 in $O(q)$. However, (8), and the effective form of (11) (valid for all points D of $J_{\mathcal{C}}$), replaces the determination of (s_1, s_2) with the tuple (m, n) of integers in $O(\sqrt{q})$. As a result, the search space is reduced from $O(q^{3/2})$ to $O(q)$. Thus the BSGS strategy can find (m, n) (which determines (s_1, s_2)) in time and space $O(\sqrt{q})$, compared with $O(q^{3/4})$ when searching directly for (s_1, s_2).

As in the general case, if one knows (m, n) modulo an integer M, then the area of the search rectangle is reduced by a factor of M^2, so we find the tuple (m, n) in $O(\sqrt{q}/M)$ group operations. Contrary to the general case of §2.5, since m and n have the same order of magnitude, the speed-up is always by a factor of M.

4 Examples of Families of Curves with Explicit RM

We now exhibit some families of curves and efficient RM endomorphisms that can be used as sources of inputs to our algorithm.

4.1 Correspondences and Endomorphisms

To give a concrete representation for endomorphisms of $J_{\mathcal{C}}$, we use *correspondences*: that is, divisors on the surface $\mathcal{C} \times \mathcal{C}$. Suppose that $\mathcal{R}$ is a curve on $\mathcal{C} \times \mathcal{C}$, and let $\pi_1 : \mathcal{R} \to \mathcal{C}$ and $\pi_2 : \mathcal{R} \to \mathcal{C}$ be the restrictions to $\mathcal{R}$ of the natural projections from $\mathcal{C} \times \mathcal{C}$ onto its first and second factors. We have a pullback homomorphism $(\pi_1)^* : \mathrm{Pic}(\mathcal{C}) \to \mathrm{Pic}(\mathcal{R})$, defined by

$$(\pi_1)^*\Big(\Big[\sum_{P \in \mathcal{C}(\overline{\mathbb{F}}_q)} n_P P\Big]\Big) = \Big[\sum_{P \in \mathcal{C}(\overline{\mathbb{F}}_q)} n_P \sum_{Q \in \pi_1^{-1}(P)} Q\Big],$$

where the preimages Q are counted with the appropriate multiplicities. (A standard moving lemma shows that we can always choose divisor class representatives so that each $\pi^{-1}(P)$ is zero-dimensional.) We also have a pushforward homomorphism $(\pi_2)_* : \mathrm{Pic}(\mathcal{R}) \to \mathrm{Pic}(\mathcal{C})$, defined by

$$(\pi_2)_*\Big(\Big[\sum_{Q \in \mathcal{R}(\overline{\mathbb{F}}_q)} n_Q Q\Big]\Big) = \Big[\sum_{Q \in \mathcal{R}(\overline{\mathbb{F}}_q)} n_Q \pi_2(Q)\Big].$$

Note that $(\pi_1)^*$ maps $\mathrm{Pic}^n(\mathcal{C})$ into $\mathrm{Pic}^{(n \deg \pi_1)}(\mathcal{R})$ and $(\pi_2)_*$ maps $\mathrm{Pic}^n(\mathcal{R})$ into $\mathrm{Pic}^n(\mathcal{C})$ for all n. Hence $(\pi_2)_* \circ (\pi_1)^*$ maps $\mathrm{Pic}^0(\mathcal{C})$ into $\mathrm{Pic}^0(\mathcal{C})$, so we have an *induced endomorphism*

$$\phi = (\pi_2)_* \circ (\pi_1)^* : J_{\mathcal{C}} \to J_{\mathcal{C}}.$$

We write x_1, y_1 and x_2, y_2 for the coordinates on the first and second factors of $\mathcal{C} \times \mathcal{C}$, respectively (so $\pi_i(x_1, y_1, x_2, y_2) = (x_i, y_i)$). In our examples, the correspondence $\mathcal{R}$ will be defined by two equations:

$$\mathcal{R} = V(A(x_1, x_2), B(x_1, y_1, x_2, y_2)) .$$

On the level of divisors, the image of a generic point $P = (x_P, y_P)$ of $\mathcal{C}$ (that is, a generic prime divisor) under the endomorphism ϕ is given by

$$\phi : (x_P, y_P) \longmapsto V(A(x_P, x), B(x_P, y_P, x, y)) .$$

Using the relations $y_P^2 = f(x_P)$ and $y^2 = f(x)$ (and the fact that correspondences cut out by principal ideals induce the zero homomorphism), we can easily replace A and B with Cantor-reducible generators to derive the Mumford representation of $\phi(P)$, and thus the ϕ-division polynomials.

4.2 A 1-dimensional Family with RM by $\mathbb{Z}[(1+\sqrt{5})/2]$

Let t be a free parameter, and suppose that q is not a power of 5. Let $\mathcal{C}_{\mathrm{T}}$ be the family of curves of genus 2 over $\mathbb{F}_q$ considered by Tautz, Top, and Verberkmoes in [19, Example 3.5], defined by

$$\mathcal{C}_{\mathrm{T}} : y^2 = x^5 - 5x^3 + 5x + t.$$

Let $\tau_5 = \zeta_5 + \zeta_5^{-1}$, where ζ_5 is a 5th root of unity in $\overline{\mathbb{F}}_q$. Let ϕ_{T} be the endomorphism induced by the (constant) family of correspondences

$$\mathcal{R}_{\mathrm{T}} = V\big(x_1^2 + x_2^2 - \tau_5 x_1 x_2 + \tau_5^2 - 4, y_1 - y_2\big) \subset \mathcal{C}_{\mathrm{T}} \times \mathcal{C}_{\mathrm{T}}.$$

(Note that $\mathcal{R}_{\mathrm{T}}$ and ϕ_{T} are defined over $\mathbb{F}_q(\tau_5)$, which is equal to $\mathbb{F}_q$ if and only if $q \not\equiv \pm 2 \bmod 5$.) The family $\mathcal{C}_{\mathrm{T}}$ has an unique point P_∞ at infinity, which we can use to define an embedding

$$P = (x_P, y_P) \longmapsto D_P := [(P) - (P_\infty)] \leftrightarrow (x - x_P, y - y_P)$$

of $\mathcal{C}_{\mathrm{T}}$ in $J_{\mathcal{C}_{\mathrm{T}}}$. With respect to this embedding, the ϕ_{T}-division polynomials are

$$d_2 = 1, \quad d_1 = -\tau_5 x, \quad d_0 = x^2 + \tau_5^2 - 4, \quad e_2 = 1, \quad e_1 = 0, \quad e_0 = 1.$$

Proposition 2. *The minimal polynomial of ϕ_{T} is T^2+T-1: that is, ϕ_{T} acts as multiplication by $-(1+\sqrt{5})/2$ on $J_{\mathcal{C}_{\mathrm{T}}}$. A prime ℓ splits into two principal ideals in $\mathbb{Z}[\phi_{\mathrm{T}}]$ if and only if $\ell \equiv \pm 1 \bmod 5$.*

Proof. The first claim is proven in [19, §3.5]. More directly, if P and Q are generic points of $\mathcal{C}_{\mathrm{T}}$, then on the level of divisors we find

$$(\phi_{\mathrm{T}}^2 + \phi_{\mathrm{T}})((P) - (Q)) = (P) - (Q) + \operatorname{div}\left(\frac{y - y(P)}{y - y(Q)}\right).$$

Hence $\mathbb{Z}[\phi_{\mathrm{T}}]$ is isomorphic to the ring of integers of $\mathbb{Q}(\sqrt{5})$. The primes ℓ splitting in $\mathbb{Q}(\sqrt{5})$ are precisely those congruent to ± 1 modulo 5; and $\mathbb{Q}(\sqrt{5})$ has class number 1, so the primes over ℓ are principal. □

The Igusa invariants of $\mathcal{C}_{\mathrm{T}}$, viewed as a point in weighted projective space, are $(140 : 550 : 20(32t^2 - 3) : 25(896t^2 - 3109) : 64(t^2 - 4)^2)$; in particular, $\mathcal{C}_{\mathrm{T}}$ has a one-dimensional image in the moduli space of curves of genus 2. The Jacobian of the curve with the same defining equation over $\mathbb{Q}(t)$ is absolutely simple (cf. [11, Remark 15]).

4.3 A 2-dimensional Family with RM by $\mathbb{Z}[(1+\sqrt{5})/2]$

Let s and t be free parameters, and consider the family of curves $\mathcal{C}_{\mathrm{H}} : y^2 = F_{\mathrm{H}}(x)$, where

$$F_{\mathrm{H}}(x) = sx^5 - (2s+t)x^4 + (s^2+3s+2t-1)x^3 - (3s+t-3)x^2 + (s-3)x + 1.$$

This family is essentially due to Humbert; it is equal to the family of Mestre [13, §2.1] with $(U,T) = (s,t)$, and the family of Wilson [20, Proposition 3.4.1] with $(A,B) = (s, -t-3s+3)$. The family has a full 2-dimensional image in the moduli space of genus 2 curves.

Let $\mathcal{R}_{\mathrm{H}}$ be the family of correspondences on $\mathcal{C}_{\mathrm{H}} \times \mathcal{C}_{\mathrm{H}}$ defined by

$$\mathcal{R}_{\mathrm{H}} = V\big(x_1^2x_2^2 + s(s-1)x_1x_2 - s^2(x_1-x_2) + s^2, y_1 - y_2\big);$$

let ϕ_{H} be the induced endomorphism. The family $\mathcal{C}_{\mathrm{H}}$ has a unique point P_∞ at infinity, which we can use to define an embedding

$$P = (x_P, y_P) \longmapsto D_P := [(P) - (P_\infty)] \leftrightarrow (x - x_P, y - y_P)$$

of $\mathcal{C}_{\mathrm{H}}$ in $J_{\mathcal{C}_{\mathrm{H}}}$. With respect to this embedding, the ϕ_{H}-division polynomials are

$$d_2 = x^2, \quad d_1 = (s^2-s)x + s^2, \quad d_0 = -s^2x + s^2, \quad e_2 = 1, \quad e_1 = 0, \quad e_0 = 1.$$

Proposition 3. *The minimal polynomial of ϕ_{H} is T^2+T-1: that is, ϕ_{H} acts as multipliction by $-(1+\sqrt{5})/2$ on $J_{\mathcal{C}_{\mathrm{H}}}$. A prime ℓ splits into two principal ideals in $\mathbb{Z}[\phi_{\mathrm{H}}]$ if and only if $\ell \equiv \pm 1 \bmod 5$.*

Proof. The first assertion is [13, Proposition 2] with $n = 5$; the rest of the proof is exactly as in Proposition 2. □

4.4 A 2-dimensional Family with RM by $\mathbb{Z}[\sqrt{2}]$

For an example with $\Delta = 8$, we present a twisted and reparametrized version of a construction due to Mestre [14]. Let s and t be free parameters, let $v(s)$ and $n(s)$ be the rational functions

$$v = v(s) := \frac{s^2+2}{s^2-2} \quad \text{and} \quad n = n(s) := \frac{4s(s^4+4)}{(s^2-2)^3},$$

and let $\mathcal{C}_{\mathrm{M}}$ be the family of curves defined by

$$\mathcal{C}_{\mathrm{M}} : y^2 = F_{\mathrm{M}}(x) := (vx-1)(x-v)(x^4 - tx^2 + vt - 1).$$

The family of correspondences on $\mathcal{C}_{\mathrm{M}} \times \mathcal{C}_{\mathrm{M}}$ defined by

$$\mathcal{R}_{\mathrm{M}} = V\begin{pmatrix} x_1^2x_2^2 - v^2(x_1^2+x_2^2) + 1, \\ y_1y_2 - n(x_1^2+x_2^2-t)(x_1x_2 - v(x_1+x_2)+1)\end{pmatrix}$$

induces an endomorphism ϕ_{M} of $J_{\mathcal{C}_{\mathrm{M}}}$.

The family $\mathcal{C}_{\mathrm{M}}$ has two points at infinity, P_∞^+ and P_∞^-. which are generically only defined over a quadratic extension of $\mathbb{F}_q(s,t)$. Let $D_\infty = (P_\infty^+) + (P_\infty^-)$ denote the divisor at infinity. We can use the rational Weierstrass point $P_v = (v,0)$ on $\mathcal{C}_{\mathrm{M}}$ to define an embedding

$$P = (x_P, y_P) \longmapsto D_P := [(P) + (P_v) - D_\infty] \leftrightarrow \Big((x-x_P)(x-v), y - \frac{y_P}{x_P - v}(x-v)\Big)$$

of $\mathcal{C}_{\mathrm{M}}$ in $J_{\mathcal{C}_{\mathrm{M}}}$ (the appropriate composition and reduction algorithms for divisor class arithmetic on genus 2 curves with an even-degree model, such as $J_{\mathcal{C}_{\mathrm{M}}}$, appear in [5].) With respect to this embedding, the ϕ_{M}-division polynomials are

$$\begin{aligned} d_2 &= x^2 - v^2, & e_2 &= (x^2 - v^2)F_{\mathrm{M}}(x), \\ d_1 &= 0, & e_1 &= n(x - v)(x^4 - tx^2 + tv^2 - 1), \\ d_0 &= -v^2x^2 + 1, & e_0 &= n(vx - 1)(x^4 - tx^2 + tv^2 - 1). \end{aligned}$$

Proposition 4. *The minimal polynomial of ϕ_{M} is $T^2 - 2$: that is, ϕ_{M} acts as multiplication by $\sqrt{2}$ on $J_{\mathcal{C}_{\mathrm{M}}}$. A prime ℓ splits into two principal ideals in $\mathbb{Z}[\phi_{\mathrm{M}}]$ if and only if $\ell \equiv \pm 1 \bmod 8$.*

Proof. Let P and Q be generic points of $\mathcal{C}_{\mathrm{M}}$. An elementary but lengthy calculation shows that on the level of divisors,

$$\phi_{\mathrm{M}}^2((P) - (Q)) = 2(P) - 2(Q) + \operatorname{div}\left(\frac{x + x(P)}{x + x(Q)}\right),$$

so $\phi_{\mathrm{M}}^2([D]) = 2[D]$ for all $[D]$ in $\mathrm{Pic}^0(\mathcal{C}_{\mathrm{M}})$. Hence $\phi_{\mathrm{M}}^2 = [2]$, and $\mathbb{Z}[\phi_{\mathrm{M}}]$ is isomorphic to the maximal order of $\mathbb{Q}(\sqrt{2})$. The primes ℓ splitting in $\mathbb{Q}(\sqrt{2})$ are precisely those congruent to ± 1 modulo 8; further, $\mathbb{Q}(\sqrt{2})$ has class number 1, so the primes over ℓ are principal. □

Remark 3. As noted above, this construction is a twisted reparametrization of a family of isogenies described by Mestre in [14, §2.1]. Let a_1 and a_2 be the roots of $T^2 - tT + v^2t - 1$ in $\overline{\mathbb{F}_q(v,t)}$. Mestre's curves C' and C are equal (over $\mathbb{F}_q(v, a_1, a_2)$) to our $\mathcal{C}_{\mathrm{M}}$ and its quadratic twist by $A = 2(v^2 - 1)(v^2 + 1)^2 = (2n)^2$, respectively. We may specialize the proofs in [14] to show that $\mathcal{C}_{\mathrm{M}}$ has a two-dimensional image in the moduli space of curves of genus 2, and that the Jacobian of the curve with the same defining equation over $\mathbb{Q}(s,t)$ is absolutely simple. Constructions of curves with RM by $\mathbb{Z}[\sqrt{2}]$ are further investigated in Bending's thesis [1].

Remark 4. The algorithms described here should be readily adaptable to work with Kummer surfaces instead of Jacobians. In the notation of [6], the Kummers with parameters (a, b, c, d) satisfying $b^2 = a^2 - c^2 - d^2$ have RM by $\mathbb{Z}[\sqrt{2}]$, which can be made explicit as follows: the doubling algorithm decomposes into two identical steps, since $(A : B : C : D) = (a : b : c : d)$, and the components after one step are the coordinates of a Kummer point. The step therefore defines an efficiently computable endomorphism which squares to give multiplication by 2.

5 Numerical Experiments

We implemented our algorithm in C++ using the NTL library. For non-critical steps, including computations in quadratic fields, we used Magma for simplicity. With this implementation, the determination of $\chi(T)$ for a curve over a 128-bit prime field takes approximately 3 hours on one core of a Core2 processor at 2.83 GHz. This provides a proof of concept rather than an optimized implementation.

5.1 Cryptographic Curve Generation

When looking for a cryptographic curve, we used an early-abort strategy, where we switch to another curve as soon as either the order of the Jacobian order or its twist can not be prime.

Using our adapted version of Schoof algorithm, we guarantee that the group orders are not divisible by any prime that splits in the real field up to the CRT bound used.

In fact, any prime that divides the group order of a curve having RM by the maximal order of $\mathbb{Q}(\sqrt{\Delta})$ must either be a split (or ramified) prime, or divide it with multiplicity 2. As a consequence, the early abort strategy works much better than in the classical Schoof algorithm, because, it suffices to test half the number of primes up to our CRT bound.

We ran a search for a secure curve over a prime field of 128 bits, using a CRT bound of 131. Our series of computations frequently aborted early, and resulted in 245 curves for which $\chi(T)$ was fully determined, and for which neither the group order nor its twist was divisible by a prime less than 131. Considering these twists, this provided 490 group orders, of which 27 were prime, and therefore suitable for cryptographic use. We give here the data for one of these curves, that was furthermore twist-secure: both the Jacobian and the twist Jacobian order are prime.

Let $\mathcal{C}/\mathbb{F}_q$, where $q = 2^{128} + 573$, be the curve in the family $\mathcal{C}_{\mathrm{T}}$ of §4.2 specialized to $t = 75146620714142230387068843744286456025$. The characteristic polynomial $\chi(T)$ is determined by

$$(s_1, s_2) = (-26279773936397091867,\ -908270641821524281611387087874126 43439),$$

giving prime group orders for the Jacobian:

$$115792089237316195432513528685912298808995809621534164533135283195301868637471,$$

and for its twist:

$$115792089237316195414628441331463517678650820031857370801365706066289379517451.$$

We note that correctness of the orders is easily verified on random points in the Jacobians.

5.2 A Kilobit Jacobian

Let q be the prime $2^{512} + 1273$, and consider the curve over $\mathbb{F}_q$ from the family $\mathcal{C}_{\mathrm{T}}$ of §4.2 specialized at

$$\begin{aligned} t = {} & 290856663337872724379982611299198017497745330036809577622 3 \\ & 256986807375270272014471477919882845604269700820270816721 5 \\ & 324349759210853165605908326591223512 78. \end{aligned}$$

This value of t was randomly chosen, and carries no special structure. We computed the values of the pair $(s_1 \bmod \ell, n \bmod \ell)$ for this curve for each split prime ℓ up to 419; this was enough to uniquely determine the true value of (s_1, n) using the Chinese Remainder Theorem. The numerical data for the curve follows:

$$\begin{aligned} \Delta &= 5 \\ s_1 &= -10535684568225216385772683270554282199378670073368228748 \\ &\quad 7810402851346035223080 \\ n &= -37786020778198256317368570028183842800473749792142072230 \\ &\quad 993549001035093288492 \\ s_2 &= (s_1^2 - n^2\Delta)/4 \\ &= 990287025215436155679872249605061232893936642355960654938 \\ &\quad 008045777052233348340624693986425546428828954551752076384 \\ &\quad 428888704295617466043679591527916629020 \end{aligned}$$

The order of the Jacobian is therefore

$$\begin{aligned} N &= (1+q)^2 - s_1(1+q) + s_2 \\ &= 17976931348623159077293051907890247336179769789423065727343008115773267580550237573705948956144184541720417180780929444962763452801227364805323818926258902074851818089888868757737237328920325315884646393462965754493894524803468668112345681706310648544084486938739666585942218663644225871268417790010511900552 0. \end{aligned}$$

The total runtime for this computation was about 80 days on a single core of a Core 2 clocked at 2.83 GHz. In practice, we use the inherent parallelism of the algorithm, running one prime ℓ on each available core.

We did not compute the characteristic polynomial modulo small prime powers (as in [9]), nor did we use BSGS to deduce the result from partial modular information as in §3.4 (indeed, we were more interested in measuring the behaviour of our algorithm for large values of ℓ). These improvements with an exponential-complexity nature bring much less than in the classical point counting algorithms, since they have to be balanced with a polynomial-time algorithm with a lower degree. For this example, we estimate that BSGS and small prime powers could have saved a factor of about 2 in the total runtime.

5.3 Degrees of Division Polynomials

For each prime ℓ splitting in $\mathbb{Z}[\phi_{\mathrm{T}}]$, we report the degree of the α-division polynomial d_2 (where α is the endomorphism of norm ℓ that was used). By Lemma 1, $\deg(d_2)$ is in $O(\ell)$; the table below gives the ratio $\deg(d_2)/\ell$, thus measuring the hidden constant in the $O()$ notation.

ℓ	11	19	29	31	41	59	61	71	79	89	101	109	131
$\deg d_2/\ell$	1.82	2.05	2.07	1.94	2.05	2.10	1.97	2.03	2.01	2.02	1.98	2.02	2.02
ℓ	139	149	151	179	181	191	199	211	229	239	241	251	269
$\deg d_2/\ell$	2.12	2.04	1.99	2.00	2.01	2.09	2.21	1.99	2.18	2.01	2.05	2.07	2.17
ℓ	271	281	311	331	349	359	379	389	401	409	419		
$\deg d_2/\ell$	2.01	1.99	2.11	2.12	2.13	2.02	2.00	2.16	2.03	2.10	2.00		

We have $\deg(d_1) = \deg(d_2) + 1$ and $\deg(d_0) = \deg(d_2) + 2$. All of these degrees depend only on the curve family $\mathcal{C}_{\mathrm{T}}$, and not on the individual curve chosen.

References

1. P. R. Bending, 'Curves of genus 2 with $\sqrt{2}$ multiplication'. Ph. D. thesis, University of Oxford (1998)
2. D. G. Cantor, 'Computing in the Jacobian of a hyperelliptic curve', *Math. Comp.* **48** no. 177 (1987) 95–101
3. D. G. Cantor, 'On the analogue of the division polynomials for hyperelliptic curves', *J. Reine Angew. Math.* **447** (1994) 91–145
4. H. Cohen and H. W. Lenstra, Jr., 'Heuristics on class groups of number fields', in *Number theory, Noordwijkerhout 1983 (Noordwijkerhout, 1983)*, *Lecture Notes in Math.*, **1068** (1984), 33–62.
5. S. D. Galbraith, M. C. Harrison, and D. J. Mireles Morales, 'Efficient hyperelliptic arithmetic using balanced representation for divisors', in *Algorithmic Number Theory: ANTS-IX*, LNCS 5011 (2008) 342–356
6. P. Gaudry, 'Fast genus 2 arithmetic based on Theta functions', *J. Math. Crypt.* **1** (2007) 243–265

7. P. Gaudry and É. Schost, 'Construction of secure random curves of genus 2 over prime fields', in *Advances in cryptology: EUROCRYPT 2004*, LNCS 3027 (2004) 239–256
8. P. Gaudry and É. Schost, 'A low-memory parallel version of Matsuo, Chao, and Tsuji's algorithm', in *Algorithmic number theory: ANTS-VI*, LNCS 3076 (2004) 208–222
9. P. Gaudry and É. Schost, Genus 2 point counting over prime fields, Preprint 2010. `http://hal.inria.fr/inria-00542650`
10. D. Gruenewald, 'Computing Humbert surfaces and applications', in *Arithmetic, Geometry, Cryptography and Codint Theory 2009*, Contemp. Math. 521 (2010), 59–69.
11. D. R. Kohel and B. Smith, 'Efficiently computable endomorphisms for hyperelliptic curves', in *Algorithmic number theory: ANTS-VII*, LNCS 4076 (2006) 495–509
12. K. Matsuo, J. Chao, and S. Tsuji, 'An improved baby step giant step algorithm for point counting of hyperelliptic curves over finite fields', in *Algorithmic number theory: ANTS-V*, LNCS 2369 (2002) 461–74
13. J.-F. Mestre, 'Familles de courbes hyperelliptiques à multiplications réelles', *Arithmetic algebraic geometry (Texel, 1989)*, *Progr. Math.* 89 (Birkhäuser Boston, 1991).
14. J.-F. Mestre, 'Couples de jacobiennes isogènes de courbes hyperelliptiques de genre arbitraire', Preprint, 2009, arXiv math.AG / 0902.3470 v1.
15. Y.H. Park, S. Jeong and J. Lim, 'Speeding up point multiplication on hyperelliptic curves with efficiently-computable endomorphisms', in *Advances in Cryptology—EUROCRYPT 2002*, LNCS 2332 (2002), 197–208.
16. J. Pila, 'Frobenius maps of abelian varieties and finding roots of unity in finite fields', *Math. Comp.* 55 (1990) no. 192, 745–763.
17. H.-G. Rück, 'Abelian surfaces and jacobian varieties over finite fields', *Compositio Math.* 76 (1990) no. 3, 351–366.
18. K. Takashima, A new type of fast endomorphisms on Jacobians of hyperelliptic curves and their cryptographic application. *IEICE Trans. Fundamentals*, **E89-A** (2006), no. 1, pp. 124-133.
19. W. Tautz, J. Top, and A. Verberkmoes, 'Explicit hyperelliptic curves with real multiplication and permutation polynomials', *Canad. J. Math.* 43 (1991) no. 5, 1055–1064.
20. J. Wilson, 'Curves of genus 2 with real multiplicaiton by a square root of 5'. Ph.D. thesis, University of Oxford (1998)