

Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication

Sean Ballentine, Aurore Guillevic, Elisa Lorenzo García, Chloe Martindale,
Maïke Massierer, Benjamin Smith, Jaap Top, Elisa García

► **To cite this version:**

Sean Ballentine, Aurore Guillevic, Elisa Lorenzo García, Chloe Martindale, Maïke Massierer, et al.. Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication. E. W. Howe; K. E. Lauter; J. L. Walker. Algebraic Geometry for Coding Theory and Cryptography, Feb 2016, Los Angeles, United States. Springer, Association for Women in Mathematics Series, 9, pp.63-94, 2017, Algebraic Geometry for Coding Theory and Cryptography. <<http://www.ipam.ucla.edu/programs/workshops/algebraic-geometry-for-coding-theory-and-cryptography/>>. <10.1007/978-3-319-63931-4>. <hal-01421031v3>

HAL Id: hal-01421031

<https://hal.inria.fr/hal-01421031v3>

Submitted on 7 Nov 2017

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.

Chapter 3

Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication

Sean Ballentine, Aurore Guillevic, Elisa Lorenzo García, Chloe Martindale,
Maïke Massierer, Benjamin Smith, and Jaap Top

Abstract Schoof’s classic algorithm allows point-counting for elliptic curves over finite fields in polynomial time. This algorithm was subsequently improved by Atkin, using factorizations of modular polynomials, and by Elkies, using a theory of explicit isogenies. Moving to Jacobians of genus-2 curves, the current state of the art for point counting is a generalization of Schoof’s algorithm. While we are currently missing the tools we need to generalize Elkies’ methods to genus 2, recently Martindale and Milio have computed analogues of modular polynomials for genus-2 curves whose Jacobians have real multiplication by maximal orders of small discriminant. In this article, we prove Atkin-style results for genus-2 Jacobians with real multiplication

Sean Ballentine

Department of Mathematics, University of Maryland, 4176 Campus Dr., College Park, MD 20742-4015, USA e-mail: seanbal@math.umd.edu

Aurore Guillevic

Inria Nancy Grand Est, Equipe Caramba, 615 rue du jardin botanique, CS 20101, 54603 Villers-lès-Nancy Cedex, France e-mail: aurore.guillevic@inria.fr

Elisa Lorenzo García

IRMAR, Université de Rennes 1, Campus de Beaulieu, 35042 Rennes Cedex, France e-mail: elisa.lorenzogarcia@univ-rennes1.fr

Chloe Martindale

Mathematisch Instituut, Universiteit Leiden, P.O. Box 9512, 2300 RA Leiden, The Netherlands e-mail: chloemartindale@gmail.com

Maïke Massierer

School of Mathematics and Statistics, University of New South Wales, Sydney NSW 2052, Australia e-mail: maïke@unsw.edu.au

Benjamin Smith

INRIA and Laboratoire d’Informatique de l’École polytechnique (LIX), 91120 Palaiseau, France e-mail: smith@lix.polytechnique.fr

Jaap Top

University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science, P.O. Box 407, 9700AK Groningen, The Netherlands e-mail: j.top@rug.nl

by maximal orders, with a view to using these new modular polynomials to improve the practicality of point-counting algorithms for these curves.

3.1 Introduction

Efficiently computing the number of points on the Jacobian of a genus 2 curve over a finite field is an important problem in experimental number theory and number-theoretic cryptography. When the characteristic of the finite field is small, Kedlaya's algorithm and its descendants provide an efficient solution (see [18], [13], and [12]), while in extremely small characteristic we have extremely fast AGM-style algorithms (see for example [25], [26], and [3]). However, the running times of these algorithms are exponential in the size of the field characteristic; the hardest case, therefore (and also the most important case for contemporary cryptographic applications) is where the characteristic is large, or even where the field is a prime field.

So let q be a power of a large prime p , and let \mathcal{C} be a genus-2 curve over \mathbb{F}_q . Our fundamental problem is to compute the number of \mathbb{F}_q -rational points on the Jacobian $J_{\mathcal{C}}$ of \mathcal{C} .

3.1.1 The state of the art

In theory, the problem is solved: we can compute $\#J_{\mathcal{C}}(\mathbb{F}_q)$ in polynomial time (that is, polynomial in $\log q$) using Pila's algorithm [31], which is the immediate generalization of Schoof's elliptic-curve point-counting algorithm [34] to higher-dimensional abelian varieties. But the exponent in Pila's polynomial time is extremely large; so, despite its theoretical importance, this algorithm is completely impractical (see §3.3.4). Indeed, to our knowledge it has never been implemented.

Gaudry and Schost have developed and successfully implemented a much more practical variant of Pila's algorithm for the case $q = p$ that runs in time $\tilde{O}(\log^8 p)$; not just polynomial time, but on the edge of practicality [8]. Still, their algorithm requires an extremely intensive calculation for cryptographic-sized Jacobians: Gaudry and Schost estimated a running time of around one core-month (in 2008) to compute $\#J_{\mathcal{C}}(\mathbb{F}_p)$ when p has around 128 bits [8].

The situation improves dramatically if $J_{\mathcal{C}}$ is equipped with an efficiently computable *real multiplication* endomorphism. For such Jacobians, Gaudry, Kohel, and Smith [7] give an algorithm to compute $\#J_{\mathcal{C}}(\mathbb{F}_q)$ in time $\tilde{O}(\log^5 q)$. This allowed the computation of $\#J_{\mathcal{C}}(\mathbb{F}_p)$ for one curve \mathcal{C} drawn from the genus-2 family in [37] with $p = 2^{512} + 1273$ in about 80 core-days (in 2011); this remains, to date, the record for genus-2 point counting over prime fields. For 128-bit fields, the cost is reduced to 3 core hours (in 2011).

All of these algorithms are generalizations of Schoof's algorithm, which computes the Frobenius trace (and hence the order $\#E(\mathbb{F}_q)$) of an elliptic curve E/\mathbb{F}_q

modulo ℓ for a series of small primes ℓ by considering the action of Frobenius on the ℓ -torsion. But Schoof’s algorithm is not the state of the art for elliptic-curve point counting: it has evolved into the much faster Schoof–Elkies–Atkin (SEA) algorithm, surveyed in [35]. Atkin’s improvements involve factoring the ℓ -th modular polynomial (evaluated at the j -invariant of the target curve) to deduce information on the Galois structure of the ℓ -torsion, which then restricts the possible values of the trace modulo ℓ (see §3.2.6). Elkies’ improvements involve computing the kernel of a rational ℓ -isogeny, which takes the place of the full ℓ -torsion; deducing the existence of the isogeny, and computing its kernel, requires finding a root of the ℓ -th modular polynomial evaluated at the j -invariant of the target curve (see §3.2.7).

3.1.2 *Our contributions, and beyond*

Our ultimate goal is to generalize Atkin’s and Elkies’ improvements to genus 2. In this article, we concentrate on generalizing Atkin’s methods to genus-2 Jacobians with known real multiplication. This project is prompted by the recent appearance of two new algorithms for computing modular ideals, the genus-2 analogue of modular polynomials: Milio [28] has computed modular ideals for general genus-2 Jacobians, while Milio [27, §5] and Martindale [24] have independently computed modular ideals for genus-2 Jacobians with RM by orders of small discriminants.

To extend Elkies’ methods to genus 2 we would need an analogue of Elkies’ algorithm [35, §§7-8], which computes defining equations for the kernel of an isogeny of elliptic curves (and the isogeny itself) corresponding to a root of the evaluated modular polynomial. We do not know of any such algorithm in genus 2. Couveignes and Ezome have recently developed an algorithm to compute explicit (ℓ, ℓ) -isogenies of genus-2 Jacobians [4], presuming that the kernel has already been constructed somehow—but kernel construction is precisely the missing step that we need.¹

In contrast, Atkin’s improvements for elliptic-curve Schoof require nothing beyond the modular polynomial itself; so we can hope to achieve something immediately in genus 2 by generalizing Atkin’s results on factorizations of modular polynomials to the decomposition of genus-2 modular ideals. This is precisely what we do in this article.

We focus on the RM case for three reasons. First, the construction of explicit modular ideals is furthest advanced in this case: Milio has constructed modular ideals for primes in $\mathbb{Q}(\sqrt{5})$ of norm up to 31, while for general Jacobians the current limit is 3. It is therefore already possible to compute nontrivial and interesting examples in the RM case. Second, point counting is currently much more efficient for Jacobians with efficiently computable RM; we hope that, at some point, our methods can help

¹ We would also like mention Bisson, Cosset, and Robert’s `AVIsogenies` software package [1], which provides some functionality in this direction. However, their methods apply to abelian surfaces with a lot of rational 2- and 4-torsion, and applying them to general genus-2 Jacobians (with or without known RM) generally requires a substantial extension of the base field to make that torsion rational. This is counterproductive in the context of point counting.

tip RM point counting from “feasible” into “routine”. Third, from a purely theoretical point of view, the RM case is more similar to the elliptic curve case in the sense that real multiplication allows us, in favorable circumstances, to split ℓ -torsion subgroups of the Jacobian into groups of the same size as encountered for elliptic curves.

After recalling the SEA algorithm for elliptic curves in §3.2, we describe the current state of genus 2 point counting, and set out our program for a generalized SEA algorithm in §3.3. We describe the modular invariants we need for this in §3.4, and the modular ideals that relate them in §3.4.2. We can then state and prove our main theoretical results, which are generalizations of Atkin’s theorems for these modular ideals, in §3.5. In §3.6 we provide some concrete details on the special case of RM by $\mathbb{Q}(\sqrt{5})$, before concluding with some experimental results in §3.7.

3.1.3 Vanilla abelian varieties

We can substantially simplify the task ahead by restricting our attention to a class of elliptic curves and Jacobians (more generally, abelian varieties) with sufficiently general CM endomorphism rings. The following definition makes this precise.

Definition 1. We say that a g -dimensional abelian variety \mathcal{A}/\mathbb{F}_q is *vanilla*² if its endomorphism algebra $\text{End}_{\overline{\mathbb{F}_q}}(\mathcal{A}) \otimes \mathbb{Q}$ (over the algebraic closure) is a CM field of degree $2g$ that does *not* contain any roots of unity other than ± 1 .

If an elliptic curve \mathcal{E}/\mathbb{F}_q is vanilla, then \mathcal{E} is nonsupersingular and $j(\mathcal{E})$ is neither 0 nor 1728: these are the conditions Schoof applies systematically in [35]. We note that in particular, vanilla abelian varieties are absolutely simple.

To fix notation, we recall that if \mathcal{A} is an abelian variety, then a *principal polarization* is an isomorphism $\xi: \mathcal{A} \rightarrow \mathcal{A}^\vee$ associated with an ample divisor class on \mathcal{A} , where $\mathcal{A}^\vee = \text{Pic}^0(\mathcal{A})$ is the dual abelian variety (see for example [29, §13]). We will be working with elliptic curves and Jacobians of genus-2 curves; these all have a canonical principal polarization. Each endomorphism ϕ of \mathcal{A} has a corresponding dual endomorphism ϕ^\vee of \mathcal{A}^\vee . If (\mathcal{A}, ξ) is a principally polarized abelian variety, then ξ induces a *Rosati involution* on $\text{End}(\mathcal{A})$, defined by

$$\phi \longmapsto \phi^\dagger := \xi^{-1} \circ \phi^\vee \circ \xi \quad \text{for } \phi \in \text{End}(\mathcal{A}).$$

In the world of elliptic curves, the Rosati involution is the familiar dual. For vanilla abelian varieties, the Rosati involution acts as complex conjugation on the endomorphism ring.

² Vanilla is the most common and least complicated flavour of abelian varieties over finite fields. Heuristically, over large finite fields, randomly sampled abelian varieties are vanilla with overwhelming probability. Indeed, being vanilla is invariant in isogeny classes, and Howe and Zhu have shown in [14, Theorem 2] that the fraction of isogeny classes of g -dimensional abelian varieties over \mathbb{F}_q that are ordinary and absolutely simple tends to 1 as $q \rightarrow \infty$. All absolutely simple ordinary abelian varieties are vanilla, except those whose endomorphism algebras contain roots of unity; but the number of such isogeny classes for fixed g is asymptotically negligible.

Fix a real quadratic field $F = \mathbb{Q}(\sqrt{\Delta})$, with fundamental discriminant $\Delta > 0$ and ring of integers \mathcal{O}_F . We write $\alpha \mapsto \bar{\alpha}$ for the involution of F over \mathbb{Q} ; we emphasize that in this article, $\bar{\cdot}$ does *not* denote complex conjugation.

From a theoretical point of view, when talking about real multiplication, our fundamental data are triples $(\mathcal{A}, \xi, \iota)$ where \mathcal{A} is an abelian surface, $\xi: \mathcal{A} \rightarrow \mathcal{A}^\vee$ is a principal polarization, and $\iota: \mathcal{O}_F \hookrightarrow \text{End}(\mathcal{A})$ is an embedding stable under the Rosati involution (that is, $\iota(\mu)^\dagger = \iota(\mu)$ for all μ in \mathcal{O}_F ; we can then think of the Rosati involution as complex conjugation on the endomorphism ring). While this notation $(\mathcal{A}, \xi, \iota)$ may seem quite heavy at first glance, we remind the reader that generally there are only two choices of embedding ι (corresponding to the two square roots of Δ), and we are only really interested in the case where \mathcal{A} is a Jacobian, in which case the polarization ξ is canonically determined.

3.2 Genus one curves: elliptic curve point counting

We begin by briefly recalling the SEA algorithm for elliptic curve point counting in large characteristic. First we describe Schoof's original algorithm [35], before outlining the improvements of Elkies and Atkin. This will provide a point of reference for comparisons with genus-2 algorithms.

Let \mathcal{E} be an elliptic curve over a finite field \mathbb{F}_q of large characteristic (or at least, with $\text{char}(\mathbb{F}_q) \gg \log q$). We may suppose that \mathcal{E} is defined by a (short) Weierstrass equation $\mathcal{E}: y^2 = x^3 + ax + b$, with a and b in \mathbb{F}_q .

Like all modern point-counting algorithms, the Schoof and SEA algorithms compute the characteristic polynomial

$$\chi_\pi(X) = X^2 - tX + q$$

of the Frobenius endomorphism π of \mathcal{E} . We call t the *trace* of Frobenius. Since the \mathbb{F}_q -rational points on \mathcal{E} are precisely the fixed points of π , we have

$$\#\mathcal{E}(\mathbb{F}_q) = \chi_\pi(1) = q + 1 - t;$$

so determining $\#\mathcal{E}(\mathbb{F}_q)$ is equivalent to determining t . Hasse's theorem tells us that

$$|t| \leq 2\sqrt{q}. \tag{3.1}$$

3.2.1 Schoof's algorithm

Schoof's basic strategy is to choose a set \mathcal{L} of primes $\ell \neq p$ such that $\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q}$. We then compute $t_\ell := t \bmod \ell$ for each of the primes ℓ in \mathcal{L} , and then recover the value of t from $\{(t_\ell, \ell) : \ell \in \mathcal{L}\}$ using the Chinese Remainder Theorem. The condition

$\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q}$ ensures that t is completely determined by the collection of t_ℓ (by Hasse's theorem, Equation (3.1)).

For Schoof's original algorithm, the natural choice is to let \mathcal{L} be the set of the first $O(\log q)$ primes, stopping when the condition $\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q}$ is satisfied. When applying Elkies' and Atkin's modifications, we will need to be more subtle with our choice of \mathcal{L} . It is also possible to replace primes with small prime powers; we will not explore this option here.

Now, let ℓ be one of our primes in \mathcal{L} ; our aim is to compute t_ℓ . We know that $\pi^2(P) - [t]\pi(P) + [q]P = 0$ for all P in \mathcal{E} , and hence

$$\pi^2(P) - [t_\ell]\pi(P) + [q \bmod \ell]P = 0 \quad \text{for all } P \in \mathcal{E}[\ell].$$

We can therefore compute t_ℓ as follows:

1. Construct a point P of order ℓ .
2. Compute $Q = \pi(P)$ and $R = \pi^2(P) + [q \bmod \ell]P$.
3. Search for $0 \leq t_\ell < \ell$ such that $[t_\ell]Q = R$, using Shanks' baby-step giant-step algorithm in the cyclic subgroup of the ℓ -torsion generated by Q .

To construct such a P , we begin by computing the ℓ -th division polynomial Ψ_ℓ in $\mathbb{F}_q[X]$, which is the polynomial whose roots in $\overline{\mathbb{F}_q}$ are precisely the x -coordinates of the nontrivial points in $\mathcal{E}[\ell]$. When ℓ is odd and prime to q , we have $\deg \Psi_\ell = (\ell^2 - 1)/2$. We then define the ring $A = \mathbb{F}_q[X, Y]/(\Psi_\ell(X), Y^2 - X^3 - aX - b)$, and take $P = (X, Y)$ in $\mathcal{E}(A)$.

In order to work efficiently with $Q = \pi(P) = (X^q, Y^q)$ in the search for t_ℓ , we need to compute a compact form for Q . This means computing reduced representatives for X^q and Y^q in the ring A —that is, reducing X^q modulo $\Psi_\ell(X)$ and Y^q modulo $(\Psi_\ell(X), Y^2 - X^3 - aX - b)$ —which costs $O(\log q)$ \mathbb{F}_q -operations.

Having computed t_ℓ for each ℓ in \mathcal{L} , we recover t (and hence χ_π) using the Chinese Remainder Theorem; this then yields $\#\mathcal{E}(\mathbb{F}_q) = q + 1 - t$. In cryptographic contexts, we are generally interested in curves of (almost) prime order. One particularly convenient feature of Schoof's algorithm is that it allows us to detect small prime factors of $\#\mathcal{E}(\mathbb{F}_q)$ early: we can determine if any ℓ in \mathcal{L} divides $\#\mathcal{E}(\mathbb{F}_q)$ by checking whether $t_\ell \equiv q + 1 \pmod{\ell}$. If we find such a factor, then we can immediately abort the calculation of t and move on to another candidate curve.

The cost to compute χ_ℓ is $\tilde{O}(\ell^2 + (\log q)\ell^2 + \sqrt{\ell}\ell^2)$ \mathbb{F}_q -operations. We can take \mathcal{L} to be a set of $O(\log q)$ primes, the largest of which is in $O(\log q)$; the total cost is therefore $\tilde{O}(\log^4 q)$ \mathbb{F}_q -operations.

3.2.2 Frobenius eigenvalues and subgroups

Fix a basis of $\mathcal{E}[\ell]$, and thus an isomorphism $\mathcal{E}[\ell] \cong \mathbb{F}_\ell^2$. Now π acts on $\mathcal{E}[\ell]$ as an element of $\text{GL}_2(\mathbb{F}_\ell)$. The local characteristic polynomial χ_ℓ is just the characteristic polynomial of this matrix.

Likewise, π permutes the ℓ -subgroups of $\mathcal{E}[\ell]$; that is, the one-dimensional subspaces of $\mathcal{E}[\ell] \cong \mathbb{F}_\ell^2$. These are the points of $\mathbb{P}(\mathcal{E}[\ell]) \cong \mathbb{P}^1(\mathbb{F}_\ell)$, and we can consider the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell) \cong \mathrm{Aut}(\mathbb{P}(\mathcal{E}[\ell]))$. The order of π as an element of $\mathrm{PGL}_2(\mathbb{F}_\ell)$ is clearly independent of the choice of basis.

Proposition 1. *Let \mathcal{E}/\mathbb{F}_q be an elliptic curve with Frobenius endomorphism π , and let $\ell \neq p = \mathrm{char}(\mathbb{F}_q)$ be an odd prime. If e is the order of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$, then the trace t of π satisfies*

$$t^2 = \eta_e q \quad \text{in } \mathbb{F}_\ell ,$$

$$\text{where } \eta_e = \begin{cases} \zeta + \zeta^{-1} + 2 \text{ with } \zeta \in \mathbb{F}_{\ell^2}^\times \text{ of order } e & \text{if } \gcd(\ell, e) = 1 , \\ 4 & \text{otherwise .} \end{cases}$$

Proof. We follow the proof of [35, Proposition 6.2] (correcting the minor error that leads in the case e even to an $e/2$ -th rather than e -th root of unity appearing in the last part of the statement). Let $\lambda_1, \lambda_2 \in \mathbb{F}_{\ell^2}$ be the eigenvalues of the image of π in $\mathrm{Aut}(\mathcal{E}[\ell]) \cong \mathrm{GL}_2(\mathbb{F}_\ell)$; then

$$\lambda_1 + \lambda_2 = t \quad \text{and} \quad \lambda_1 \lambda_2 = q \quad \text{in } \mathbb{F}_\ell .$$

In case $\lambda_1 = \lambda_2$ we have $e \mid \ell$ and the assertion follows. In case $\lambda_1 \neq \lambda_2$ the given e is the minimal integer > 0 with $\lambda_1^e = \lambda_2^e$. In particular $\gcd(e, \ell) = 1$ and $\lambda_2 = \lambda_1 \zeta$ for some primitive e -th root of unity ζ (in \mathbb{F}_{ℓ^2} ; in fact $e \mid \ell - 1$ in case the eigenvalues are in \mathbb{F}_ℓ and $e \mid \ell + 1$ otherwise). Hence $q = \lambda_1 \lambda_2 = \lambda_1^2 \zeta$ which implies

$$t^2 = (\lambda_1 + \lambda_2)^2 = \lambda_1^2 (1 + \zeta)^2 = q \zeta^{-1} (\zeta^2 + 2\zeta + 1) = (\zeta + \zeta^{-1} + 2)q .$$

3.2.3 Modular polynomials and isogenies

The order- ℓ subgroups of $\mathcal{E}[\ell]$ are precisely the kernels of ℓ -isogenies from \mathcal{E} to other elliptic curves, and the set of all such ℓ -isogenies (up to isomorphism) corresponds to the set of roots of $\Phi_\ell(j(\mathcal{E}), x)$ in $\overline{\mathbb{F}}_q$. The classical modular polynomial $\Phi_\ell(X, Y)$, of degree $\ell + 1$ (in X and Y) over \mathbb{Z} , is defined by the property that $\Phi_\ell(j(\mathcal{E}_1), j(\mathcal{E}_2)) = 0$ precisely when there exists an ℓ -isogeny $\mathcal{E}_1 \rightarrow \mathcal{E}_2$. For ℓ in $O(\log q)$, one can compute $\Phi_\ell(j(\mathcal{E}), x)$ in $\tilde{O}(\ell^3)$ \mathbb{F}_q -operations using Sutherland's algorithm [36]. Alternatively, we can use precomputed databases of modular polynomials over \mathbb{Z} , reducing them modulo p and specializing them at $j(\mathcal{E})$.

The Galois orbits of the roots of $\Phi_\ell(j(\mathcal{E}), x)$ correspond to orbits of ℓ -isogeny kernels under π , and to orbits of points of $\mathbb{P}^1(\mathbb{F}_\ell)$ under the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$. If $j(\mathcal{E}_1)$ and $j(\mathcal{E}_2)$ are both in \mathbb{F}_{q^k} , then the isogeny is defined over \mathbb{F}_{q^k} (up to a possible twist); in particular, its kernel is defined over \mathbb{F}_{q^k} . More precisely, we have the following key lemma:

Lemma 1 (Proposition 6.1 of [35]). Let \mathcal{E}/\mathbb{F}_q be a vanilla elliptic curve with Frobenius endomorphism π .

1. The polynomial $\Phi_\ell(j(\mathcal{E}), x)$ has a root in \mathbb{F}_{q^e} if and only if the kernel of the corresponding ℓ -isogeny is a one-dimensional eigenspace of π^e in $\mathcal{E}[\ell]$.
2. The polynomial $\Phi_\ell(j(\mathcal{E}), x)$ splits completely over \mathbb{F}_{q^d} if and only if π^d acts as a scalar matrix on $\mathcal{E}[\ell]$; that is, if and only if d is a multiple of the order e of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$. In particular, the minimal such d is e .

3.2.4 Elkies, Atkin, and volcanic primes

The primes $\ell \neq p$ are divided into 3 classes, or types, with respect to a given \mathcal{E}/\mathbb{F}_q : *Elkies*, *Atkin*, and *volcanic*. The type of ℓ simultaneously reflects the factorization of $\Phi_\ell(j(\mathcal{E}), x)$ and the Galois structure of the ℓ -subgroups of $\mathcal{E}[\ell]$. Here we recall a number of facts about these classes, all of which are proven in [35, §6]; see also [38, §12.4].

A prime ℓ is **Elkies** if the ideal (ℓ) is split in $\mathbb{Z}[\pi]$; or, equivalently, if $t^2 - 4q$ is a nonzero square modulo ℓ . Each of the two prime ideals over (ℓ) defines the kernel of an ℓ -isogeny, $\phi_i: \mathcal{E} \rightarrow \mathcal{E}_i$ for $i = 1, 2$, say. This means that $j(\mathcal{E}_1)$ and $j(\mathcal{E}_2)$ must be roots in \mathbb{F}_q of $\Phi_\ell(j(\mathcal{E}), x)$. Lemma 1 then implies that

$$\Phi_\ell(j(\mathcal{E}), x) = (x - j(\mathcal{E}_1))(x - j(\mathcal{E}_2)) \prod_{i=1}^{(\ell-1)/e} f_i(x)$$

where each of the f_i are irreducible of degree e , and $e > 1$ is the order of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$, which must divide $\ell - 1$ in this case.

A prime ℓ is **Atkin** if the ideal (ℓ) is inert in $\mathbb{Z}[\pi]$; or, equivalently, if $t^2 - 4q$ is *not* a square modulo ℓ . There are *no* \mathbb{F}_q -rational ℓ -isogenies from \mathcal{E} , and no \mathbb{F}_q -rational ℓ -subgroups of $\mathcal{E}[\ell]$. Looking at the modular polynomial, Lemma 1 implies

$$\Phi_\ell(j(\mathcal{E}), x) = \prod_{i=1}^{(\ell+1)/e} f_i(x),$$

where each of the f_i is an irreducible polynomial of degree e , and $e > 1$ is the order of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$, which must divide $\ell + 1$ in this case.

Finally, a prime ℓ is **volcanic** if the ideal (ℓ) is ramified in $\mathbb{Z}[\pi]$; or, equivalently, if ℓ divides $t^2 - 4q$. Applying Lemma 1, either

$$\Phi_\ell(j(\mathcal{E}), x) = \prod_{i=1}^{\ell+1} (x - j_i)$$

with all of the j_i in \mathbb{F}_q (so there are $\ell + 1$ rational ℓ -isogenies, and $\ell + 1$ rational ℓ -subgroups of $\mathcal{E}[\ell]$); or

$$\Phi_\ell(j(\mathcal{E}), x) = (Y - j_1) \cdot f(x) ,$$

with f irreducible of degree ℓ (so there is a single rational ℓ -isogeny, and one rational ℓ -subgroup of $\mathcal{E}[\ell]$). In either situation, $\pi|_{\mathcal{E}[\ell]}$ acts on $\mathcal{E}[\ell]$ with eigenvalues $\lambda_1 = \lambda_2$, so its image in $\mathrm{PGL}_2(\mathbb{F}_\ell)$ therefore has order $e \mid \ell$.

We note an interesting and useful fact in passing: if \mathcal{E}/\mathbb{F}_q is vanilla, $\ell \neq p$ is an odd prime, and r is the number of irreducible factors of $\Phi_\ell(j(\mathcal{E}), x)$, then

$$(-1)^r = \left(\frac{q}{\ell}\right) \tag{3.2}$$

(cf. [35, Proposition 6.3]; the proof generalizes easily from $q = p$ to general prime powers).

3.2.5 Computing the type of a prime

The type of a given prime ℓ for \mathcal{E} (that is, being volcanic, Atkin, or Elkies) is defined in terms of the structure of $\mathbb{Z}[\pi]$ and the trace t . When we are point-counting, these are unknown quantities; but we can still determine the type of ℓ *without* knowing t or $\mathbb{Z}[\pi]$, by factoring $\Phi_\ell(j(\mathcal{E}), x)$ and comparing with the possible factorization types above. This, in turn, gives us useful information about t and $\mathbb{Z}[\pi]$. Determining the type of ℓ in this way costs $\tilde{O}(\ell^2 + (\log q)\ell)$ \mathbb{F}_q -operations.

In fact, computing the type of ℓ for \mathcal{E} is a good way of checking the correctness of a claimed modular polynomial. Suppose somebody has computed a polynomial $F(J_1, J_2)$, and claims it is equal to Φ_ℓ . The factorization patterns for modular polynomials corresponding to the prime types above are so special that there is very little hope of getting these patterns for $F(j(\mathcal{E}), x)$ for varying \mathcal{E} and p unless F and Φ_ℓ define the same variety in the (J_1, J_2) -plane. We will use the genus-2 analogue of this observation in §3.7 to check the correctness of some of Martindale’s modular polynomials.

3.2.6 Atkin’s improvement

Atkin’s contribution to the SEA algorithm was to exploit the factorization type of the modular polynomial to restrict the possible values of $t \pmod{\ell}$. While this does not improve the asymptotic complexity of Schoof’s algorithm, it did allow significant practical progress before the advent of Elkies’ improvements.

For example: if ℓ is volcanic, then by definition

$$t^2 = 4q \quad \text{in } \mathbb{F}_\ell , \tag{3.3}$$

which determines t_ℓ up to sign: $t \equiv \pm 2\sqrt{q} \pmod{\ell}$. Note that this is also a consequence of Proposition 1, which we will now apply to the other two prime types.

If ℓ is Elkies or Atkin for \mathcal{E} , then Proposition 1 tells us that

$$t^2 = (\zeta + \zeta^{-1} + 2)q \quad \text{in } \mathbb{F}_\ell \quad (3.4)$$

for some primitive e -th root of unity ζ in \mathbb{F}_{ℓ^2} , where $e \mid \ell - 1$ if ℓ is Elkies and $e \mid \ell + 1$ if ℓ is Atkin. The number of possible values of t_ℓ^2 is therefore half the number of primitive e -th roots in these cases. Note that modular polynomials can only give us information about t_ℓ^2 —that is, t_ℓ up to sign—since their solutions tell us about isogenies only up to quadratic twists, and twisting changes the sign of the trace.

Obviously, the smaller the degree e of the non-linear factors of $\Phi_\ell(j(\mathcal{E}), x)$, the fewer the values that t_ℓ can possibly take. For example, if $e = 2$ then $t_\ell = 0$; if $e = 3$, then $t_\ell = \pm\sqrt{q}$ in \mathbb{F}_ℓ ; and if $e = 4$, then $t_\ell = \pm\sqrt{2q}$ in \mathbb{F}_ℓ .

The challenging part of Atkin's technique is making use of these extra modular congruences. Atkin's *match-and-sort* algorithm (see for example [23, §11.2]) is a sort of sophisticated baby-step giant-step in $\mathcal{E}(\mathbb{F}_q)$ exploiting this modular information. Alternatively, we can use Joux and Lercier's *Chinese-and-match* algorithm [17].

3.2.7 Elkies' improvement

Elkies' contribution to the SEA algorithm was to note that when computing t_ℓ , we can replace $\mathcal{E}[\ell]$ with the kernel of a rational ℓ -isogeny, if it exists. Looking at the classification of primes, we see that there exists a rational ℓ -isogeny precisely when ℓ is volcanic or Elkies (whence the terminology). Of course, as we saw above, if ℓ is one of the rare volcanic primes then t_ℓ is already determined up to sign; it remains to see what can be done for Elkies primes.

Let ℓ be an Elkies prime for \mathcal{E} , and let ϕ_1 and ϕ_2 be ℓ -isogenies corresponding to the two roots of $\Phi_\ell(j(\mathcal{E}), x)$ in \mathbb{F}_q . First, we note that $\pi(P_i) = [\lambda_i]P_i$ for P_i in $\ker \phi_i$, and $\lambda_1 + \lambda_2 \equiv t \pmod{\ell}$. We only need to compute one of the λ_i , since then the other is determined by the relation $\lambda_1 \lambda_2 = q$.

So let ϕ be one of the two ℓ -isogenies; we want to compute its eigenvalue λ . The nonzero elements (x, y) of $\ker \phi$ satisfy $f_\phi(x) = 0$, where f_ϕ is a polynomial of degree $(\ell - 1)/2$ (if ℓ is odd; if $\ell = 2$, then $\deg f_\phi = 1$). To compute λ , we define the ring $A = \mathbb{F}_q[X, Y]/(f_\phi(X), Y^2 - X^3 - aX - b)$, set $P = (X, Y)$ in $\mathcal{E}(A)$, then compute $Q = \pi(P)$ and solve for λ in $Q = [\lambda]P$; then $t_\ell \equiv \lambda + q/\lambda \pmod{\ell}$.

This approach is substantially faster than Schoof's algorithm for Elkies ℓ , because the degree of f_ϕ is only $(\ell - 1)/2$, whereas the degree of Ψ_ℓ is $(\ell^2 - 1)/2$; so each operation in $\mathcal{E}(A)$ costs much less than it would if we used Ψ_ℓ instead of f_ϕ . (In practice, it is also nice to be able to reduce the number of costly Frobenius computations, since we only need to compute $\pi(P)$ and not $\pi(\pi(P))$.)

The crucial step is computing f_ϕ given only \mathcal{E} and the corresponding root j_i of $\Phi_\ell(j(\mathcal{E}), X)$. We can do this using Elkies' algorithm, which is explained in [35,

§§7–8]. The total cost of computing t_ℓ is then $\tilde{O}(\log^3 q)$ \mathbb{F}_q -operations: that is, a whole factor of $\log q$ faster compared to Schoof’s algorithm.

Ideally, then, we should choose \mathcal{L} to only contain Elkies and volcanic primes: that is, non-Atkin primes. The usual naive heuristic on prime classes is to suppose that as $q \rightarrow \infty$, the number of Atkin and non-Atkin primes less than B for \mathcal{E}/\mathbb{F}_q is approximately equal when $B \sim \log q$; under this heuristic, taking \mathcal{L} to contain only non-Atkin primes, the SEA algorithm computes t in $\tilde{O}(\log^4 q)$ \mathbb{F}_q -operations.

While the heuristic holds on the average, assuming the GRH, Galbraith and Satoh have shown that it can fail for some curves [33, Appendix A]: there exist curves \mathcal{E}/\mathbb{F}_q such that if we try to compute t_ℓ using ℓ in the smallest possible set \mathcal{L} containing only non-Atkin primes, then \mathcal{L} must contain primes in $\Omega(\log^2 q)$.

Remark 1. It is important to note that Elkies’ technique applies only to primes ℓ where there exists a rational ℓ -isogeny: that is, only Elkies and volcanic primes. Atkin’s technique for restricting the possible values of t_ℓ applies to *all* primes—not only Atkin primes.

3.3 The genus 2 setting

Let \mathcal{C} be a genus-2 curve defined over \mathbb{F}_q (again, for q odd). We suppose that \mathcal{C} is defined by an equation of the form $y^2 = f(x)$, where f is squarefree of degree 5.³ The curve \mathcal{C} then has a unique point at infinity, which we denote ∞ .

3.3.1 The Jacobian

We write $J_{\mathcal{C}}$ for the Jacobian of \mathcal{C} . Our main algorithmic handle on $J_{\mathcal{C}}$ is Mumford’s model for hyperelliptic Jacobians, which represents the projective $J_{\mathcal{C}}$ as a disjoint union of three affine subsets. In this model, points of $J_{\mathcal{C}}$ correspond to pairs of polynomials $\langle a(x), b(x) \rangle$ where a is monic, $\deg b < \deg a \leq 2$, and $b^2 \equiv f \pmod{a}$ (we call $\langle a, b \rangle$ the *Mumford representation* of the Jacobian point). Mumford’s coordinates on the affine subsets of $J_{\mathcal{C}}$ are the coefficients of the polynomials a and b (and in particular, a point $\langle a, b \rangle$ of $J_{\mathcal{C}}$ is defined over \mathbb{F}_q if and only if a and b have coefficients in \mathbb{F}_q). The three affine subsets are

$$\begin{aligned} W_2 &:= \{ \langle a, b \rangle \in J_{\mathcal{C}} \mid \deg(a) = 2 \} && \text{ (“general” elements) ,} \\ W_1 &:= \{ \langle a, b \rangle \in J_{\mathcal{C}} \mid \deg(a) = 1 \} && \text{ (“special” elements) ,} \\ W_0 &:= \{ 0_{J_{\mathcal{C}}} = \langle 1, 0 \rangle \} && \text{ (the trivial element) ,} \end{aligned}$$

³ For full generality, we should also allow $\deg f = 6$; the curve \mathcal{C} then has two points at infinity. This substantially complicates the formulae without significantly modifying the algorithms or their asymptotic complexity, so we will not treat this case here.

and $J_{\mathcal{C}} = W_2 \sqcup W_1 \sqcup W_0$. The group law on $J_{\mathcal{C}}$ can be explicitly computed on Mumford representatives using Cantor's algorithm [2].

The point of $J_{\mathcal{C}}$ corresponding to a general divisor class $[(x_P, y_P) + (x_Q, y_Q) - 2\infty]$ on \mathcal{C} is represented by $\langle a, b \rangle$ where $a(x) = (x - x_P)(x - x_Q)$ and b is the linear polynomial such that $b(x_P) = y_P$ and $b(x_Q) = y_Q$. Special classes $[(x_P, y_P) - \infty]$ are represented by $\langle a, b \rangle = \langle x - x_P, y_P \rangle$, while $0_{J_{\mathcal{C}}} = [0]$ is represented by $\langle a, b \rangle = \langle 1, 0 \rangle$.

3.3.2 Frobenius and endomorphisms of $J_{\mathcal{C}}$

The characteristic polynomial χ_{π} of the Frobenius endomorphism π has the form

$$\chi_{\pi}(X) = X^4 - tX^3 + (2q + s)X^2 - tqX + q^2,$$

where s and t are integers satisfying the inequalities (cf. [32])

$$|s| < 4q, \quad |t| \leq 4\sqrt{q}, \quad t^2 > 4s, \quad s + 4q > 2|t|\sqrt{q}.$$

We have

$$\#J_{\mathcal{C}}(\mathbb{F}_q) = \chi_{\pi}(1) = 1 - t + 2q + s - tq + q^2,$$

as well as $\#\mathcal{C}(\mathbb{F}_q) = 1 - t + q$ and $\#\mathcal{C}(\mathbb{F}_{q^2}) = 1 - t^2 + 4q + 2s + q^2$. In genus 2, therefore, the point counting problem is to determine the integers s and t .

3.3.3 Real multiplication

We are interested in Jacobians $J_{\mathcal{C}}$ with real multiplication by a fixed order \mathcal{O} in a quadratic real field $F := \mathbb{Q}(\sqrt{\Delta})$; that is, such that there is an embedding $\iota: \mathcal{O} \rightarrow \text{End}(J_{\mathcal{C}})$. In this article, we will further restrict to the case where \mathcal{O} is the maximal order \mathcal{O}_F of F ; note that if \mathcal{O} is an order in F that is not locally maximal at a prime ℓ , then there exist no isogenies of degree ℓ that preserve the polarization (see Definition 2). These Jacobians can be constructed either from points in their moduli spaces (as in §3.4), or from a few known explicit families (as in §3.7).

The fixed field $\mathbb{Q}(\pi + \pi^{\dagger})$ of the Rosati involution on $\mathbb{Q}(\pi)$ is a real quadratic field, and $\mathbb{Z}[\pi + \pi^{\dagger}]$ is a suborder of \mathcal{O}_F . The characteristic polynomial of $\pi + \pi^{\dagger}$ is

$$\chi_{\pi + \pi^{\dagger}}(X) = (X^2 - tX + s)^2,$$

so determining $\chi_{\pi + \pi^{\dagger}}$ also solves the point counting problem for $J_{\mathcal{C}}$.

Later, we will be particularly interested in \mathcal{C} such that $J_{\mathcal{C}}$ has real multiplication by an order of small discriminant. While such curves are special, from a cryptographic perspective they are not "too special". From an arithmetic point of view, all curves (with ordinary simple Jacobians) over \mathbb{F}_q have real multiplication. Here, we simply

require that real multiplication to have small discriminant; the discriminant of the entire endomorphism ring of $J_{\mathcal{C}}$ can still be just as large as for a general choice of curve over the same field. From a geometric point of point view, the moduli of these \mathcal{C} live on two-dimensional Humbert surfaces inside the three-dimensional moduli space of genus-2 curves. In concrete terms, this means that when selecting random curves over a fixed \mathbb{F}_q , only $\sim 1/q$ of them have real multiplication by a fixed order; but if we restrict our choice to those curves then there are still $O(q^2)$ of them to choose from.

3.3.4 From Schoof to Pila

The Schoof–Pila algorithm deals with higher dimensions [34, 31]. Its input is a set of defining equations for a projective model of the abelian variety, and its group law. Jacobians of genus-2 curves are abelian varieties, and we can apply Pila’s algorithm to them using the defining equations computed by Flynn [5] or Grant [10]. However, the complexity of Pila’s algorithm is $O((\log q)^\Delta)$, where Δ (and the big-O constant) depends on the number of variables (i.e., the dimension of the ambient projective space) and the degree and number of the defining equations. Pila derives an upper bound for Δ in [31, §4], but when we evaluate this bound in the parameters of Flynn’s model for $J_{\mathcal{C}}$ (72 quadratic forms in 16 variables) we get a 30-bit Δ ; Grant’s model (13 quadratic and cubic forms in 9 variables) yields a 23-bit Δ .⁴ While these are only upper bounds, we are clearly in the realm of the impractical here.

3.3.5 The Gaudry–Schost approach

Pila’s algorithm requires a concrete (and necessarily complicated) nonsingular projective model for $J_{\mathcal{C}}$. The Gaudry–Schost algorithm applies essentially the same ideas to Mumford’s affine models for subsets of $J_{\mathcal{C}}$.

Our first problem is to find an analogue for $J_{\mathcal{C}}$ of the elliptic division polynomials Ψ_ℓ . Ultimately, we want an ideal $I_\ell = (F_0, \dots, F_r) \subset \mathbb{F}_q[A_1, A_0, B_1, B_0]$ such that $\langle a, b \rangle = \langle x^2 + a_1x + a_0, b_1x + b_0 \rangle$ is in $J_{\mathcal{C}}[\ell]$ if and only if (a_1, a_0, b_1, b_0) is in the variety of I_ℓ : that is,

$$[\ell]\langle x^2 + a_1x + a_0, b_1x + b_0 \rangle = 0 \iff F(a_1, a_0, b_1, b_0) = 0 \text{ for all } F \in I_\ell .$$

Then, the image of $\langle x^2 + A_1x + A_0, B_1x + B_0 \rangle$ in $J_{\mathcal{C}}(\mathbb{F}_q[A_1, A_0, B_1, B_0]/I_\ell)$ is an element of order ℓ that we can use for a Schoof-style computation of $\chi(T) \pmod{\ell}$.

The simplest approach here would be to take a general Mumford representative $\langle x^2 + A_1x + A_0, B_1x + B_0 \rangle$, compute $L = [\ell]\langle x^2 + A_1x + A_0, B_1x + B_0 \rangle$, and then equate coefficients in $L = 0_{J_{\mathcal{C}}}$ to derive the relations in I_ℓ . But we cannot do this,

⁴ With polynomial time estimates like these, who needs enemies?

because L is in $W_2(\mathbb{F}_q(A_1, A_0, B_1, B_0))$ (that is, its a -polynomial has degree 2, and its b -polynomial degree 1), while $0_{J_\ell} = \langle 1, 0 \rangle$ is in W_0 : these elements are not in the same affine subvariety, and cannot be directly compared or equated in this form.

Gaudry and Harley [6] neatly stepped around this problem by observing that any element of J_ℓ can be written as the difference of two elements of W_1 (which may be defined over a quadratic extension). They therefore start with $D = [(x_P, y_P) + (x_Q, y_Q) - 2\infty] = [(x_P, y_P) - (x_Q, -y_Q)]$ in J_ℓ , and find polynomial relations on $x_P, y_P, x_Q,$ and y_Q such that $[\ell]D = 0$ by computing $[\ell]\langle x - x_P, y_P \rangle$ and $[\ell]\langle x - x_Q, -y_Q \rangle$, and equating coefficients in $[\ell]\langle x - x_P, y_P \rangle = [\ell]\langle x - x_Q, -y_Q \rangle$. There is a quadratic level of redundancy in these relations, which is a direct result of the redundancy in the initial representation of D : the involution $(x_P, y_P) \leftrightarrow (x_Q, y_Q)$ fixes D .

Gaudry and Schost remove this redundancy by resymmetrizing the relations with respect to this involution, re-expressing them in terms of $A_1 = -(x_P + x_Q), A_0 = x_P x_Q, B_1 = (y_P - y_Q)/(x_P - x_Q),$ and $B_0 = (x_P y_Q - x_Q y_P)/(x_P - x_Q),$ and computing a triangular basis for the resulting *division ideal* I_ℓ . Their algorithm yields a triangular basis for I_ℓ , which facilitates fast reduction modulo I_ℓ .

Once we have I_ℓ , we can compute $t \pmod{\ell}$ and $s \pmod{\ell}$ as follows:

1. Construct the symbolic ℓ -torsion point

$$P := \langle x^2 + A_1 x + A_0, B_1 x + B_0 \rangle \in J_\ell(\mathbb{F}_q[A_1, A_0, B_1, B_0]/I_\ell);$$

2. Compute the points

$$\begin{aligned} Q_s &:= \pi^2(P), \\ Q_t &:= \pi(\pi^2(P) + [q \bmod \ell]\pi(P)), \\ R &:= \pi^4(P) + [2q \bmod \ell]\pi^2(P) + [q^2 \bmod \ell]P \end{aligned}$$

using Cantor arithmetic, with reduction of coefficients modulo I_ℓ ;

3. Search for $0 \leq s_\ell, t_\ell < \ell$ such that

$$[t_\ell]Q_t - [s_\ell]Q_s = R$$

(using, say, a two-dimensional baby-step giant-step algorithm).

The result is an algorithm that runs in time $\tilde{O}(\log^8 q)$. Of course, once t has been determined, we can simplify Steps (2) and (3) above to find s_ℓ more quickly for the remaining ℓ , but this does not change the asymptotic complexity. In practice, the algorithm has been used to construct cryptographically secure curves: Gaudry and Schost computed a generic genus-2 curve over $\mathbb{F}_{2^{127}-1}$ such that both the Jacobian and its quadratic twist have prime order [8]. Instances of the discrete logarithm problem in this Jacobian offer a claimed security level of roughly 128 bits, which is the current minimum for serious cryptosystems. This computation also represents the current record for point counting for general genus-2 curves.

The Gaudry–Schost computation illustrates not only the state-of-the-art of genus-2 point counting, but also the practical challenge involved in producing cryptographically strong genus-2 Jacobians. The Schoof-like point counting algorithm was only

applied using the prime powers 2^{17} , 3^9 , 5^4 , and 7^2 , and the primes 11 through 31. Combining the information given by these prime powers completely determines t , but not s ; but it still gives us enough modular information about s to be able to recover its precise value using Pollard’s kangaroo algorithm in a reasonable time (≈ 2 hours, in this case). The kangaroo algorithm is exponential, and would not be practical for computing this Jacobian order alone without the congruence data generated by the Schoof-like computations. Gaudry and Schost estimated the average cost of these calculations as one core-month (in 2008) per curve.

3.3.6 Point counting with efficiently computable RM

In [7], Gaudry, Kohel, and Smith described a number of improvements to the Gaudry–Schost algorithm that apply when $J_{\mathcal{C}}$ is equipped with an explicit and efficiently computable endomorphism ϕ generating a real quadratic subring of $\text{End}(J_{\mathcal{C}})$. When we say that ϕ is *explicit* we mean that we can compute the images under ϕ of divisor classes on $J_{\mathcal{C}}$, including symbolic Mumford representatives for generic divisor classes. When we say that ϕ is *efficiently computable*, we mean that these images can be computed for a cost comparable with a few group operations: that is, from an algorithmic point of view, we may view evaluation of ϕ as an elementary group operation like adding or doubling.

Suppose that $\mathbb{Z}[\pi + \pi^\dagger]$ is contained in $\mathbb{Z}[\phi]$ (this is reasonable, since in the examples we know, $\mathbb{Z}[\phi]$ is a maximal order), and let Δ be the discriminant of $\mathbb{Z}[\phi]$. Then $\pi + \pi^\dagger = m\phi + n$ for some m and n , which completely determine s and t : if the characteristic polynomial of ϕ is $(X^2 - t_\phi X + s_\phi)^2$, then $t = 2m + nt_\phi$ and $s = (t^2 - s_\phi^2 \Delta)/4$. It follows that m and n are both in $O(\sqrt{q})$.

We can compute m and n using a technique similar to Gaudry–Schost. Multiplying the relation $\pi + \pi^\dagger = m\phi + n$ through by π , we have $\pi^2 - (m\phi + n)\pi + q = 0$. Imitating Schoof’s algorithm, we can compute $m_\ell := m \pmod{\ell}$ and $n_\ell := n \pmod{\ell}$ by taking a generic element D of $J_{\mathcal{C}}[\ell]$ (as in Gaudry–Schost), computing $(\pi^2 + q)(D)$, $\pi(D)$, and $\phi\pi(D)$ (using two applications of π), and then solving for m_ℓ and n_ℓ .

We can do even better by exploiting split primes in $\mathbb{Z}[\phi]$. If $\ell = \iota_1 \iota_2$ is split, then the ℓ -torsion decomposes as $J_{\mathcal{C}}[\iota_1] \oplus J_{\mathcal{C}}[\iota_2]$, and once we have found a short generator (or generators) for ι_i we can take D to be an element of $J_{\mathcal{C}}[\iota_i]$ instead of $J_{\mathcal{C}}[\ell]$. Such generators can be found with coefficients in $O(\sqrt{\ell})$; the result is that we work modulo a much smaller ideal, of degree $O(\ell^2)$ rather than $O(\ell^4)$.

But going further, $\pi + \pi^\dagger$ acts as a scalar on $J_{\mathcal{C}}[\iota_i]$, and so we can compute its eigenvalue to determine m_ℓ and n_ℓ . The total cost of computing m_ℓ and n_ℓ , and hence t_ℓ and s_ℓ , is then $\tilde{O}(\log^5 q)$ [7, Theorem 1], a substantial improvement on Gaudry–Schost’s $\tilde{O}(\log^8 q)$.

The computation resembles what we would do for an Elkies prime in the elliptic case, except that there is no need for modular polynomials to compute the prime type, or for an analogue of Elkies’ algorithm: we know in advance which primes split in $\mathbb{Z}[\phi]$, and we can compute the kernel using the decomposition. But if we did have an

analogue of Elkies’ algorithm, then we could further reduce the complexity by further decomposing some of the $J_{\mathcal{C}}[l_i]$ into cyclic factors, and thus working modulo ideals of degree $O(\ell)$. If we have an analogue of Atkin’s algorithm, then we can restrict the possible values of m_ℓ and n_ℓ ; this would not change the asymptotic complexity of the algorithm, but it could have a significant practical impact.

3.3.7 Generalizing Elkies’ and Atkin’s improvements to genus 2

Ultimately, we would like to generalize the SEA algorithm to genus 2. The first requirement is a genus-2 analogue of elliptic modular polynomials; so assume for the moment that we have a modular ideal relating suitable invariants of genus-2 curves.

To generalize Elkies’ improvements to genus 2, we need an analogue of Elkies’ algorithm: that is, an algorithm which, given two general moduli points corresponding to isogenous Jacobians, constructs defining polynomials for (the kernel of) the isogeny. The most convenient such presentation would be as an ideal cutting out the intersection of the kernel with W_2 , since then the Gaudry–Schost approach could be adapted without too much difficulty (at least in theory). Unfortunately, at present, no such algorithm is known.

In contrast, Atkin’s techniques for elliptic curves require only the factorization of (specializations of) elliptic modular polynomials; we deduce possible congruences on the trace from the degrees of the factors. It is clear how we should generalize Atkin’s techniques to genus 2: we should deduce possible congruences on s and t from the degrees of primary components of specialized modular ideals.

The following sections make this concrete. In §3.4, we define the appropriate analogues of the elliptic j -invariant for genus-2 curves with real multiplication. We can then define real-multiplication analogues of the elliptic modular polynomials in §3.4.2, before investigating their factorization in §3.5.

3.3.8 μ -isogenies

Before defining any generalized invariants or modular polynomials, we must define an appropriate class of isogenies in genus 2: that is, isogenies that are compatible with the real multiplication structure. (This is not an issue for elliptic curves, because the elliptic analogue of the real endomorphism subring is just \mathbb{Z} —and everything is compatible with integer multiplications.)

Definition 2. Let $(\mathcal{A}, \xi, \iota)$ and $(\mathcal{A}', \xi', \iota')$ be triples encoding principally polarized abelian surfaces with real multiplication by \mathcal{O}_F . Here $\xi: \mathcal{A} \rightarrow \mathcal{A}^\vee$ and $\xi': \mathcal{A}' \rightarrow (\mathcal{A}')^\vee$ are principal polarizations, and $\iota: \mathcal{O}_F \hookrightarrow \text{End}(\mathcal{A})$ and $\iota': \mathcal{O}_F \hookrightarrow \text{End}(\mathcal{A}')$ are embeddings that are stable under the Rosati involution. If μ is a totally positive element of F , then a μ -isogeny $(\mathcal{A}, \xi, \iota) \rightarrow (\mathcal{A}', \xi', \iota')$ is an isogeny $f: \mathcal{A} \rightarrow \mathcal{A}'$ such that the diagrams

$$\begin{array}{ccc}
\mathcal{A} & \xrightarrow{\iota(\mu)} \mathcal{A} & \xrightarrow{\xi} \mathcal{A}^\vee \\
\downarrow f & & \uparrow f^\vee \\
\mathcal{A}' & \xrightarrow{\xi'} & (\mathcal{A}')^\vee
\end{array}
\quad \text{and} \quad
\begin{array}{ccc}
F & \xrightarrow{\iota} & \text{End}(\mathcal{A}) \otimes \mathbb{Q} \\
\searrow \iota' & & \downarrow \phi \\
& & \text{End}(\mathcal{A}') \otimes \mathbb{Q}
\end{array}$$

commute, where ϕ is the map induced by f on endomorphism algebras.

If $f: (\mathcal{A}, \xi, \iota) \rightarrow (\mathcal{A}', \xi', \iota')$ is a μ -isogeny, then the polarization ξ' pulls back via f to $\xi \circ \iota(\mu)$. For comparison, an elliptic ℓ -isogeny is an $f: \mathcal{E} \rightarrow \mathcal{E}'$ such that the canonical polarization on \mathcal{E}' pulls back via f to ℓ times the polarization on \mathcal{E} (in more concrete terms: the identity point $0_{\mathcal{E}'}$ on \mathcal{E}' pulls back via f to a divisor on \mathcal{E} equivalent to $\ell \cdot 0_{\mathcal{E}}$).

3.4 Invariants

Elliptic modular polynomials relate isogenous elliptic curves in terms of their j -invariants; their genus-2 analogues must relate invariants of genus-2 Jacobians. This section describes and relates the various invariants that we will need. Since we are dealing with classical constructions in this section, we work over a field $k \subseteq \mathbb{C}$. However, the resulting algebraic expressions carry over to the case where $k = \mathbb{F}_q$ (at least for large enough p). All of the results in this section are well-known, and are shown here for completeness and easy reference; we refer the reader to [19], [21], [22], and [24] for further detail.

3.4.1 Invariants for RM abelian surfaces

Let F be a real quadratic field with ring of integers \mathcal{O}_F . We need RM analogues of the elliptic j -invariant and elliptic modular polynomials for μ -isogenies of abelian surfaces with RM by \mathcal{O}_F . Our first step is to define appropriate replacements for the j -invariant that classify our triples $(\mathcal{A}, \xi, \iota)$ up to isomorphism. Instead of a single j -invariant, we will have a triple (J_1, J_2, J_3) of *RM invariants*, which are functions on the corresponding Hilbert modular surface.

The invariants (J_1, J_2, J_3) are constructed as follows. For a field k , we consider the coarse moduli space $\mathcal{H}_F(k)$ of triples $(\mathcal{A}, \xi, \iota)$ (where as before, \mathcal{A}/k is an abelian variety with a principal polarization $\xi: \mathcal{A} \rightarrow \mathcal{A}^\vee$ and an embedding $\iota: \mathcal{O}_F \hookrightarrow \text{End}_k(\mathcal{A})$ stable under the Rosati involution). Then $\mathcal{H}_F(k)$ is coarsely represented by the Hilbert modular space $\text{SL}_2(\mathcal{O}_F \oplus \mathcal{O}_F) \backslash (F \otimes \mathbb{H})$ (see [9]), where $F \otimes \mathbb{H} := \{\tau \in F \otimes \mathbb{C} : \Im(\tau) > 0\}$ and for any fractional ideal \mathfrak{f} of F ,

$$\text{SL}_2(\mathcal{O}_F \oplus \mathfrak{f}) := \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}_2(F) : a, d \in \mathcal{O}_F, b \in \mathfrak{f}, c \in \mathfrak{f}^{-1} \right\}$$

acts on $F \otimes \mathbb{H}$ by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau = \frac{a\tau + b}{c\tau + d}.$$

Proposition 2. *Let V be the Baily–Borel compactification of $SL_2(\mathcal{O}_F) \backslash (F \otimes \mathbb{H})$, and $\mathbb{C}(V)$ the function field of V . There exist rational functions J_1, J_2 , and J_3 on V such that*

$$\mathbb{C}(V) = \mathbb{C}(J_1, J_2, J_3).$$

Proof. The transcendence degree of $\mathbb{C}(V)$ over \mathbb{C} is 2, so there exist 2 algebraically independent functions J_1, J_2 in $\mathbb{C}(V)$. Furthermore, $\mathbb{C}(V)$ is a finite separable field extension of $\mathbb{C}(J_1, J_2)$, so it is generated by at most one further element, J_3 .

Definition 3. Fixing a choice of rational functions J_1, J_2 , and J_3 as in Proposition 2, we call (J_1, J_2, J_3) the *RM invariants* for F .

3.4.2 Hilbert modular polynomials for RM abelian surfaces

We are now ready to define modular polynomials for abelian surfaces with RM structure. For elliptic curves we have a single j -invariant, and we can relate ℓ -isogenous j -invariants using a single bivariate polynomial $\Phi_\ell(X, Y)$. For our abelian surfaces, we have a tuple of three invariants (J_1, J_2, J_3) , and to relate μ -isogenous tuples of invariants we need a *modular ideal* of polynomials in $\mathbb{Q}[X_1, X_2, X_3, Y_1, Y_2, Y_3]$, such that when we specialize the first three variables in the (J_1, J_2, J_3) corresponding to the isomorphism class of some triple (A, ξ, ι) , the result is an ideal cutting out the moduli points (J'_1, J'_2, J'_3) for triples (A', ξ', ι') that are μ -isogenous to (A, ξ, ι) .

The *Hilbert modular polynomials* below represent a particularly convenient basis for this ideal. We refer the reader to [24, Chapter 2] for theoretical details and proofs, as well as algorithms for computing the polynomials. Alternatively, Milio's algorithm can be used to compute Hilbert modular polynomials $\Phi_\ell(X, \mathfrak{J}_1, \mathfrak{J}_2)$ and $\Psi_\ell(X, \mathfrak{J}_1, \mathfrak{J}_2)$, in time $O(d_T d_{\mathfrak{J}_2}) \tilde{O}(\ell N) + 4(\ell + 1) \tilde{O}(d_T d_{\mathfrak{J}_2} N) \subseteq \tilde{O}(d_T d_{\mathfrak{J}_2} \ell N)$ [27, Theorem 5.4.4], where N is the precision and $d_T, d_{\mathfrak{J}_2}$ are degrees involved in the computation, see [27, §5.4].

Definition 4. The *Hilbert modular polynomials*

$$\begin{aligned} G_\mu(X_1, X_2, X_3, Y_1), \\ H_{\mu,2}(X_1, X_2, X_3, Y_1, Y_2) &= H_{\mu,2}^{(1)}(X_1, X_2, X_3, Y_1) Y_2 + H_{\mu,2}^{(0)}(X_1, X_2, X_3, Y_1), \\ H_{\mu,3}(X_1, X_2, X_3, Y_1, Y_3) &= H_{\mu,3}^{(1)}(X_1, X_2, X_3, Y_1) Y_3 + H_{\mu,3}^{(0)}(X_1, X_2, X_3, Y_1) \end{aligned}$$

in $\mathbb{Q}[X_1, X_2, X_3, Y_1, Y_2, Y_3]$ are defined such that for all triples $(\mathcal{A}, \xi, \iota)$ and $(\mathcal{A}', \xi', \iota')$ representing points τ and τ' in a certain Zariski-open subset⁵ of the Baily–Borel

⁵ See [24, Chapter 2, Section 2] for details on this subset. For point counting over large finite fields, it is enough to note that since the subset is Zariski open, randomly sampled Jacobians with real multiplication by \mathcal{O}_F have their RM invariants in this subset with overwhelming probability.

compactification of $\mathrm{SL}_2(\mathcal{O}_F \oplus \mathfrak{f}) \backslash (F \otimes \mathbb{H})$, there exists a μ -isogeny $f: (\mathcal{A}, \xi, \iota) \rightarrow (\mathcal{A}', \xi', \iota')$ if and only if

$$\begin{aligned} G_\mu(J_1(\tau), J_2(\tau), J_3(\tau), J_1(\tau')) &= 0, \\ H_{\mu,2}(J_1(\tau), J_2(\tau), J_3(\tau), J_1(\tau'), J_2(\tau')) &= 0, \\ H_{\mu,3}(J_1(\tau), J_2(\tau), J_3(\tau), J_1(\tau'), J_3(\tau')) &= 0. \end{aligned}$$

The special form of G_μ , $H_{2,\mu}$, and $H_{3,\mu}$ are very convenient for computations. If (J_1, J_2, J_3) is a fixed moduli point, then each root α of $G(J_1, J_2, J_3, x)$ corresponds to a unique μ -isogenous moduli point

$$(J'_1, J'_2, J'_3) = \left(\alpha, -\frac{H_{\mu,2}^{(0)}(J_1, J_2, J_3, \alpha)}{H_{\mu,2}^{(1)}(J_1, J_2, J_3, \alpha)}, -\frac{H_{\mu,3}^{(0)}(J_1, J_2, J_3, \alpha)}{H_{\mu,3}^{(1)}(J_1, J_2, J_3, \alpha)} \right).$$

We observe that the action of Galois on the set of μ -isogenies from an RM abelian variety representing (J_1, J_2, J_3) is completely described by the action of Galois on the roots of $G_\mu(J_1, J_2, J_3, x)$; in particular, over \mathbb{F}_q , rational cycles of μ -isogenies under Frobenius correspond to irreducible factors of $G_\mu(J_1, J_2, J_3, x)$. From the point of view of Atkin generalizations, therefore, we only really need G_μ to replace Φ_ℓ .

3.4.3 Invariants for curves and abelian surfaces

We need to relate the RM invariants (J_1, J_2, J_3) to the invariants for plain old principally polarized abelian surfaces, and in particular Jacobians of genus 2 curves without any special RM structure. The moduli space \mathcal{A}_2 of principally polarized abelian surfaces is coarsely represented by the Siegel modular space $\mathrm{Sp}_2(\mathbb{Z}) \backslash \mathbb{H}_2$, where

$$\mathbb{H}_2 := \left\{ \tau = \begin{pmatrix} \tau_1 & \tau_2 \\ \tau_2 & \tau_3 \end{pmatrix} \in \mathrm{Sym}_2(\mathbb{C}) : \Im(\tau) > 0 \right\},$$

and the symplectic group

$$\mathrm{Sp}_2(\mathbb{Z}) = \left\{ g \in \mathrm{GL}_4(\mathbb{Z}) : g \begin{pmatrix} 0 & I_2 \\ -I_2 & 0 \end{pmatrix} g^t = \begin{pmatrix} 0 & I_2 \\ -I_2 & 0 \end{pmatrix} \right\}$$

acts on \mathbb{H}_2 via

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau = \frac{a\tau + b}{c\tau + d}.$$

Every rational function on $\mathrm{Sp}_2(\mathbb{Z}) \backslash \mathbb{H}_2$ is a quotient of elements of the graded ring of holomorphic Siegel modular forms for $\mathrm{Sp}_2(\mathbb{Z})$. Igusa proved in [15] that this ring is generated by ψ_4 , ψ_6 , χ_{10} , and χ_{12} , where

$$\psi_k(\tau) = \sum_{\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in P \backslash \mathrm{Sp}_2(\mathbb{Z})} \det(c\tau + d)^{-k}$$

is the normalized Eisenstein series of weight k for even integers $k \geq 4$ (here P is the standard Siegel parabolic subgroup of $\mathrm{Sp}_2(\mathbb{Z})$), and

$$\begin{aligned} \chi_{10} &= -2^{-12} \cdot 3^{-5} \cdot 5^{-2} \cdot 7^{-1} \cdot 43867(\psi_4\psi_6 - \psi_{10}), \\ \chi_{12} &= 2^{-13} \cdot 3^{-7} \cdot 5^{-3} \cdot 7^{-2} \cdot 337^{-1} \cdot 131 \cdot 593(3^2 \cdot 7^2 \psi_4^3 + 2 \cdot 5^3 \psi_6^2 - 691 \psi_{12}) \end{aligned}$$

are Siegel modular cusp forms of weight 10 and 12 respectively.

Curves of genus 2 are typically classified up to isomorphism by their Igusa invariants (j_1, j_2, j_3) , or by their Igusa–Clebsch invariants (A, B, C, D) . Since the map $\mathcal{C} \mapsto J_{\mathcal{C}}$ is an open immersion of the (coarse) moduli space of genus-2 curves \mathcal{M}_2 into \mathcal{A}_2 , the Igusa invariants j_i can be written as rational functions of $\psi_4, \psi_6, \chi_{10}$ and χ_{12} as follows [16]:

$$\begin{aligned} j_1(\tau) &= 2 \cdot 3^5 \cdot \chi_{12}^5 \chi_{10}^{-6}, \\ j_2(\tau) &= 2^{-3} \cdot 3^3 \cdot \psi_4 \chi_{12}^3 \chi_{10}^{-4}, \\ j_3(\tau) &= 2^{-5} \cdot 3 \cdot (\psi_6 \chi_{12}^2 \chi_{10}^{-3} + 2^2 \cdot 3 \cdot \psi_4 \chi_{12}^3 \chi_{10}^{-4}). \end{aligned}$$

Here $j_i(\tau) = j_i(\mathcal{C})$ if there is a genus 2 curve \mathcal{C}/\mathbb{C} such that $J_{\mathcal{C}}$ is isomorphic to the abelian surface $\mathbb{C}^2/(\mathbb{Z}^2\tau + \mathbb{Z}^2)$. If there is no such \mathcal{C} , which happens exactly when $\chi_{10}(\tau) = 0$, then $j_i(\tau)$ is not well-defined. The Igusa–Clebsch invariants are related to the Siegel modular forms by

$$(\psi_4, \psi_6, \chi_{10}, \chi_{12}) = (2^{-2}B, 2^{-3}(AB - 3C), -2^{-14}C, 2^{-17}3^{-1}AD). \quad (3.5)$$

3.4.4 Pulling back curve invariants to RM invariants

The natural maps $\mathbb{H}^2 \rightarrow \mathbb{H}_2$, $\mathrm{SL}_2(F) \rightarrow \mathrm{Sp}_2(\mathbb{Q})$, and $(\mathcal{O}_F/2\mathcal{O}_F)^2 \rightarrow (\mathbb{Z}/2\mathbb{Z})^4$ induce an embedding

$$\phi: \mathcal{H}_F(k) \hookrightarrow \mathcal{A}_2(k),$$

which we can use to pull back Igusa invariants to RM invariants, thus expressing the j_i in terms of the J_i . We will see detailed formulæ for this pullback for $F = \mathbb{Q}(\sqrt{5})$ in Proposition 7.

This pullback from curves and their invariants to RM invariants is essential for our computations: after all, in point counting one usually starts from a curve. In our applications, we are given the equation of a curve \mathcal{C}/\mathbb{F}_q drawn from a family of curves with known RM by \mathcal{O}_F . Having computed the Igusa or Igusa–Clebsch invariants of \mathcal{C} , we can pull them back to RM invariants (J_1, J_2, J_3) . This pullback is possible, because \mathcal{C} was chosen from an appropriate family, but choosing a preimage (J_1, J_2, J_3) implicitly involves choosing one of the two embeddings of \mathcal{O}_F

into $\text{End}(J_{\mathcal{C}})$. This choice cannot always be made over the ground field: a point in $\mathcal{A}_2(k)$ may not pull back to a pair of points in $\mathcal{H}_F(k)$, but rather a conjugate pair of points over a quadratic extension of k . Proposition 8 makes this subtlety explicit in the case $F = \mathbb{Q}(\sqrt{5})$.

3.5 Atkin theorems in genus 2

We are now ready to state some Atkin-style results for μ -isogenies in genus 2.

Let $(\mathcal{A}, \xi, \iota)$ be a triple describing a vanilla abelian surface over \mathbb{F}_q with real multiplication by \mathcal{O}_F , and let μ be a totally positive element of \mathcal{O}_F of norm ℓ . Then $\iota(\mu)$ is an endomorphism of degree ℓ^2 , and we have a subgroup⁶

$$\mathcal{A}[\mu] := \ker(\iota(\mu)) \subset \mathcal{A}[\ell].$$

If $(\bar{\mu}) \neq (\mu)$ (that is, $(\ell) \neq (\mu^2)$), then we have a decomposition $\mathcal{A}[\ell] = \mathcal{A}[\mu] \oplus \mathcal{A}[\bar{\mu}]$. The one-dimensional subspaces of $\mathcal{A}[\mu]$ are the kernels of μ -isogenies.

In §3.2 we used the elliptic modular polynomial Φ_ℓ to study the structure of $\mathcal{E}[\ell]$. Here, we will use the Hilbert modular polynomial G_μ to study the structure of $\mathcal{A}[\mu]$. The propositions of this section are generalizations for curves of genus 2 to Schoof's Propositions 6.1, 6.2 and 6.3 for elliptic curves in [35].

3.5.1 Roots of G_μ and the order of Frobenius

Our first result relates the order of Frobenius acting on $\mathbb{P}(\mathcal{A}[\mu])$ to the extensions of \mathbb{F}_q generated by roots of specialized Hilbert modular polynomials.

Proposition 3. *Let \mathcal{A}/\mathbb{F}_q be a vanilla abelian surface with RM by \mathcal{O}_F and RM invariants (J_1, J_2, J_3) in \mathbb{F}_q^3 , and with Frobenius endomorphism π . Let μ be a totally positive element of \mathcal{O}_F of prime norm $\ell = \mu\bar{\mu}$.*

1. *The polynomial $G_\mu(J_1, J_2, J_3, x)$ has a zero \tilde{J}_1 in \mathbb{F}_{q^ℓ} if and only if the kernel of the corresponding μ -isogeny $\mathcal{A} \rightarrow \tilde{\mathcal{A}}$ is a 1-dimensional eigenspace of π^ℓ in $\mathcal{A}[\mu]$.*
2. *The polynomial $G_\mu(J_1, J_2, J_3, x)$ splits completely in $\mathbb{F}_{q^\ell}[x]$ if and only if π^ℓ acts as a scalar matrix on $\mathcal{A}[\mu]$.*

Proof. The proof follows that of [35, Proposition 6.1] (stated as Lemma 1 here).

For (1): Let $f: \mathcal{A} \rightarrow \tilde{\mathcal{A}}$ be a μ -isogeny with kernel S , and let $(\tilde{J}_1, \tilde{J}_2, \tilde{J}_3)$ be the RM invariants of $\tilde{\mathcal{A}}$. If S is an eigenspace of π^ℓ , then the quotient $\mathcal{A} \rightarrow \mathcal{A}/S$ is defined over \mathbb{F}_{q^ℓ} . The Igusa invariants of \mathcal{A}/S are therefore all in \mathbb{F}_{q^ℓ} , and since \mathcal{A}/S is isomorphic to $\tilde{\mathcal{A}}$ as a principally polarized abelian surface, the Igusa invariants

⁶ We emphasize that the subgroup $\mathcal{A}[\mu]$ depends on ι , but we have chosen to write $\mathcal{A}[\mu]$ instead of the more cumbersome $\mathcal{A}[\iota(\mu)]$.

of $\tilde{\mathcal{A}}$ are all in \mathbb{F}_{q^e} . To conclude that \tilde{J}_1 is in \mathbb{F}_{q^e} , we need to show that the injection $\tilde{\iota}: \mathcal{O}_F \hookrightarrow \text{End}(\tilde{\mathcal{A}})$ is defined over \mathbb{F}_{q^e} ; but this follows from the commutativity of the second diagram in Definition 2.

Conversely: suppose $G_\mu(J_1, J_2, J_3, \tilde{J}_1) = 0$ for some \tilde{J}_1 in \mathbb{F}_{q^e} . Then the fact that each of the $H_{\mu,i}$ is a linear polynomial in Y_i with coefficients in $\mathbb{F}_q[J_1, J_2, J_3, \tilde{J}_1] = \mathbb{F}_{q^e}$ shows that there exist \tilde{J}_2 and \tilde{J}_3 in \mathbb{F}_{q^e} such that $(\tilde{J}_1, \tilde{J}_2, \tilde{J}_3)$ are the RM invariants of a triple $(\tilde{\mathcal{A}}, \tilde{\xi}, \tilde{\iota})$ that is μ -isogenous to $(\mathcal{A}, \xi, \iota)$. This means that there is an $\overline{\mathbb{F}}_q$ -isomorphism $(\tilde{\mathcal{A}}, \tilde{\xi}, \tilde{\iota}) \rightarrow (\mathcal{A}', \xi', \iota')$ where $(\mathcal{A}', \xi', \iota')$ is defined over \mathbb{F}_{q^e} . Let $f: \mathcal{A} \rightarrow \mathcal{A}'$ be the composite μ -isogeny. Its kernel S is a one-dimensional subspace of $\mathcal{A}[\ell]$. It remains to show that S is an eigenspace of π^e ; this is the case if and only if f is defined over \mathbb{F}_{q^e} . The \mathbb{Z} -module $\text{Hom}_{\overline{\mathbb{F}}_q}(\mathcal{A}, \mathcal{A}')$ is free of rank 4 (because \mathcal{A} is vanilla); and its submodule $\text{Hom}_{\mathbb{F}_{q^e}}(\mathcal{A}, \mathcal{A}')$ of \mathbb{F}_{q^e} -isogenies is either 0 or equal to $\text{Hom}_{\overline{\mathbb{F}}_q}(\mathcal{A}, \mathcal{A}')$. Hence, f is defined over \mathbb{F}_{q^e} if $\text{Hom}_{\mathbb{F}_{q^e}}(\mathcal{A}, \mathcal{A}') \neq 0$; and $\text{Hom}_{\mathbb{F}_{q^e}}(\mathcal{A}, \mathcal{A}') \neq 0$ if and only if the Frobenius endomorphisms of $\mathcal{A}/\mathbb{F}_{q^e}$ and \mathcal{A}' have the same characteristic polynomial.

Since \mathcal{A} is vanilla, and \mathcal{A}' is $\overline{\mathbb{F}}_q$ -isogenous to \mathcal{A} , we have $\text{End}_{\overline{\mathbb{F}}_q}(\mathcal{A}') \otimes \mathbb{Q} \cong \text{End}_{\overline{\mathbb{F}}_q}(\mathcal{A}) \otimes \mathbb{Q} \cong K$ for some quartic CM-field K . So let ψ and ψ' be the images in K of the Frobenius endomorphisms of $\mathcal{A}/\mathbb{F}_{q^e}$ and \mathcal{A}' , respectively (note that $\psi = \pi^e$). Now up to complex conjugation, we have $\psi^s = (\psi')^s$ in K for some $s > 0$. If $\psi = \psi'$, then \mathcal{A} and \mathcal{A}' are \mathbb{F}_{q^e} -isogenous, and we are done. If $\psi = -\psi'$, then we replace $(\mathcal{A}', \xi', \iota')$ by its quadratic twist; and then \mathcal{A} and \mathcal{A}' are \mathbb{F}_{q^e} -isogenous. Otherwise, if $\psi \neq \pm\psi'$, then ψ/ψ' must be a root of unity of order at least 3 in K , which is impossible because \mathcal{A} is vanilla. Hence $\psi = \psi'$, so ψ and ψ' have the same characteristic polynomial, and therefore f is defined over \mathbb{F}_{q^e} .

For (2): If all of the zeroes of $G_\mu(J_1, J_2, J_3, x)$ are contained in \mathbb{F}_{q^e} , then all of the 1-dimensional subspaces of $\mathcal{A}[\mu]$ are eigenspaces of π^e by Part (1). This implies that π^e acts as a scalar matrix on $\mathcal{A}[\mu]$.

Remark 2. As an example of what can go wrong if the vanilla condition is dropped, consider the curve

$$\mathcal{C}: y^2 = x^5 + 1.$$

The Jacobian $J_{\mathcal{C}}$ of this curve has complex multiplication by $\mathbb{Q}(\zeta_5)$, so it is not vanilla. While $J_{\mathcal{C}}$ has real multiplication by the maximal order of $\mathbb{Q}(\sqrt{5})$, the Siegel modular form ψ_4 is zero for this curve. Proposition 8 below gives explicit formulæ for J_1 , J_2 , and J_3^2 for Jacobians with maximal real multiplication by $\mathbb{Q}(\sqrt{5})$; and when we look at those formulæ, we see that J_1 is not well-defined when $\psi_4 = 0$.

3.5.2 The factorization of G_μ

The Frobenius endomorphism π of \mathcal{A} commutes with $\iota(\mu)$ (since \mathcal{A} is vanilla), so it restricts to an endomorphism of $\mathcal{A}[\mu]$.

Lemma 2. Let \mathcal{A}/\mathbb{F}_q be a vanilla abelian surface with Frobenius endomorphism π , and let ℓ be an odd prime.

1. If ℓ splits in $\mathbb{Z}[\pi + \pi^\dagger]$ (or equivalently, if $t^2 - 4s$ is a square in \mathbb{F}_ℓ), then $\chi_\pi(T) \equiv (T^2 - uT + q)(T^2 - u'T + q) \pmod{\ell}$ for some u and u' in $\mathbb{Z}/\ell\mathbb{Z}$.
2. If ℓ is ramified in $\mathbb{Z}[\pi + \pi^\dagger]$ (or equivalently, if ℓ divides $t^2 - 4s$), then $\chi_\pi(T) \equiv (T^2 - uT + q)^2 \pmod{\ell}$ where $u = t/2$ in $\mathbb{Z}/\ell\mathbb{Z}$.
3. If ℓ is inert in $\mathbb{Z}[\pi + \pi^\dagger]$ (or equivalently, if $t^2 - 4s$ is a square in \mathbb{F}_ℓ), then $\chi_\pi(T) \not\equiv (T^2 - uT + q)(T^2 - u'T + q) \pmod{\ell}$ for any $u, u' \in \mathbb{Z}/\ell\mathbb{Z}$.

Proof. This is a direct consequence of [20, Chapter 1, Proposition 25].

Lemma 3. Let $(\mathcal{A}, \xi, \iota)$ be a triple describing a vanilla abelian surface over \mathbb{F}_q with real multiplication by \mathcal{O}_F , and let μ be a totally positive element of \mathcal{O}_F of prime norm $\mu\bar{\mu} = \ell$. The restriction of the Frobenius endomorphism π to $\mathcal{A}[\mu]$ has characteristic polynomial

$$\chi_{\pi, \mu}(T) \equiv T^2 - uT + q \pmod{\ell} \quad \text{for some } u \in \mathbb{Z}/\ell\mathbb{Z}.$$

Proof. By definition, $\ell = \mu\bar{\mu}$ splits in \mathcal{O}_F , so it either splits or ramifies in the suborder $\mathbb{Z}[\pi + \pi^\dagger] \subseteq \mathcal{O}_F$; we are therefore in Case (1) or (2) of Lemma 2. In particular, both π and π^\dagger restrict to endomorphisms of $\mathcal{A}[\mu]$, and they have the same eigenvalues λ and q/λ ; so the characteristic polynomial of π is $T^2 - (\lambda + q/\lambda)T + q$. The result follows with $u = \lambda + q/\lambda$.

Proposition 4 uses the factorization of the modular polynomial G_μ , specialized at the RM invariants of \mathcal{A} , to derive information $\chi_{\pi, \mu}(T) \pmod{\ell}$.

Proposition 4. Let $(\mathcal{A}, \xi, \iota)$ be a triple describing a vanilla abelian surface over \mathbb{F}_q with real multiplication by \mathcal{O}_F and with RM invariants (J_1, J_2, J_3) , and let μ be a totally positive element of \mathcal{O}_F of prime norm $\mu\bar{\mu} = \ell$. Let π be the Frobenius endomorphism of \mathcal{A} , with $\chi_{\pi, \mu}(T) = T^2 - uT + q$ the characteristic polynomial of the restriction of π to $\mathcal{A}[\mu]$, and let e be the order of π in $\text{Aut}(\mathbb{P}(\mathcal{A}[\mu])) \cong \text{PGL}_2(\mathbb{F}_\ell)$.

The polynomial $G_\mu(J_1, J_2, J_3, x)$ has degree $\ell + 1$ in $\mathbb{F}_q[x]$, and its factorization type is as follows:

1. If $u^2 - 4q$ is not a square in \mathbb{F}_ℓ , then $e > 1$ and the factorization type is

$$(e, \dots, e) \quad \text{where } e \mid \ell + 1.$$

2. If $u^2 - 4q$ is a nonzero square in \mathbb{F}_ℓ , then the factorization type is

$$(1, 1, e, \dots, e) \quad \text{where } e \mid \ell - 1.$$

3. If $u^2 - 4q = 0$ in \mathbb{F}_ℓ , then the factorization type is

$$(1, e) \quad \text{where } e = \ell.$$

Proof. By Lemma 2, the endomorphism π acts on $\mathcal{A}[\mu]$ as a 2×2 matrix in $\mathrm{GL}_2(\mathbb{F}_\ell)$ with characteristic polynomial $T^2 - uT + q = 0$. If the matrix has two conjugate eigenvalues λ_1, λ_2 in \mathbb{F}_{ℓ^2} , then we are in Case (1): there are no 1-dimensional eigenspaces of π in $\mathcal{A}[\mu]$, and all irreducible factors of $G_\mu(J_1, J_2, J_3, x)$ have degree e , where e is the smallest exponent such that λ_i^e is in \mathbb{F}_ℓ .

If the matrix has two eigenvalues in \mathbb{F}_ℓ and is diagonalizable, then the discriminant $t^2 - 4s$ is a square modulo ℓ : we are in Case (2). This time $\mathcal{A}[\mu]$ is the direct product of two 1-dimensional eigenspaces, which account for two linear factors of $G_\mu(J_1, J_2, J_3, x)$. The remaining factors have degree e , where e is the smallest positive integer such that π^e acts as a scalar matrix.

If the matrix has a double eigenvalue and is not diagonalizable, then we are in Case (3): there is only one 1-dimensional eigenspace, and the matrix of π^ℓ is scalar.

3.5.3 The characteristic polynomial of Frobenius

Now that we can compute the order of Frobenius, we want to use this to derive information on the characteristic polynomial. Proposition 5 generalizes Proposition 1 to genus 2.

Proposition 5. *Let $(\mathcal{A}/\mathbb{F}_q, \xi, \iota)$ be a triple describing a vanilla abelian surface with real multiplication by \mathcal{O}_F , and let μ be a totally positive element of prime norm $\ell = \mu\bar{\mu} \notin \{2, p\}$. Let π be the Frobenius endomorphism of \mathcal{A} , and $\chi_{\pi, \mu}(T) = T^2 - uT + q$ the characteristic polynomial of its restriction to $\mathcal{A}[\mu]$. If e is the order of the image of π in $\mathrm{Aut}(\mathbb{P}(\mathcal{A}[\mu])) \cong \mathrm{PGL}_2(\mathbb{F}_\ell)$, then*

$$u^2 = \eta_e q \quad \text{in } \mathbb{F}_\ell,$$

$$\text{where } \eta_e = \begin{cases} \zeta + \zeta^{-1} + 2 & \text{with } \zeta \in \mathbb{F}_{\ell^2}^\times \text{ of order } e \text{ if } \gcd(\ell, e) = 1, \\ 4 & \text{otherwise.} \end{cases}$$

Proof. The proof is identical to that of Proposition 1.

Coming back to point counting: suppose we have a Jacobian $J_{\mathcal{C}}$ with real multiplication by \mathcal{O}_F ; we want to compute the characteristic polynomial

$$\chi_\pi(T) = T^4 - tT^3 + (2q + s)T^2 - tqT + q^2.$$

If we have a totally positive element μ in \mathcal{O}_F such that $\mu\bar{\mu} = \ell$, then we know that $\chi_\pi(T) \pmod{\ell}$ splits into two quadratic factors:

$$\chi_\pi(T) \equiv \chi_{\pi, \mu}(T)\chi_{\pi, \bar{\mu}}(T) \equiv (T^2 - uT + q)(T^2 - u'T + q) \pmod{\ell},$$

so

$$t \equiv u + u' \pmod{\ell} \quad \text{and} \quad s \equiv uu' - 2q \pmod{\ell}. \quad (3.6)$$

Given precomputed Hilbert modular polynomials G_μ and $G_{\bar{\mu}}$, then, we can specialize them at the RM invariants of J_ℓ and factor to determine the order of Frobenius on $J_\ell[\mu]$ and on $J_\ell[\bar{\mu}]$ using Proposition 4. We can then apply Proposition 5 and Equations (3.6) to restrict the possible values of s and t modulo ℓ .

The question of how best to exploit this extra modular information remains open. Atkin's match-and-sort and Joux and Lercier's Chinese-and-match algorithms for elliptic curves cannot be re-used directly here, because they were designed to solve the one-dimensional problem of determining the elliptic trace, while here we have the two-dimensional problem of determining (s, t) .

3.5.4 Prime types for real multiplication by \mathcal{O}_F

The factorization patterns in Proposition 4 are the same as those we saw for specialized elliptic modular polynomials in §3.2.4. This leads us to define an analogous classification of prime types, for totally positive elements in \mathcal{O}_F of prime norm.

Definition 5. Let μ be a totally positive element of \mathcal{O}_F such that $\mu\bar{\mu} = (\ell)$ for some prime $\ell \neq 2, p$. We say that

- μ is \mathcal{O}_F -**Elkies** for a vanilla triple (\mathcal{A}, ξ, t) with RM invariants (J_1, J_2, J_3) if the factorization type of $G_\mu(J_1, J_2, J_3, x)$ is $(1, 1, e, \dots, e)$ with $e > 1$;
- μ is \mathcal{O}_F -**Atkin** for a vanilla triple (\mathcal{A}, ξ, t) with RM invariants (J_1, J_2, J_3) if the factorization type of $G_\mu(J_1, J_2, J_3, x)$ is (e, \dots, e) with $e > 1$; and
- μ is \mathcal{O}_F -**volcanic** for a vanilla triple (\mathcal{A}, ξ, t) with RM invariants (J_1, J_2, J_3) if the factorization type of $G_\mu(J_1, J_2, J_3, x)$ is $(1, e)$ or $(1, \dots, 1)$.

If $K \cong \text{End}_{\mathbb{F}_q}(\mathcal{A}) \otimes \mathbb{Q}$ is Galois then the type of μ completely determines the type of $\bar{\mu}$ (and vice versa). For general K , however, this does not hold: the type of $\bar{\mu}$ is not determined by the type of μ .

3.5.5 The parity of the number of factors of G_μ

The following proposition is the genus-2 real multiplication analogue of Equation (3.2) (cf. [35, Proposition 6.3]).

Proposition 6. Let (\mathcal{A}, ξ, t) be a triple describing a vanilla abelian surface over \mathbb{F}_q with real multiplication by \mathcal{O}_F , and with RM invariants (J_1, J_2, J_3) . Let μ be a totally positive element of \mathcal{O}_F of prime norm $\mu\bar{\mu} = \ell$, let $\chi_{\pi, \mu}(T) = T^2 - uT + q$ be the characteristic polynomial of π restricted to $\mathcal{A}[\mu]$, and let r denote the number of irreducible factors in the factorization of $G_\mu(J_1, J_2, J_3, x)$. Then

$$(-1)^r = \left(\frac{q}{\ell}\right).$$

Proof. If ℓ divides $u^2 - 4q$ and π has order ℓ in Case (3) of Proposition 4, then the result is true. Suppose therefore that $u^2 - 4q \not\equiv 0 \pmod{\ell}$, that is, we are in Cases (1) or (2) of Proposition 4, and let $\mathcal{T} \subseteq \mathrm{GL}_2(\mathbb{F}_\ell)$ be a maximal torus containing π . In other words, we take $\mathcal{T} = \{\mathrm{diag}(\alpha, \beta) : \alpha, \beta \in \mathbb{F}_\ell^\times\}$ split in Case (2), and \mathcal{T} non-split (i.e., isomorphic to $\mathbb{F}_{\ell^2}^\times$) in Case (1). The image $\overline{\mathcal{T}}$ of \mathcal{T} in $\mathrm{PGL}_2(\mathbb{F}_\ell)$ is cyclic of order $\ell + 1$ in Case (1) and $\ell - 1$ in Case (2). The determinant induces an isomorphism $\det: \overline{\mathcal{T}}/\overline{\mathcal{T}}^2 \rightarrow \mathbb{F}_\ell^\times/(\mathbb{F}_\ell^\times)^2$. The action of π is via $\det(\pi) = q$, and we obtain an isomorphism $\det: \overline{\mathcal{T}}/\langle \overline{\mathcal{T}}^2, \pi \rangle \rightarrow \mathbb{F}_\ell^\times/(\langle (\mathbb{F}_\ell^\times)^2, q \rangle)$. This shows that the index $[\overline{\mathcal{T}} : \pi]$ is odd if and only if q is not a square mod ℓ . Since the number r of irreducible factors of $G_\mu(J_1, J_2, J_3, x)$ over \mathbb{F}_q is equal to $r = (l + 1)/e$ or $r = 2 + (l - 1)/e = [\overline{\mathcal{T}} : \pi]$, the proposition follows.

3.6 The case $F = \mathbb{Q}(\sqrt{5})$: Gundlach–Müller invariants

All of the theory above can be made much more explicit in the case where $F = \mathbb{Q}(\sqrt{5})$, where the invariants J_1, J_2 , and J_3 are known as Gundlach–Müller invariants [11, 30]. Our computational results are based on this case, so we will work out the details here, following the treatment in [22].

Fixing a square root of 5 in \mathbb{C} , we set $\varepsilon := (1 + \sqrt{5})/2$ and $\bar{\varepsilon} := (1 - \sqrt{5})/2$; each is the image of the fundamental unit of $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ under one of its two embeddings into \mathbb{C} . Let

$$q_1 := e\left(\frac{\varepsilon z_1 - \bar{\varepsilon} z_2}{\sqrt{5}}\right) \quad \text{and} \quad q_2 := e\left(\frac{z_2 - z_1}{\sqrt{5}}\right) \quad \text{for} \quad z = (z_1, z_2) \in \mathbb{H}^2.$$

The Eisenstein series of even weight $k \geq 2$ are defined by

$$g_k(z) = 1 + \sum_{t=a+b\bar{\varepsilon} \in \mathcal{O}_F^+} b_k(t) q_1^a q_2^b,$$

where the coefficients $b_k(t)$ are defined by

$$b_k(t) = \kappa_k \sum_{(\mu) \supseteq (t)} N(\mu)^{k-1} \quad \text{with} \quad \kappa_k = \frac{(2\pi)^{2k} \sqrt{5}}{(k-1)! 2^k \zeta_F(k)} \in \mathbb{Q}$$

(here $N(\mu)$ is the norm $\#\mathcal{O}_F/(\mu)$). The Hilbert modular forms s_6, s_{10}, s_{12} , and s_{15} of respective weight 6, 10, 12, and 15 for $\mathcal{H}_{\mathbb{Q}(\sqrt{5})}$ are defined by

$$\begin{aligned}
s_6 &:= -\frac{67}{2^5 \cdot 3^3 \cdot 5^2} (g_6 - g_2^3), \\
s_{10} &:= \frac{1}{2^{10} \cdot 3^5 \cdot 5^5 \cdot 7} \left(191 \cdot 2161 g_{10} - 5 \cdot 67 \cdot 2293 g_2^2 g_6 + 2^2 \cdot 3 \cdot 7 \cdot 4231 g_2^5 \right), \\
s_{12} &:= \frac{1}{2^2} (s_6^2 - g_2 s_{10}), \\
s_5^2 &:= s_{10}, \\
s_{15}^2 &:= 5^5 s_{10}^3 - \frac{5^3 g_2^2 s_6 s_{10}^2}{2} + \frac{g_2^5 s_{10}^2}{2^4} + \frac{3^2 \cdot 5^2 g_2 s_6^3 s_{10}}{2} - \frac{g_2^4 s_6^2 s_{10}}{2^3} - 2 \cdot 3^3 s_6^5 + \frac{g_2^3 s_6^4}{2^4}.
\end{aligned}$$

Finally, the Gundlach–Müller invariants for $\mathbb{Q}(\sqrt{5})$ are

$$J_1 := s_6/g_2^3, \quad J_2 := g_2^5/s_5^2, \quad \text{and} \quad J_3 := s_5^3/s_{15}.$$

The Hilbert modular polynomials for $\mathbb{Q}(\sqrt{5})$ are too large to reproduce here, but they can be downloaded from `martindale.info`.⁷

Proposition 7 ([22, Proposition 4.5] with correction to $\phi^*(j_1)$). *For $F = \mathbb{Q}(\sqrt{5})$, the Igusa invariants pull back to*

$$\begin{aligned}
\phi^*(j_1) &= 4J_2(3J_1^2J_2 - 2)^5, \\
\phi^*(j_2) &= \frac{1}{2}J_2(3J_1^2J_2 - 2)^3, \\
\phi^*(j_3) &= 2^{-3}J_2(2J_1^2J_2 - 2)^2(4J_1^2J_2 + 2^5 \cdot 3^2J_1 - 3).
\end{aligned}$$

For our computations, we want to write J_1, J_2 and J_3 in terms of the Siegel modular forms $\psi_4, \psi_6, \chi_{10}$ and χ_{12} . (For a canonical way of writing J_1, J_2 and J_3 in terms of Igusa–Clebsch invariants, we refer to [24, Example 2.5.4].)

Proposition 8 ([24, Example 2.5.4]). *For $F = \mathbb{Q}(\sqrt{5})$, we have*

$$\begin{aligned}
J_2 &= \phi^*((\psi_4\psi_6/\chi_{10} - 3^5 2^{12})(-2 - 2(\psi_6^2 - 2^{12} 3^6 \chi_{12})/\psi_4^3)^{-1}), \\
J_1 &= 3^2 2^5 J_2^{-1} + \phi^*(2^{-6} 3^{-3} (1 - (\psi_6^2 - 2^{12} 3^6 \chi_{12})/\psi_4^3)), \\
J_3 &= 5^5 - 2^{-1} 5^3 J_1 J_2 + 2^{-4} J_2 + 2^{-1} 3^2 5^2 J_2^2 J_1^3 - 2^{-3} J_1^2 J_2^2 - 2 \cdot 3^3 J_2^3 J_1^5 + 2^{-4} J_2^3 J_1^4.
\end{aligned}$$

The choice of square root for J_3 corresponds to the choice of embedding ι .

Proposition 8 can be used to find RM invariants for curves drawn from families with known real multiplication, before factoring specialized Hilbert modular polynomials in those RM invariants to derive information on Frobenius. However, it also crystallizes the rationality question alluded to at the end of §3.4.4: as we see, a set of values of the Hilbert modular forms over \mathbb{F}_q (or, equivalently, a tuple of Igusa or Igusa–Clebsch invariants over \mathbb{F}_q) only determine J_1, J_2 , and J_3^2 over \mathbb{F}_q .

⁷ The polynomials $H_{\mu,3}$ do not appear there, but only G_μ is required to apply our results in §3.5.

To get J_3 , we need to choose a square root of J_3^2 ; but J_3^2 is not guaranteed to be a square in \mathbb{F}_q . If J_3^2 is not a square in \mathbb{F}_q , then we cannot apply Propositions 3 or 4—not even if J_3 does not appear unsquared in the specialized polynomial G_μ .

3.7 Experimental results

In order to validate the factorization patterns of Proposition 4, we ran a series of experiments for $F = \mathbb{Q}(\sqrt{5})$, using the family of curves [37]

$$\mathcal{C}_a : y^2 = x^5 - 5x^3 + 5x + a$$

whose Jacobians all have real multiplication by $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$. This family was used in the point-counting records of [7]. The Igusa–Clebsch invariants of \mathcal{C}_a are

$$(A, B, C, D) = \left(2^5 \cdot 5^2 \cdot 7, 2^{10} \cdot 5^4, -2^{13} \cdot 5^5 \cdot (9a^2 - 236), 2^{20} \cdot 5^5 \cdot (a^2 - 4)^2 \right).$$

Our experiments treated

1. the ramified prime $\ell = 5$, with $\mu = (5 + \sqrt{5})/2$, and the modular polynomial G_μ from `martindale.info`;
2. the split prime $\ell = 11$, with $\mu = (7 + \sqrt{5})/2$, and the modular polynomial G_μ from `martindale.info`.

We collected statistics on the factorization patterns for 10000 tests. For each test, we chose a random prime q of ten decimal digits, and we chose a randomly from \mathbb{F}_q subject to the requirement that \mathcal{C}_a be nonsingular, which is $a^2 \neq 4$. We then applied the formulæ of Equation (3.5) and Proposition 8 to obtain the RM invariants J_2 and J_1 for the Jacobian of \mathcal{C}_a , as well as the squared invariant J_3^2 .

In half the cases on average, J_3^2 had a square root in \mathbb{F}_q ; in these cases we could obtain J_3 , and proceed to factor $G_\mu(J_1, J_2, J_3, x)$. The average frequencies of the resulting factorization patterns appear in Tables 3.1 and 3.2 (here we take the averages over the roughly 5000 tests where J_3^2 has a root in \mathbb{F}_q ; for the two roots J_3 and $-J_3$ in \mathbb{F}_q , we always obtained the same factorization pattern).

According to Proposition 4, we would expect that $1/\ell$ of the time μ should be $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -volcanic, $(\ell - 1)/2\ell$ of the time μ should be $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Elkies, and $(\ell - 1)/2\ell$ of the time μ should be $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Atkin. The summary of our above results in Table 3.3 appears to confirm this. This gives us considerable confidence that the Hilbert modular polynomials computed in [24, Chapter 2] are correct.

Finally, we ran the same tests on Milio’s modular polynomial⁸ $\Phi(\mathfrak{J}_1, \mathfrak{J}_2, X)$ for $\ell = 5$ and $\mu = (5 + \sqrt{5})/2$, where $\mathfrak{J}_1 = J_2$ and $\mathfrak{J}_2 = J_1 J_2$. We obtained exactly the same factorization patterns each time J_3 was in \mathbb{F}_q .

⁸ Available from <https://members.loria.fr/EMilio/modular-polynomials/>

Factorization pattern, type of μ	Number found	Percentage
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Elkies: $(1, 1, e, \dots, e)$ with $e > 1$	total 1835	total 36.8%
	(1, 1, 4)	1266 25.4%
	(1, 1, 2, 2)	569 11.4%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Atkin: (e, \dots, e) with $e > 1$	total 2049	total 41.1%
	(6)	844 16.9%
	(3, 3)	794 15.9%
	(2, 2, 2)	411 8.2%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Volcanic: $(1, e)$ or $(1, \dots, 1)$	total 1105	total 22.1%
	(1, 5)	1058 21.2%
	(1, 1, 1, 1, 1)	47 0.9%

Table 3.1 Factorization pattern frequencies for the modular polynomial $G_\mu(J_1, J_2, J_3, x)$ for $\mu = (5 + \sqrt{5})/2$ of norm $\ell = 5$. The degree of $G_\mu(J_1, J_2, J_3, x)$ in x is 6. We only factored when J_3^2 was a square in \mathbb{F}_q , which happened in 4989 of the 10000 trials (49.9%).

Factorization pattern, type of μ	Number found	Percentage
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Elkies: $(1, 1, e, \dots, e)$ with $e > 1$	total 2262	total 44.7%
	(1, 1, 10)	994 19.7%
	(1, 1, 5, 5)	1040 20.6%
	(1, 1, 2, 2, 2, 2, 2)	228 4.5%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Atkin: (e, \dots, e) with $e > 1$	total 2329	total 46.1%
	(12)	859 17.0%
	(6, 6)	404 8.0%
	(4, 4, 4)	424 8.4%
	(3, 3, 3, 3)	429 8.5%
	(2, 2, 2, 2, 2, 2)	213 4.2%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -volcanic: $(1, e)$ or $(1, \dots, 1)$	total 466	total 9.2%
	(1, 11)	461 9.1%
	(1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1)	5 0.1%

Table 3.2 Factorization pattern frequencies for the modular polynomial $G_\mu(J_1, J_2, J_3, x)$ for $\mu = (7 + \sqrt{5})/2$ of norm $\ell = 11$. The degree of $G_\mu(J_1, J_2, J_3, x)$ in x is 12. We only factored when J_3^2 was a square in \mathbb{F}_q , which happened in 5057 of the 10000 trials (50.6%).

		Prime type frequencies for μ		
		$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -volcanic	$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Elkies	$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ -Atkin
$\mu = \frac{5-\sqrt{5}}{2}$	Theory	20.0%	40.0%	40.0%
	Experiments	22.1%	36.8%	41.1%
$\mu = \frac{7+\sqrt{5}}{2}$	Theory	9.1%	45.5%	45.5%
	Experiments	9.2%	44.7%	46.1%

Table 3.3 Experimental evidence supporting the correctness of Martindale's Hilbert modular polynomials.

Acknowledgements This article reports on work carried out at the workshop *Algebraic Geometry for Coding Theory and Cryptography* at the Institute for Pure and Applied Mathematics (IPAM), University of California, Los Angeles, February 22–26, 2016. The authors thank IPAM for its generous support. Chloe Martindale was supported by an ALGANT-doc scholarship in association with Universiteit Leiden and Université de Bordeaux. Maïke Massierer was supported by the Australian Research Council (DP150101689).

References

1. G. Bisson, R. Cosset, and D. Robert. AVIsogenies: a library for computing isogenies between abelian varieties. <http://avisogenies.gforge.inria.fr>.
2. D. G. Cantor. Computing in the Jacobian of a hyperelliptic curve. *Math. Comp.*, 48(177):95–101, 1987.
3. R. Carls. *A generalized arithmetic geometric mean*. PhD thesis, University of Groningen, The Netherlands, 2004.
4. J.-M. Couveignes and T. Ezome. Computing functions on Jacobians and their quotients. *LMS J. Comput. Math.*, 18(1):555–577, 2015.
5. E. V. Flynn. The Jacobian and formal group of a curve of genus 2 over an arbitrary ground field. *Math. Proc. Cambridge Philos. Soc.*, 107(3):425–441, 1990.
6. P. Gaudry and R. Harley. Counting points on hyperelliptic curves over finite fields. In W. Bosma, editor, *Algorithmic Number Theory, 4th International Symposium, ANTS-IV (Leiden, The Netherlands)*, volume 1838 of *Lecture Notes in Computer Science*, pages 313–332. Springer, Berlin, 2000.
7. P. Gaudry, D. Kohel, and B. Smith. Counting points on genus 2 curves with real multiplication. In D. H. Lee and X. Wang, editors, *Advances in Cryptology—ASIACRYPT 2011 (Seoul, South Korea)*, volume 7073 of *Lecture Notes in Computer Science*, pages 504–519. Springer, Heidelberg, 2011.
8. P. Gaudry and E. Schost. Genus 2 point counting over prime fields. *J. Symbolic Comput.*, 47(4):368–400, 2012.
9. G. van der Geer. *Hilbert Modular Surfaces*, volume 16 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer, Berlin, 1988.
10. D. Grant. Formal groups in genus two. *J. Reine Angew. Math.*, 411:96–121, 1990.
11. K.-B. Gundlach. Die Bestimmung der Funktionen zur Hilbertschen Modulgruppe des Zahlkörpers $\mathbb{Q}(\sqrt{5})$. *Math. Ann.*, 152(3):226–256, 1963.
12. M. C. Harrison. An extension of Kedlaya’s algorithm for hyperelliptic curves. *J. Symbolic Comput.*, 47(1):89–101, 2012.
13. D. Harvey. Kedlaya’s algorithm in larger characteristic. *Int. Math. Res. Not. IMRN*, 2007(22):Art. ID rnm095, 29, 2007.
14. E. W. Howe and H. J. Zhu. On the existence of absolutely simple abelian varieties of a given dimension over an arbitrary field. *J. Number Theory*, 92(1):139–163, 2002.
15. J.-i. Igusa. On Siegel modular forms of genus two. *Amer. J. Math.*, 84:175–200, 1962.
16. J.-i. Igusa. Modular forms and projective invariants. *Amer. J. Math.*, 89:817–855, 1967.
17. A. Joux and R. Lercier. “Chinese & match”, an alternative to Atkin’s “match and sort” method used in the SEA algorithm. *Math. Comp.*, 70(234):827–836, 2001.
18. K. S. Kedlaya. Counting points on hyperelliptic curves using Monsky–Washnitzer cohomology. *J. Ramanujan Math. Soc.*, 16(4):323–338, 2001.
19. S. Lang. *Introduction to Algebraic and Abelian Functions*, volume 89 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1982.
20. S. Lang. *Algebraic Number Theory*, volume 16 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1986.
21. K. Lauter, M. Naehrig, and T. Yang. Hilbert theta series and invariants of genus 2 curves. *J. Number Theory*, 161:146–174, 2016.

22. K. Lauter and T. Yang. Computing genus 2 curves from invariants on the Hilbert moduli space. *J. Number Theory*, 131(5):936–958, 2011.
23. R. Lercier. *Algorithmique des courbes elliptiques dans les corps finis*. PhD thesis, École Polytechnique, Palaiseau, France, 1997.
24. C. Martindale. *Isogeny Graphs, Modular Polynomials, and Applications*. PhD thesis, Universiteit Leiden, 2017. in preparation.
25. J.-F. Mestre. Lettre à Gaudry et Harley. <https://webusers.imj-prg.fr/~jean-francois.mestre/lettreGaudryHarley.ps>, 2001.
26. J.-F. Mestre. Algorithme pour compter des points de courbes en petite caractéristique et petit genre. <https://webusers.imj-prg.fr/~jean-francois.mestre/rennescrypto.ps>, 2002. notes from a talk given at the Rennes cryptography seminar.
27. E. Milio. *Computing modular polynomials in dimension 2*. PhD thesis, Université de Bordeaux, Dec. 2015.
28. E. Milio. A quasi-linear time algorithm for computing modular polynomials in dimension 2. *LMS J. Comput. Math.*, 18(1):603–632, 2015.
29. J. S. Milne. Abelian varieties. In *Arithmetic Geometry (Storrs, Connecticut, 1984)*, pages 103–150. Springer, New York, 1986.
30. R. Müller. Hilbertsche Modulformen und Modulfunktionen zu $\mathbb{Q}(\sqrt{5})$. *Arch. Math. (Basel)*, 45(3):239–251, 1985.
31. J. Pila. Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Math. Comp.*, 55(192):745–763, 1990.
32. H.-G. Rück. Abelian surfaces and Jacobian varieties over finite fields. *Compositio Math.*, 76(3):351–366, 1990.
33. T. Satoh. On p -adic point counting algorithms for elliptic curves over finite fields. In C. Fieker and D. R. Kohel, editors, *Algorithmic Number Theory (Sydney, 2002)*, volume 2369 of *Lecture Notes in Computer Science*, pages 43–66. Springer, Berlin, 2002.
34. R. Schoof. Elliptic curves over finite fields and the computation of square roots mod p . *Math. Comp.*, 44(170):483–494, 1985.
35. R. Schoof. Counting points on elliptic curves over finite fields. *J. Théor. Nombres Bordeaux*, 7(1):219–254, 1995.
36. A. V. Sutherland. On the evaluation of modular polynomials. In *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium (San Diego, California)*, volume 1 of *Open Book Series*, pages 531–555. Mathematical Sciences Publishers, Berkeley, CA, 2013.
37. W. Tautz, J. Top, and A. Verberkmoes. Explicit hyperelliptic curves with real multiplication and permutation polynomials. *Canad. J. Math.*, 43(5):1055–1064, 1991.
38. L. C. Washington. *Elliptic Curves: Number Theory and Cryptography*, volume 50 of *Discrete Mathematics and its Applications*. Chapman & Hall/CRC, Boca Raton, FL, second edition, 2008.