MATHEMATICS OF COMPUTATION
S 0025-5718(04)01682-5
Article electronically published on May 25, 2004

MODULAR EQUATIONS FOR HYPERELLIPTIC CURVES

P. GAUDRY AND É. SCHOST

ABSTRACT. We define modular equations describing the ℓ-torsion subgroups of the Jacobian of a hyperelliptic curve. Over a finite base field, we prove factorization properties that extend the well-known results used in Atkin's improvement of Schoof's genus 1 point counting algorithm.

INTRODUCTION

Modular equations relating invariants of ℓ-isogenous elliptic curves are a fundamental tool in computational arithmetic geometry. A great effort has been devoted to obtaining equations sparser or with smaller coefficients than the classical polynomials Φ_ℓ [12], so nowadays these equations can be computed efficiently even for quite large ℓ. One of their important applications is the determination of the cardinality of an elliptic curve defined over a finite field [24]: the best method to date, at least for prime finite fields, is the Schoof-Elkies-Atkin algorithm, in which the ℓ-torsion structure is widely used.

Nevertheless, very little is known about similar equations for higher genus curves. Since the hyperelliptic case is the best suited for computations, we restrict to this situation. Our goal in this article is then twofold:

- We define modular equations for hyperelliptic curves, without using modular forms. In the particular case of genus 1, our equations coincide with those introduced by Charlap, Coley and Robbins in [11].
- When the base field is finite, we prove that the well-known factorization properties of genus 1 modular equations extend to our higher genus construction. This makes them amenable for use in higher genus extensions of the Atkin improvement of Schoof's initial algorithm [27].

Here is a brief overview of our construction. Consider a hyperelliptic curve $\mathcal{C}$ of genus g, $\mathbf{Jac}(\mathcal{C})$ its Jacobian, and ℓ a prime. The quotient of the Jacobian by a subgroup of order ℓ is an abelian variety ℓ-isogenous to $\mathbf{Jac}(\mathcal{C})$, but in genus greater than 1 it is in general not the Jacobian of a curve. General abelian varieties are more intricate to handle than Jacobians of curves, for which invariants can be easily computed, so we rather study directly the ℓ-torsion subgroup of the Jacobian. Our modular equations are thus defined using the group structure of the ℓ-torsion subgroup.

Received by the editor July 15, 2002 and, in revised form, August 16, 2003.
2000 *Mathematics Subject Classification.* Primary 11Y40; Secondary 11G20, 11Y16.
Key words and phrases. Modular equations, hyperelliptic curves, Schoof-Elkies-Atkin algorithm.

©2004 Copyright held by the authors

More precisely, these equations are univariate polynomials whose roots are in correspondence with the cyclic subgroups of the ℓ-torsion group. This definition avoids the use of modular forms, so it is valid over any perfect field. The construction is very similar to that of resolvents in Galois theory; as such, when the base field is finite, the factorization patterns of the modular equations are very specific, and carry enough information to be of use in higher genus Schoof point-counting algorithms.

As an example, we have detailed the relationship between the 3-torsion modular equation of a genus 2 curve and the cardinality of its Jacobian modulo 3. This equation is now used within Magma's hyperelliptic curve package [1] as part of the point-counting algorithm, since in many cases the Jacobian order modulo 3 can be deduced quickly using this equation. For large finite fields of cryptographic size, the gain brought by this method is marginal, as the computation modulo 3 becomes a tiny part of the whole computation. Yet, in a generalist system such as Magma, it is also important to optimize point counting algorithms for smaller base fields. For such situations, for fields of order up to about 10^6, using the 3-torsion modular equation yields a significant speed-up.

The paper is organized as follows. In Section 1, we precise the notation used in the sequel. The modular equations are defined in Section 2, where we also give their basic properties and detail the example of genus 1. In Section 3, we prove that the modular equations have the expected specialization properties. This is crucial for the computational point of view, which is studied in Section 4. In Section 5, we finally consider the finite field case, and show how the factorization patterns of our modular equations extend the well-known case of genus 1; we apply this for the point-counting problem.

Acknowledgments. We thank François Morain for his numerous comments and suggestions. We are grateful to John Boxall for giving us references about the Manin-Mumford conjecture. The heaviest computations were done on the machines of the CNRS–École polytechnique MEDICIS computation center [2], using the Magma computer algebra system [1]. The second author is a member of the TERA project [3].

1. Notation

Let k be a perfect field of characteristic different from 2 and $\mathcal{C}$ a genus g hyperelliptic curve defined over k. We suppose that the affine part of $\mathcal{C}$ is defined by the equation $y^2 = f(x)$, with f monic of degree $2g+1$, and for simplicity we shall say that $\mathcal{C}$ is the curve defined by $y^2 = f(x)$. The unique point at infinity on $\mathcal{C}$ is denoted by ∞.

We also assume that the characteristic of k is different from $2g+1$, so that we can transform $f(x)$ into a polynomial whose coefficient in x^{2g} is zero. This simplification is similar to what is often done in genus 1 when taking an equation of the form $y^2 = x^3 + Ax + B$. Our results also hold in characteristic $2g+1$, but with different equations.

We denote the Jacobian of $\mathcal{C}$ by $\mathbf{Jac}(\mathcal{C})$. This is a projective variety defined over k; the canonical injection $\mathcal{C} \to \mathbf{Jac}(\mathcal{C})$ associates to $P \in \mathcal{C}$ the divisor class of $P - \infty$; it is also defined over k.

If K is an extension field of k, we may distinguish the curves defined on k and K by $y^2 = f(x)$, by denoting them $\mathcal{C}/k$ and $\mathcal{C}/K$. Then the injection

$\mathcal{C}/K \to \mathbf{Jac}(\mathcal{C}/K)$ extends the injection $\mathcal{C}/k \to \mathbf{Jac}(\mathcal{C}/k)$, and the group law on $\mathbf{Jac}(\mathcal{C}/K)$ extends that of $\mathbf{Jac}(\mathcal{C}/k)$.

In particular, let $\overline{k}$ be an algebraic closure of k. Then for a prime ℓ, we will denote by $\mathbf{Jac}[\ell]$ the subgroup of ℓ-torsion elements of $\mathbf{Jac}(\mathcal{C}/\overline{k})$.

Let τ be the hyperelliptic involution on $\mathcal{C}/\overline{k}$, and let ι denote the injection $\mathcal{C}/\overline{k} \to \mathbf{Jac}(\mathcal{C}/\overline{k})$. As a consequence of the Riemann-Roch theorem, any element in $\mathbf{Jac}(\mathcal{C}/\overline{k})$ can be uniquely represented by a divisor of the form $D = \sum_{1\le j\le r} \iota(P_j)$ with the following properties:

(1) all P_j are points on the affine part of $\mathcal{C}/\overline{k}$,
(2) $P_j \neq \tau(P_{j'})$ for all $j \neq j'$,
(3) r is at most g.

The integer r is called the *weight* of D.

Let D and $\{P_j\}_{1\le j\le r}$ be as above; since the points P_j are not at infinity, we may take $P_j = (x_j, y_j, 1)$. Then the *Mumford-Cantor* representation of D [25, 9] is defined by

$$D = \langle u(x), v(x)\rangle = \langle x^r + u_{r-1}x^{r-1} + \cdots + u_0, v_{r-1}x^{r-1} + \cdots + v_0\rangle,$$

where $u = \prod_{1\le j\le r}(x - x_j)$ and $v(x_j) = y_j$ holds with suitable multiplicities, so that u divides $v^2 - f$. Since k is perfect, the divisor D is defined over a field K containing k if and only if the polynomials u and v have coefficients in K.

For j in $0, \ldots, r-1$, we will denote by $u_j(D)$ (resp. $v_j(D)$) the coefficient u_j (resp. v_j) in this representation.

2. Modular equations

2.1. **Definitions.** Let ℓ be an odd prime different from the characteristic of k. In this subsection, we define the ℓ-th modular equation of a genus g hyperelliptic curve $\mathcal{C}$ defined over k.

To this end, we consider the ℓ-torsion divisors in $\mathbf{Jac}(\mathcal{C}/\overline{k})$. The assumption that ℓ differs from the characteristic of k implies that the number of ℓ-torsion divisors of nonzero weight is $\ell^{2g} - 1$ [22]. From now on, we assume that all these divisors have weight exactly g; see subsection 2.3 for the relevance of this assumption.

Genericity assumption. *All nonzero ℓ-torsion divisors in $\mathbf{Jac}(\mathcal{C}/\overline{k})$ have weight g.*

Let D be an ℓ-torsion divisor. The divisors

$$\langle D\rangle = \left\{ -\left[\frac{\ell-1}{2}\right] D, \ldots, -D, 0, D, \ldots, \left[\frac{\ell-1}{2}\right] D \right\}$$

form a cyclic subgroup of cardinality ℓ in $\mathbf{Jac}[\ell]$. Our objective is to be able to "separate" these subgroups, using only algebraic constructions. To this effect we choose a function $t_\ell(D)$ with values in $\overline{k}$, which takes a constant value on each of the subgroups $\langle D\rangle$. Our modular equations may then be thought as a minimal polynomial of t_ℓ.

Precisely, we define t_ℓ as the following sum:

$$t_\ell(D) = \sum_{1\le i\le \frac{\ell-1}{2}} u_{g-1}\big([i]D\big). \tag{1}$$

Our genericity assumption implies that this sum is well-defined for all nonzero ℓ-torsion divisors D. Note that $[-i]D$ and $[i]D$ have the same u_{g-1}-coordinate, so

even though we restrict the number of summands to $(\ell-1)/2$, $t_\ell(D)$ depends only on the subgroup generated by D, as requested.

We next define the polynomial $\chi_\ell \in \overline{k}[T]$, whose roots are the values taken by t_ℓ on the nonzero ℓ-torsion divisors:

$$\chi_\ell = \prod_{D \in \mathbf{Jac}[\ell]\setminus\{0\}} \big(T - t_\ell(D)\big).$$

The polynomial χ_ℓ is an $(\ell-1)$-th power in $\overline{k}[T]$. Indeed $\mathbf{Jac}[\ell]\setminus\{0\}$ can be written as the disjoint union of the $\frac{\ell^{2g}-1}{\ell-1}$ sets $\langle D\rangle \setminus \{0\}$, and the function $t_\ell(D)$ takes a constant value on each part of this partition.

We now show that χ_ℓ is actually in $k[T]$. Let σ be in $\mathrm{Gal}(\overline{k}/k)$. If D is any divisor, $\sigma\big(u_{g-1}(D)\big) = u_{g-1}\big(\sigma(D)\big)$. Also, σ commutes with the group law, whence $\sigma\big([i]D\big) = [i]\sigma(D)$, so σ induces a permutation among the nonzero ℓ-torsion divisors. If D is such a divisor, then the equality

$$\sigma\big(t_\ell(D)\big) \quad = \quad t_\ell\big(\sigma(D)\big)$$

obviously holds. Since σ permutes the ℓ-torsion divisors, this equality shows that χ_ℓ is left invariant by σ, so χ_ℓ is in $k[T]$. Since k is a perfect field, and χ_ℓ is an $(\ell-1)$-th power in $\overline{k}[T]$, there exists a polynomial Ξ_ℓ with coefficients in k such that $\chi_\ell = \Xi_\ell^{\ell-1}$.

Definition 1. The unique monic polynomial Ξ_ℓ such that $\chi_\ell = \Xi_\ell^{\ell-1}$ is called the ℓ-th *modular equation* of $\mathcal{C}$.

The polynomial Ξ_ℓ has degree $\frac{\ell^{2g}-1}{\ell-1}$. To emphasize the dependence on the curve $\mathcal{C}$, it may also be denoted by $\Xi_\ell(\mathcal{C})$.

The rest of this article is devoted to describing the main properties of these equations, how to compute them and how to use them for cardinality computation, in the case when k is a finite field.

Remark 1. Our choice of the function t_ℓ is arbitrary. In Section 5, we show that the interesting case is when Ξ_ℓ is squarefree, which happens when t_ℓ takes distinct values on distinct cyclic subgroups. Unfortunately, this will not be the case for all curves; for such curves, an alternative choice of t_ℓ may solve the problem:

Instead of considering the sum of the u_{g-1}-coordinates of half of the divisors in the subgroup, we choose some integer k and form the sum of the k-th power of any linear combination of all the coordinates (u, v). Then, we might have to extend the summation in equation (1) to all elements in the subgroup $\langle D\rangle$, since not all coordinates are negation-invariant. The subsequent results follow in a similar manner for such alternative constructions.

Yet in practice, choosing the coordinate u_{g-1} yields the polynomial with smallest coefficients when working over $\mathbb{Q}$, and in most of our experiments in genus 1 and 2, this polynomial turned out to be squarefree, as requested.

Remark 2. In the sequel, we will often consider curves with generic coefficients. Thus we define for once and for all the *generic curve of genus g* as the curve of equation

$$\mathcal{C}_g \ : y^2 = x^{2g+1} + F_{2g-1}x^{2g-1} + \cdots + F_0,$$

over the rational function field $\mathbb{Q}(F_0, \ldots, F_{2g-1})$. In this case, the polynomial Ξ_ℓ belongs to $\mathbb{Q}(F_0, \ldots, F_{2g-1})[T]$, and satisfies the following homogeneity property.

Theorem 1. *The ℓ-th modular equation of the curve $\mathcal{C}_g$ is weighted homogeneous, when giving weight 1 to T and weight $2g+1-i$ to F_i for $i = 0, \ldots, 2g-1$.*

Proof. Let λ be a nonzero rational, and let $\widetilde{\mathcal{C}_g}$ be the curve defined by

$$y^2 = x^{2g+1} + \widetilde{F_{2g-1}} x^{2g-1} + \cdots + \widetilde{F_0},$$

where $\widetilde{F_i} = \lambda^{2g+1-i} F_i$, for $i = 1, \ldots, 2g-1$. Then the map $\varphi : \mathcal{C}_g \to \widetilde{\mathcal{C}_g}$ defined by $\varphi(x,y) = (\lambda x, \lambda^{2g+1} y)$ is an isomorphism between $\mathcal{C}_g$ and $\widetilde{\mathcal{C}_g}$. This isomorphism extends to an isomorphism between $\mathbf{Jac}(\mathcal{C}_g)$ and $\mathbf{Jac}(\widetilde{\mathcal{C}_g})$, which acts as follows in the Mumford-Cantor representation:

$$(u_0, \ldots, u_{g-1}, v_0, \ldots, v_{g-1}) \mapsto (\lambda^g u_0, \ldots, \lambda u_{g-1}, \lambda^{2g+1} v_0, \ldots, \lambda^{g+2} v_{g-1}).$$

Given an ℓ-torsion divisor D on $\mathbf{Jac}(\mathcal{C}_g)$, the value $t_\ell(D)$ is sent to $\lambda t_\ell(D)$. Thus

$$\Xi_\ell(F_0, \ldots, F_{2g-1}, t_\ell) = 0 \iff \Xi_\ell(\lambda^{2g+1} F_0, \ldots, \lambda^2 F_{2g-1}, \lambda t_\ell) = 0.$$

This proves the theorem. □

The weighted homogeneity implies that not all monomials appear in the modular equation for the generic curve. As a consequence, our modular equations are somewhat sparse, and we shall see below that for elliptic curves they provide a much smaller alternative to the classical modular polynomials Φ_ℓ.

Remark 3. In our formalism, the modular equation for 2-torsion Ξ_2 is ill-defined in genus greater than 1. Indeed, the genericity assumption for 2-torsion is never satisfied, since the $2g+1$ roots of the defining polynomial $f(x)$ give the abscissae of $2g+1$ weight 1 divisors of 2-torsion. In the particular case of elliptic curves, we can set $\Xi_2 = f$.

2.2. **The elliptic case.** We illustrate our definition on an elliptic curve $\mathcal{E}$, given by an equation $y^2 = f(x)$, with f monic of degree 3. In genus 1, the genericity assumption is always satisfied, since the only divisor whose weight is not maximal is zero.

If $P = (x, y)$ is a point on $\mathcal{E}$ and i a positive integer, the coordinates of $[i]P$ are rational functions of P, see [30]:

$$[i]P = \left(\frac{\phi_i(P)}{\psi_i(P)^2}, \frac{\omega_i(P)}{\psi_i(P)^3} \right).$$

The polynomials $\phi_i(P)$, $\psi_i(P)^2$, and also $\psi_i(P)$ if i is odd, are polynomials in x only. To follow the notation of the previous subsection, we see them as polynomials in the variable T.

Given an odd prime ℓ, the abscissae of the ℓ-torsion points are the roots of ψ_ℓ. Let P be such a point; for i in $1, \ldots, \frac{\ell-1}{2}$, the denominator in the rational function

$$\left(\frac{\phi_i(P)}{\psi_i(P)^2} \right)$$

is coprime to ψ_ℓ. The image of this rational function modulo ψ_ℓ is a polynomial $h_{i,\ell}$ in $k[T]$ which gives the abscissa of $[i]P$ in terms of the abscissa of P, for P of ℓ-torsion. Then, for all ℓ-torsion points P, $t_\ell(P)$ is given by the sum

$$t_\ell(P) = \sum_{1 \le i \le \frac{\ell-1}{2}} h_{i,\ell}\big(x(P)\big).$$

The polynomial χ_ℓ is thus the characteristic polynomial of $\sum_{1\leq i\leq \frac{\ell-1}{2}} h_{i,\ell}$ modulo ψ_ℓ, and the modular equation $\Xi_\ell \in k[T]$ is the $(\ell-1)$-th root of χ_ℓ.

Let us take $f = x^3+F_1x+F_0$, defining what we called the generic curve of genus 1 over $\mathbb{Q}(F_0,F_1)$. Then the first values of Ξ_ℓ are

$$\begin{aligned}
\Xi_3 &= T^4 + 2F_1T^2 + 4F_0T - \tfrac{1}{3}F_1^2,\\
\Xi_5 &= T^6 + 20F_1T^4 + 160F_0T^3 - 80F_1^2T^2 - 128F_1F_0T - 80F_0^2,\\
\Xi_7 &= T^8 + 84F_1T^6 + 1512F_0T^5 - 1890F_1^2T^4 - 9072F_1F_0T^3\\
&\quad + (-21168F_0^2 + 644F_1^3)T^2 + 5832F_1^2F_0T - 567F_1^4,\\
\Xi_{11} &= T^{12} + 550F_1T^{10} + 27500F_0T^9 - 103125F_1^2T^8 - 1650000F_1F_0T^7\\
&\quad + (-13688400F_0^2 + 645700F_1^3)T^6 + 20625000F_1^2F_0T^5\\
&\quad + (35793120F_1F_0^2 - 11407385F_1^4)T^4\\
&\quad + (34041920F_0^3 - 58614160F_1^3F_0)T^3\\
&\quad + (-175832976F_1^2F_0^2 - 2177802F_1^5)T^2\\
&\quad + (-235016704F_1F_0^3 + 1351692F_1^4F_0)T\\
&\quad - 110680064F_0^4 + 6297984F_1^3F_0^2 - 321651F_1^6.
\end{aligned}$$

These polynomials were already considered by Charlap, Coley and Robbins in [11], where the authors constructed them via modular forms. Our modular equations are a generalization to higher genus.

Remark 4. Except for Ξ_3, for which a factor $\frac{1}{3}$ occurs, there are no denominators in the coefficients of the modular equations of the generic elliptic curve. This fact is proven in [11] using properties of modular forms. In higher genus we do not know a priori whether there are denominators in the modular equations of the generic curves. The computation in Section 4 shows that the modular equation Ξ_3 of the genus 2 generic curve does not have any denominator, but we do not expect this to be true in general.

2.3. **Relevance of the genericity assumption.** As mentioned in the previous subsection, the genericity assumption is satisfied in genus 1 for all curves, for all torsion indices coprime to the characteristic of the base field.

This condition is also satisfied for all genus 2 curves for 3-torsion. To see this, consider a genus 2 curve $\mathcal{C}$. A divisor with nonmaximal weight is of the form $P-\infty$ for some point $P\in\mathcal{C}$. Then the equality [3] $(P-\infty) = 0$ can be rewritten as [2] $(P-\infty) = -(P-\infty)$, which implies that $P=\infty$ by the Riemann-Roch theorem. Thus, except for zero, all 3-torsion divisors have weight 2.

The genericity assumption is closely related to the Manin-Mumford conjecture which states that the Jacobian of a curve over the complex field contains only finitely many torsion elements of weight 1. More generally, Lang's conjecture, which is now known to be true [17, p. 435], implies that the Jacobian of a given curve over the complex field contains only finitely many torsion elements of nonmaximal weight, as soon as this Jacobian is simple. As a consequence, for a given curve with simple Jacobian, the number of primes ℓ for which the genericity assumption does not hold is finite, hence the name.

Note finally that this condition is true for all ℓ for the curve of genus 2 defined by $y^2 = x^5 + 5x^3 + x$, see [7]. Using the specialization theorem given in Section 3, we deduce that the genericity assumption is also true for all ℓ for the generic curve of genus 2.

3. Specialization properties

In the elliptic curve point-counting methods, a widely used strategy is to compute modular equations over the rationals and then reduce them modulo the characteristic of the base field. In this section, we want to legitimate this approach for our modular equations. Our purpose is thus to prove the intuitive result that the reductions of the modular equations of a curve $\mathcal{C}$ coincide with the modular equations of the reduction of $\mathcal{C}$.

We are interested both in specializing the coefficients of a curve defined over a rational function field, and also in reducing the coefficients of a curve defined over a number field. Thus we work in the setting of local and global fields, which will help encompass both notions. Throughout this section, *we assume once and for all that all the fields are perfect.*

We recall below the definition of reduction and good reduction of a curve defined on a local or a global field. Then Theorem 2 proves that if $\mathcal{C}$ is a curve with good reduction, and if the reduced curve satisfies the genericity assumption for some prime ℓ, then this is also the case for $\mathcal{C}$, and its ℓ-modular equation specializes as expected; the main ingredient of the proof is the injectivity of the reduction of the ℓ-torsion, as proven for instance in [17]. We will use Theorem 2 for computational purposes in Section 4.

First, some notation is necessary. If (K, v) is a non-Archimedean local field, we denote by

$$R_K = \{a \in K, v(a) \leq 1\}, \quad \mathfrak{m}_K = \{a \in K, v(a) < 1\}$$

respectively the ring of integers of K and its maximal ideal.

With this notation, let $\mathcal{C} : y^2 = f(x)$ be a hyperelliptic curve defined over K, such that f has its coefficients in R_K. We say that $\mathcal{C}$ has *good reduction* if $2\mathrm{disc}(f)$ is not in $\mathfrak{m}_K$. In this case, the curve defined by the reduction of $y^2 = f(x)$ modulo $\mathfrak{m}_K$ has the same genus as $\mathcal{C}$. This curve is called the reduction of $\mathcal{C}$; its Jacobian is the reduction of the Jacobian of $\mathcal{C}$ modulo $\mathfrak{m}_K$, see [23].

If now $\mathcal{K}$ is a global field and v a non-Archimedean valuation of $\mathcal{K}$, the completion K of $\mathcal{K}$ at v is a local field. Thus the above discussion enables us to define the notion of good reduction at v of a curve $\mathcal{C}$ defined over $\mathcal{K}$: if $\mathcal{C}$ is defined by an equation with coefficients in $\mathcal{K} \cap R_K$, $\mathcal{C}$ has good reduction at v if $\mathcal{C}/K$ has good reduction. Then the specialization properties are stated as follows.

Theorem 2. *Let $\mathcal{K}$ be a global field, and v a non-Archimedean valuation of $\mathcal{K}$. Let K be the completion of $\mathcal{K}$ at v, R_K its ring of integers, $\mathfrak{m}_K$ its maximal ideal and k the residual field $R_K/\mathfrak{m}_K$. Let $\mathcal{C}$ be a genus g hyperelliptic curve defined over K, and ℓ an odd prime different from the characteristic of k.*

Assume that $\mathcal{C}$ is defined by an equation with coefficients in $\mathcal{K} \cap R_K$, that $\mathcal{C}$ has good reduction at v, and that the reduced curve $\overline{\mathcal{C}}$ satisfies the genericity assumption for the prime ℓ. Then $\mathcal{C}$ satisfies the genericity assumption for ℓ, all coefficients of $\Xi_\ell(\mathcal{C})$ are in $\mathcal{K} \cap R_K$, and $\Xi_\ell(\overline{\mathcal{C}}) = \Xi_\ell(\mathcal{C})$ modulo $\mathfrak{m}_K$.

Before proving the theorem, a few comments are in order. Theorem 2 applies to specializations at non-Archimedean places of number fields, so for instance it allows for the reduction modulo prime numbers of modular equations with rational coefficients. The second obvious application is the specialization of modular equations with coefficients in *univariate* function fields defined over a perfect field.

We considered in Section 2 the case of the *generic curve of genus* g, which is defined over the rational function field $\mathbb{Q}(F_0, \ldots, F_{2g-1})$ by the equation

$$y^2 = x^{2g+1} + F_{2g-1}x^{2g-1} + \cdots + F_0.$$

Its modular equations have coefficients in $\mathbb{Q}(F_0, \ldots, F_{2g-1})$. For $g > 1$, the fraction field of $\mathbb{Q}[[F_0, \ldots, F_{2g-1}]]$ is not a local field, so Theorem 2 does not apply directly, i.e., it does not allow one to specialize $F_0, \ldots, F_{2g-1}$ at once. Yet we may circumvent this difficulty. Consider the isomorphism

$$\mathbb{Q}(F_0, \ldots, F_{2g-1}) \simeq \mathbb{Q}(F_0, \ldots, F_{2g-2})(F_{2g-1}).$$

The right-hand side is a univariate function field over $\mathbb{Q}(F_0, \ldots, F_{2g-2})$, for which Theorem 2 applies; i.e., specializing F_{2g-1} is allowed. Iterating this process allows us to successively specialize $F_{2g-2}, \ldots, F_0$, as requested.

Proof of Theorem 2. Since K is an extension field of $\mathcal{K}$, the modular equations $\Xi_\ell(\mathcal{C}/\mathcal{K})$ and $\Xi_\ell(\mathcal{C}/K)$ coincide. Thus it is enough to prove the result for the curve $\mathcal{C}/K$, defined over the local field K.

Let $\mathfrak{K}$ be an algebraic closure of K, let $\mathbf{Jac}[\ell]$ be the ℓ-torsion divisors on $\mathbf{Jac}(\mathcal{C}/\mathfrak{K})$ and let $\mathfrak{I}$ be the canonical injection $\mathcal{C}/\mathfrak{K} \to \mathbf{Jac}(\mathcal{C}/\mathfrak{K})$. Each divisor $D \neq 0$ in $\mathbf{Jac}[\ell]$ can be uniquely written $D = \sum_{1 \leq j \leq r(D)} \mathfrak{I}(P_j^D)$, where $P_1^D, \ldots, P_{r(D)}^D$ are points in $\mathcal{C}/\mathfrak{K}$, not at infinity and not pairwise conjugate, and where $r(D)$ is at most g.

We let L be a finite extension of K, such that all ℓ-torsion divisors D and all points $P_1^D, \ldots, P_{r(D)}^D$ are L-rational. Then L is a non-Archimedean local field for a valuation that extends that of K, and it is still denoted by v. We denote by R_L and $\mathfrak{m}_L$ the ring of integers of L and its maximal ideal, and the reduction modulo $\mathfrak{m}_L$ is denoted by a bar. We still denote by $\mathfrak{I}$ the canonical injection $\mathcal{C}/L \to \mathbf{Jac}(\mathcal{C}/L)$.

The curve $\mathcal{C}/L$ has good reduction; the reduced curve, defined over the residual field of L, is still denoted by $\overline{\mathcal{C}}$; the canonical injection $\overline{\mathcal{C}} \to \mathbf{Jac}(\overline{\mathcal{C}})$ is denoted by ι. Note that the residual fields of K and L have the same characteristic.

Let D be an ℓ-torsion divisor on $\mathbf{Jac}(\mathcal{C}/L)$. We simplify the notation $D = \sum_{1 \leq j \leq r(D)} \mathfrak{I}(P_j^D)$, writing $D = \sum_{1 \leq j \leq r} \mathfrak{I}(P_j)$ instead. Then, due to the good reduction of $\mathcal{C}$, the following holds:

$$\overline{D} = \overline{\sum_{1 \leq j \leq r} \mathfrak{I}(P_j)} = \sum_{1 \leq j \leq r} \overline{\mathfrak{I}(P_j)} = \sum_{1 \leq j \leq r} \iota\left(\overline{P_j}\right).$$

Moreover, $\overline{D}$ is an ℓ-torsion divisor on $\mathbf{Jac}(\overline{\mathcal{C}})$. Since the characteristic of k is different from ℓ, Theorem C.2.6 in [17] shows that $\overline{D}$ is not zero. Using the genericity assumption on $\overline{\mathcal{C}}$, this shows that $\overline{D}$ has weight g, which implies that $r = g$, so $\mathcal{C}$ satisfies the genericity assumption too. This also implies that all points $\overline{P_j}$ are on the affine part of $\overline{\mathcal{C}}$.

Let us write $P_j = (x_j, y_j, z_j)$, with all coordinates in the ring of integers R_L, and at least one of them not in $\mathfrak{m}_L$. By the above remark, z_j does not reduce to zero modulo $\mathfrak{m}_L$, so $v(z_j) = 1$. Dividing by z_j, we write P_j as $(X_j, Y_j, 1)$, with coordinates in R_L; then $\overline{P_j}$ is given by $(\overline{X_j}, \overline{Y_j}, 1)$. Taking all points P_j into account, this shows that $u_{g-1}(D)$ is in R_L, and a short calculation also gives that

$u_{g-1}\left(\overline{D}\right) = \overline{u_{g-1}\left(D\right)}$. Recall now that the function $t_\ell(D)$ is defined as

$$t_\ell(D) = \sum_{1 \leq i \leq \frac{\ell-1}{2}} u_{g-1}\big([i]D\big),$$

so that $t_\ell(D)$ is in R_L. Since $\overline{[i]D} = [i]\overline{D}$, we finally conclude that $\overline{t_\ell(D)} = t_\ell(\overline{D})$.

The set of nonzero ℓ-torsion divisors on $\mathcal{C}$ reduces to the set of nonzero ℓ-torsion divisors on $\overline{\mathcal{C}}$, so $\chi_\ell(\mathcal{C}) = \prod_{D \in \mathbf{Jac}[\ell]\setminus\{0\}} \big(T - t_\ell(D)\big)$ is in $R_L[T]$, and $\big(\chi_\ell(\mathcal{C}) \mod \mathfrak{m}_L\big) = \chi_\ell(\overline{\mathcal{C}})$. Since $\mathcal{C}$ is defined over K, $\chi_\ell(\mathcal{C})$ is actually in $R_K[T]$, and $\big(\chi_\ell(\mathcal{C}) \mod \mathfrak{m}_K\big) = \chi_\ell(\overline{\mathcal{C}})$.

From the definition $\chi_\ell(\mathcal{C}) = \Xi_\ell(\mathcal{C})^{\ell-1}$, we see that $\Xi_\ell(\mathcal{C})$ is also in $R_K[T]$ by Gauss' lemma. The above reduction can then be written $\big(\Xi_\ell(\mathcal{C})^{\ell-1} \mod \mathfrak{m}_K\big) = \Xi_\ell(\overline{\mathcal{C}})^{\ell-1}$. Since both polynomials are monic, $\big(\Xi_\ell(\mathcal{C}) \mod \mathfrak{m}_K\big) = \Xi_\ell(\overline{\mathcal{C}})$. □

4. Algorithms

We consider now the question of computing the modular equations, with an emphasis on the genus 2 case. In the first subsection, we show how Cantor's division polynomials [10] can be used to give a description of the ℓ-torsion subgroup of a hyperelliptic Jacobian in genus 2, from which its ℓ-th modular equation can be deduced; we illustrate this with examples for 3- and 5-torsion. We also present a solution by specialization techniques for 3-torsion in genus 2. The arbitrary genus case is finally addressed, using Adleman and Huang's algorithm for computing the ℓ-torsion points on a hyperelliptic Jacobian [4].

4.1. Computing Ξ_ℓ in genus 2.

Computing Ξ_ℓ is a two-stage process: first compute a representation of the ℓ-torsion divisors, then compute the modular equation using this information. This strategy was already hinted at in subsection 2.2, where we addressed the elliptic case. In genus 1, the ℓ-torsion divisors are the roots of the elliptic division polynomials, and the modular equations come from characteristic polynomial computations modulo these division polynomials. In genus 2, the torsion divisors can be characterized using Cantor's division polynomials [10]; the subsequent characteristic polynomial computations are analogous to the elliptic case.

In this subsection, $\mathcal{C}$ is a genus 2 curve defined over a field k by the equation $y^2 = f(x)$, with $f(x) = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$. The solution we present here demands further conditions on the ℓ-torsion divisors, which are generically satisfied. Explicitly, there exists a nonempty Zariski-open subset W of $\overline{k}^4$ such that the following analysis applies to the curve $\mathcal{C}$ as soon as f_0, f_1, f_2, f_3 belong to W.

4.1.1. *First step: computing the torsion divisors.* We use the strategy from [13]. For a weight 2 divisor $D = P_1 + P_2 - 2\infty$, the condition $[\ell]D = 0$ can be restated as

$$[\ell](P_1 - \infty) = -[\ell](P_2 - \infty).$$

Let x_1, x_2, y_1, y_2 be new variables, and take $P_1 = (x_1, y_1)$, $P_2 = (x_2, y_2)$. Then the Mumford-Cantor coordinates of $[\ell](P_1 - \infty)$ and $-[\ell](P_2 - \infty)$ are rational functions of x_1, x_2, y_1, y_2 with coefficients in k. These rational functions can be

derived efficiently by recursive formulae given in [10]:

$$[\ell]P_1 = \left\langle x^2 + \frac{d_1^{(\ell)}(x_1)}{d_0^{(\ell)}(x_1)}x + \frac{d_2^{(\ell)}(x_1)}{d_0^{(\ell)}(x_1)},\ \frac{y_1 e_1^{(\ell)}(x_1)}{e_0^{(\ell)}(x_1)}x + \frac{y_1 e_2^{(\ell)}(x_1)}{e_0^{(\ell)}(x_1)} \right\rangle,$$

where $d_0^{(\ell)}, d_1^{(\ell)}, d_2^{(\ell)}, e_0^{(\ell)}, e_1^{(\ell)}, e_2^{(\ell)}$ are univariate polynomials of degrees respectively $2\ell^2-3$, $2\ell^2-2$, $2\ell^2-1$, $3\ell^2-2$, $3\ell^2-3$, $3\ell^2-2$. From now on we assume that ℓ is fixed, and drop the superscript (ℓ) for simplicity.

The genericity assumption on $\mathcal{C}$ can then be made explicit: in [10] it is shown that a weight 1 divisor $D = (x_1, y_1) - \infty$ is an ℓ-torsion divisor if and only if $d_1(x_1) = d_0(x_1) = 0$. Hence the genericity assumption is valid on $\mathcal{C}$ if and only if d_1 and d_0 are coprime polynomials.

If this is the case, we form the equations in x_1, x_2, y_1, y_2 expressing the equality $[\ell](P_1 - \infty) = -[\ell](P_2 - \infty)$:

$$\begin{aligned}
\frac{d_1(x_1)}{d_0(x_1)} &= \frac{d_1(x_2)}{d_0(x_2)},\\
\frac{d_2(x_1)}{d_0(x_1)} &= \frac{d_2(x_2)}{d_0(x_2)},\\
\frac{y_1 e_1(x_1)}{e_0(x_1)} &= -\frac{y_2 e_1(x_2)}{e_0(x_2)},\\
\frac{y_1 e_2(x_1)}{e_0(x_1)} &= -\frac{y_2 e_2(x_2)}{e_0(x_2)}.
\end{aligned}$$

Further equations and inequalities must be added to this system:

$$\begin{aligned}
&y_1^2 = f(x_1), \quad y_2^2 = f(x_2),\\
&x_1 \neq x_2, \quad d_0(x_1) e_0(x_1) \neq 0.
\end{aligned}$$

The equations specify that P_1 and P_2 are indeed points on the curve. The first inequality is necessary to discard the obvious solutions $P_1 = -P_2$, but may also eliminate points P_1 such that $[\ell](P_1 - \infty)$ is of 2-torsion. The second inequality allows one to clean the denominators, and it is equivalent to the assumption that $[\ell](P_1 - \infty)$ has weight 2.

Due to the symmetry in (P_1, P_2), for a general curve, this system has $2(\ell^4 - 1)$ solutions, but this number may drop if $[\ell](P_1 - \infty)$ is of 2-torsion or has weight 1, for some ℓ-torsion divisor of the form $D = P_1 + P_2 - 2\infty$. Letting $I_\ell \subset k[x_1, x_2, y_1, y_2]$ be the ideal defining the solutions of the above system, we thus have to check that I_ℓ has the maximal number of solutions, i.e., $2(\ell^4-1)$. We assume from now on that we are in this favorable case. Furthermore, as an outcome of an elimination procedure, we assume that a basis of the quotient $k[x_1, x_2, y_1, y_2]/I_\ell$ and the corresponding multiplication table are available.

4.1.2. *Second step: deducing Ξ_ℓ.* In genus 2, the function $t_\ell(D)$ defined in Section 2 becomes

$$t_\ell(D) = \sum_{1 \leq i \leq \frac{\ell-1}{2}} u_1\big([i]D\big).$$

Roughly speaking, we want to compute t_ℓ modulo I_ℓ. To this effect, for $i = 1, \ldots, (\ell-1)/2$, we let $h_{i,\ell}$ be the polynomial in $k[x_1, x_2, y_1, y_2]$ which takes the value $u_1\big([i]D\big)$ when evaluated on an ℓ-torsion divisor D. The reader may refer to subsection 2.2, where similar polynomials are derived in genus 1.

To obtain $h_{i,\ell}$, we compute the Mumford-Cantor coordinates of $[i](P_1 - \infty) + [i](P_2 - \infty)$, where P_1 (resp. P_2) is the point of coordinates (x_1, y_1), resp. (x_2, y_2), all computations being done modulo I_ℓ. These computations can be done using Cantor's division polynomials and Cantor's addition algorithm [9]. As such, they involve divisions modulo I_ℓ, which are possible in general, but may fail in unlucky cases.

We assume from now on that all these divisions can be done. In this lucky case, due to the symmetry in (P_1, P_2), the polynomial χ_ℓ defined in Section 2 is the square root of the characteristic polynomial of $T_\ell := \sum_{1 \leq i \leq \frac{\ell-1}{2}} h_{i,\ell}$ modulo I_ℓ. Thus, knowing the characteristic polynomial of T_ℓ, we deduce the modular equation Ξ_ℓ by taking its $2(\ell-1)$-th root.

4.1.3. *Complexity estimates.* To perform the above tasks, we can use standard effective elimination algorithms such as Gröbner bases [8] or geometric resolution procedures [14, 15]. Yet, the very specific shape of the system defining the ideal I_ℓ enables us to estimate precisely the number of operations necessary to compute Ξ_ℓ.

Explicitly, we now show that, up to logarithmic factors, Ξ_ℓ can be computed in $O(\ell^8)$ base field operations. To this effect, we assume that fast algorithms for polynomial multiplication, gcd, ..., are used and neglect all logarithmic factors: the notation $g \in O\tilde{}(f)$ means that g belongs to $O(f \log(f)^a)$, for some constant a. All relevant references can be found in [31].

First, following [13], we note that for a generic choice of the curve $\mathcal{C}$, the minimal polynomial R of x_1 modulo I_ℓ has the maximal possible degree, that is, $\ell^4 - 1$. The article [13] already shows how to compute this polynomial: using fast resultant and gcd computations, this can be done in $O\tilde{}(\ell^6)$ base field operations. In fact, by a subresultant computation, x_2 can be expressed in terms of x_1, and expressions for y_1 and y_2 follow at no higher cost. Thus we obtain a Gröbner basis of I_ℓ for a lexicographic order. Then, performing a multiplication or a division in the above basis can be done within $O\tilde{}(\ell^4)$ base field operations.

Let us turn to the second step, the deduction of Ξ_ℓ. We follow the algorithm given in the previous paragraphs. Performing Cantor's addition algorithm requires a fixed number of operations modulo I_ℓ. Since $O(\ell)$ such additions are necessary, all polynomials $h_{i,\ell}$ defined above, and thus T_ℓ, can be computed in $O\tilde{}(\ell^5)$ base field operations.

The characteristic polynomial of T_ℓ can then be deduced in $O\tilde{}(\ell^8)$ base field operations. To this effect, we first compute the first $2(\ell^4 - 1)$ powers of T_ℓ, for a cost of $O\tilde{}(\ell^8)$ base field operations. Then we compute the trace of each of these polynomials in the extension $k \to k[x_1, x_2, y_1, y_2]$ without affecting the complexity. Finally, we recover the characteristic polynomial of T_ℓ using Newton's relations for the same cost. Yun's squarefree decomposition algorithm [32] finally gives the polynomial Ξ_ℓ, without increasing the asymptotic complexity.

Note that the above complexity estimates depend only on ℓ. Therefore in the case of a finite base field of cardinality q, the number of bit-operations grows like $O\tilde{}(\ell^8 \log q)$. Hence the size of the base field is not a limiting factor.

4.1.4. *First example: 3-torsion.* For the particular case of the 3-torsion in genus 2, we already noted that the genericity assumption is always satisfied. Further simplifications arise in this case.

Consider again the equations in x_1, x_2, y_1, y_2 expressing that $[3](P_1 - \infty) = -[3](P_2 - \infty)$. The computation shows that all denominators in these equations are

powers of $f(x_1)$ and $f(x_2)$; the Riemann-Roch theorem shows that these are nonzero quantities if $P_1 + P_2 - 2\infty$ is of 3-torsion. In other words, the ideal I_3 described above always has the maximal number of solutions, i.e., $2 \times (3^4 - 1) = 160$.

Cantor's division polynomials d_0, d_1, d_2, e_0, e_1, e_2 for 3-torsion have degrees 17, 16, 15, 25, 25, 24. We define the ideal $I_3 \subset k[x_1, x_2, y_1, y_2]$ using these polynomials; then the definitions in Section 2 show that for 3-torsion, $t_3(D)$ coincides with $u_1(D)$, so χ_3 is the square root of the characteristic polynomial of $x_1 + x_2$ modulo I_3, and Ξ_3 is the square root of χ_3.

As an example, if $\mathcal{C}$ is the genus 2 curve defined over $\mathbb{F}_p$ with $p = 101009$ by the equation

$$y^2 = x^5 + x^3 - 2233x^2 + 1944x + 21551,$$

then its modular equation for 3-torsion is

$$\begin{aligned}\Xi_3(T) = {} & T^{40} + 312T^{38} + 74066T^{37} + 21549T^{36} + 98476T^{35} + 77413T^{34} \\ & + 11876T^{33} + 87429T^{32} + 62497T^{31} + 26387T^{30} + 53806T^{29} + 55344T^{28} \\ & + 3150T^{27} + 85491T^{26} + 39525T^{25} + 66578T^{24} + 43546T^{23} + 37423T^{22} \\ & + 17475T^{21} + 96511T^{20} + 45626T^{19} + 38344T^{18} + 36106T^{17} + 88202T^{16} \\ & + 80287T^{15} + 16826T^{14} + 27075T^{13} + 64347T^{12} + 84421T^{11} \\ & + 99539T^{10} + 78579T^9 + 29813T^8 + 30858T^7 + 58361T^6 + 37204T^5 \\ & + 39712T^4 + 98618T^3 + 92702T^2 + 47474T + 57754.\end{aligned}$$

4.1.5. *Second example: 5-torsion.* We illustrate the case of 5-torsion with the same curve. The ideal $I_5 \subset \mathbb{F}_p[x_1, x_2, y_1, y_2]$ is generated using the univariate polynomials d_0, d_1, d_2, e_0, e_1, e_2 for 5-torsion, of respective degrees 49, 48, 47, 73, 73, 72. As indicated above, we apply the resolution techniques of [13], so as to obtain a basis of $\mathbb{F}_p[x_1, x_2, y_1, y_2]/I_5$.

For 5-torsion in genus 2, $t_5(D)$ equals $u_1(D) + u_1\big([2]D\big)$. Taking $P_1 = (x_1, y_1)$, $P_2 = (x_2, y_2)$ and $D = P_1 + P_2 - 2\infty$, we compute the Mumford-Cantor coordinates of $[2]D$ modulo I_5, from which we deduce t_5 as a function defined modulo I_5, its characteristic polynomial χ_5, and finally Ξ_5:

$$\begin{aligned}\Xi_5(T) = {} & T^{156} + 100894T^{155} + 45811T^{154} + 44233T^{153} + 39740T^{152} + 95818T^{151} + 4458T^{150} \\ & + 93570T^{149} + 28550T^{148} + 1113T^{147} + 63748T^{146} + 65156T^{145} + 19143T^{144} \\ & + 19730T^{143} + 33367T^{142} + 20907T^{141} + 63820T^{140} + 51940T^{139} + 42326T^{138} \\ & + 29817T^{137} + 75942T^{136} + 59745T^{135} + 97234T^{134} + 85218T^{133} + 24915T^{132} \\ & + 16689T^{131} + 93260T^{130} + 46818T^{129} + 27999T^{128} + 93775T^{127} + 2219T^{126} \\ & + 19973T^{125} + 1129T^{124} + 52225T^{123} + 6886T^{122} + 85816T^{121} + 77152T^{120} \\ & + 12511T^{119} + 64657T^{118} + 14966T^{117} + 42288T^{116} + 90382T^{115} + 60923T^{114} \\ & + 40482T^{113} + 43464T^{112} + 61885T^{111} + 2196T^{110} + 60160T^{109} + 78999T^{108} \\ & + 88624T^{107} + 86206T^{106} + 1602T^{105} + 73726T^{104} + 27596T^{103} + 4276T^{102} \\ & + 93140T^{101} + 58403T^{100} + 1234T^{99} + 88895T^{98} + 59594T^{97} + 13604T^{96} + 6248T^{95} \\ & + 30964T^{94} + 40104T^{93} + 58036T^{92} + 38120T^{91} + 69691T^{90} + 75816T^{89} + 22051T^{88} \\ & + 36408T^{87} + 59984T^{86} + 96659T^{85} + 35117T^{84} + 96676T^{83} + 62595T^{82} + 51710T^{81} \\ & + 38161T^{80} + 68329T^{79} + 99009T^{78} + 10635T^{77} + 97403T^{76} + 58767T^{75} + 74987T^{74} \\ & + 26947T^{73} + 100504T^{72} + 6849T^{71} + 50414T^{70} + 38143T^{69} + 82683T^{68} + 24600T^{67} \\ & + 83911T^{66} + 80258T^{65} + 98589T^{64} + 76450T^{63} + 89676T^{62} + 9956T^{61} + 58260T^{60} \\ & + 92379T^{59} + 57187T^{58} + 40082T^{57} + 33146T^{56} + 20951T^{55} + 77435T^{54} + 3376T^{53} \\ & + 67384T^{52} + 96487T^{51} + 74090T^{50} + 65498T^{49} + 61846T^{48} + 62046T^{47} + 30397T^{46} \\ & + 18364T^{45} + 55016T^{44} + 32487T^{43} + 94900T^{42} + 33300T^{41} + 25231T^{40} + 71704T^{39} \\ & + 59710T^{38} + 44750T^{37} + 31125T^{36} + 10050T^{35} + 7371T^{34} + 29794T^{33} + 72166T^{32} \\ & + 14168T^{31} + 53045T^{30} + 80342T^{29} + 20690T^{28} + 79145T^{27} + 74121T^{26} + 1983T^{25} \\ & + 37232T^{24} + 76446T^{23} + 6132T^{22} + 98206T^{21} + 94392T^{20} + 81694T^{19} + 43792T^{18} \\ & + 60209T^{17} + 98392T^{16} + 60483T^{15} + 71502T^{14} + 94495T^{13} + 77466T^{12} + 61989T^{11} \\ & + 51160T^{10} + 65902T^9 + 28152T^8 + 70740T^7 + 94924T^6 + 4993T^5 + 30256T^4 \\ & + 62697T^3 + 27820T^2 + 55949T + 63692.\end{aligned}$$

4.1.6. *The particular case of the generic curve.* We have a special interest in the generic curve of genus 2 defined over $\mathbb{Q}(F_0, F_1, F_2, F_3)$ by the equation

$$\mathcal{C}_2 \ : y^2 = x^5 + F_3 x^3 + F_2 x^2 + F_1 x + F_0,$$

and its modular equations in $\mathbb{Q}(F_0, F_1, F_2, F_3)[T]$. Computing these polynomials is a one-time job, which becomes useful in conjunction with the specialization theorem of Section 3, as illustrated below.

The first case of interest is for 3-torsion. To compute $\Xi_3(\mathcal{C}_2)$, the elimination techniques mentioned above become cumbersome: the base field is now a rational function field, and we have no control on the degrees in (F_0, F_1, F_2, F_3) of the intermediate expressions we have to handle. Thus we used alternative resolution techniques, based on the work of the TERA group [3]: the algorithm we used is based on a symbolic Newton operator, which enables us to compute successive approximations of $\Xi_3(\mathcal{C}_2)$, with increasing precisions in (F_0, F_1, F_2, F_3), starting from the data of Ξ_3 for a curve with rational coefficients.

Using these techniques, computing $\Xi_3(\mathcal{C}_2)$ requires approximately 4 hours of CPU time using Magma, on a 4 GB, 500 MHz Compac DS20 processor. We mention that a direct approach based on a Gröbner basis computation fails, by lack of memory. For more details, we refer to [29].

The resulting polynomial $\Xi_3(\mathcal{C}_2)$ can be downloaded from the web-page of the second author `http://www.medicis.polytechnique.fr/~schost/`. It has 1747 monomials with nonzero coefficients, belongs to $\mathbb{Z}[F_0, F_1, F_2, F_3][T]$, and is monic in T. To have an estimate of its size, here are its first coefficients:

$$\begin{aligned}\Xi_3 = T^{40} &+ 312F_3T^{38} - 7904F_2T^{37} + (-16344F_3^2 + 183488F_1)T^{36}\\ &+ (337536F_3F_2 - 4179456F_0)T^{35}\\ &+ (345456F_3^3 + 464064F_3F_1 - 2381088F_2^2)T^{34}\\ &+ (-6044256F_3^2F_2 - 196322688F_3F_0 + 33608832F_2F_1)T^{33} + \cdots.\end{aligned}$$

The largest coefficient in absolute value is 304960695828480.

Now let p be a prime greater than 3 and $\mathcal{C}$ a genus 2 curve defined over $\mathbb{Z}/p\mathbb{Z}$ by the equation

$$\mathcal{C} : y^2 = x^5 + f_3 x^3 + f_2 x^2 + f_1 x + f_0.$$

Theorem 2 shows that the polynomial $\Xi_3(\mathcal{C}) \in \mathbb{Z}/p\mathbb{Z}[T]$ is obtained by substituting the values f_i for the indeterminates F_i in $\Xi_3(\mathcal{C}_2)$ after reducing its coefficient modulo p. We note that this solution is used within Magma's hyperelliptic curve package [1].

Until now, we always excluded the case $\ell = p$. To get a hint about the problems that can arise, we can reduce all the coefficients of $\Xi_3(\mathcal{C}_2)$ modulo 3, expecting some degeneracy similar to the Kronecker relations for the elliptic modular polynomials Φ_ℓ. We obtain

$$\begin{aligned}\Xi_3(\mathcal{C}_2) \quad &\equiv \quad \left(T^4 + f_1T - f_2\right)\left(T^{12} + f_1T^9 - f_0f_2T^6 + (f_0^3f_1 - f_0f_1f_2 - f_1^3)T^3\right.\\ &\quad \left. + f_0^3f_1^2 - f_0^2f_2^2 - f_0f_1^2f_2 - f_0^4f_2 - f_1^4 - f_2^3\right)^3 \mod 3\\ &\equiv \quad \left(T^4 + f_1T - f_2\right)\left(T^4 + \sqrt[3]{f_1}T^3 - \sqrt[3]{f_0f_2}T^2 + \sqrt[3]{f_0^3f_1 - f_0f_1f_2 - f_1^3}T\right.\\ &\quad \left. + \sqrt[3]{f_0^3f_1^2 - f_0^2f_2^2 - f_0f_1^2f_2 - f_0^4f_2 - f_1^4 - f_2^3}\right)^9 \mod 3.\end{aligned}$$

In characteristic 3, the Jacobian of a curve of genus 2 has at most 9 points of 3-torsion. Therefore there are at most 4 distinct values for the u_1-coordinate of

3-torsion elements, whereas Ξ_3 modulo 3 has 8 roots. We deduce that Theorem 2 cannot be generalized easily to the case $p = \ell$.

The next interesting case is that of $\Xi_5(\mathcal{C}_2)$ for 5-torsion. We estimate that this polynomial has several million monomials, so new techniques will be needed to store it, relying on the evaluation philosophy of [16].

4.2. **Computing the modular equation: the general case.** The solution of subsection 4.1 is specific to genus 2, and requires additional nonvanishing assumptions. We turn here to the general case of a genus g hyperelliptic curve $\mathcal{C}$ defined over a field k, and indicate how to compute its ℓ-th modular equation. This process is very similar to the one described above: first compute a representation of the ℓ-torsion divisors on $\mathbf{Jac}(\mathcal{C})$, then compute the modular equation. The previous additional assumptions are ruled out now, at the cost notably of a more delicate pre-computation.

As in the genus 2 case, the Mumford-Cantor coordinates of the ℓ-torsion divisors are algebraic over k. For the general case, Adleman and Huang propose in [4] an algorithm that computes a representation of these numbers by the following objects, whose existence is guaranteed if $\mathcal{C}$ satisfies the genericity assumption for ℓ:

(1) A function q in $k[u_0, \ldots, u_{g-1}, v_0, \ldots, v_{g-1}]$, for instance a linear form. The function q must take $\ell^{2g} - 1$ distinct values on $\mathbf{Jac}[\ell] \setminus \{0\}$, i.e., q must separate the points of $\mathbf{Jac}[\ell] \setminus \{0\}$.
(2) The squarefree polynomial $Q = \prod_{D \in \mathbf{Jac}[\ell]\setminus\{0\}} \big(S - q(D)\big)$ in $k[S]$.
(3) The interpolating polynomials $U_0, \ldots, U_{g-1}, V_0, \ldots, V_{g-1}$ in $k[S]$ of degree less than $\ell^{2g} - 1$ such that for all ℓ-torsion divisors D, U_i (resp. V_i) takes the value $u_i(D)$ (resp. $v_i(D)$) when evaluated at $q(D)$.

Adleman and Huang's algorithm works in a complex setting of semi-algebraic maps, which reflects the fact that in the general case, the multiplication by an ℓ-map cannot be represented by a single rational function in Mumford-Cantor's coordinates.

Once the above objects are known, they let us compute the polynomial χ_ℓ; then the modular equation Ξ_ℓ is deduced just as in subsection 4.1. The computation of χ_ℓ follows the same inspiration as in the previous subsection. Again, recall that the function $t_\ell(D)$ is defined as

$$t_\ell(D) = \sum_{1 \le i \le \frac{\ell-1}{2}} u_{g-1}\big([i]D\big).$$

For simplicity, assume first that the polynomial Q is irreducible, so $A := k[S]/(Q)$ is a field, and let s be the image of S in A. Since A is a field, for any i, we can apply Cantor's algorithm [9] of multiplication by i in $A[x]$ to the divisor

$$D(s) = \langle x^g + U_{g-1}(s)x^{g-1} + \cdots + U_0(s), V_{g-1}(s)x^{g-1} + \cdots + V_0(s)\rangle.$$

Let $h_{i,\ell} \in A$ be the u_{g-1}-coordinate of $[i]D(s)$. If we consider $h_{i,\ell}$ in $k[S]$, then $u_{g-1}\big([i]D\big) = h_{i,\ell}\big(q(D)\big)$ holds for all ℓ-torsion divisors D. This implies that

$$t_\ell(D) = \sum_{1 \le i \le \frac{\ell-1}{2}} h_{i,\ell}\big(q(D)\big)$$

also holds for all ℓ-torsion divisors D. Thus the polynomial χ_ℓ is the characteristic polynomial of $\sum_{1 \le i \le \frac{\ell-1}{2}} h_{i,\ell}$ modulo Q.

In the general case, the polynomial Q is squarefree but not irreducible; then A is not a field, but a product of fields: $A \simeq \prod A_j$. This might make it impossible to apply the multiplication by i algorithm to the divisor $D(s)$, since this algorithm requires divisions in A. A first solution to obtain χ_ℓ consists in factoring Q: applying the process described above in each field A_j yields a polynomial $\chi_{\ell,j}$, and the product of all $\chi_{\ell,j}$ is χ_ℓ.

Of course, we want to avoid the factorization. Then a better solution is dynamic evaluation techniques [19]: the iterates of $D(s)$ are computed as if $A = k[S]/(Q)$ were a field. The divisions in A are performed using Extended GCD computations. When a division occurs where the dividend is not invertible in A, we have found a new factor of Q. Then we pursue the computations modulo each factor and finally multiply all results, as above. Thus the factorization of Q is not required; only the necessary factors come out in the Extended GCD computations.

5. Factorization patterns of Ξ_ℓ over a finite field

Let $\mathcal{C}$ be a hyperelliptic curve of genus g over a finite field $\mathbb{F}_q$. Let ℓ be a prime coprime to q and let us assume that the genericity assumption holds for the ℓ-torsion on $\mathcal{C}$. The polynomial Ξ_ℓ can be factored over $\mathbb{F}_q$. Due to the Galoisian properties of the polynomial, the possible patterns of factorization are very specific. The general result is stated in the first subsection; then we give the explicit examples of genus 1 and genus 2, and the application to point-counting in the last subsection.

5.1. Link with the Frobenius action.

Let $\pi : x \mapsto x^q$ denote the q-th power Frobenius action on $\overline{\mathbb{F}_q}$, extended to $\mathcal{C}$ and to $\mathbf{Jac}(\mathcal{C})$. The ℓ-torsion subgroup $\mathbf{Jac}[\ell]$ can be viewed as an $\mathbb{F}_\ell$-vector space of dimension $2g$ on which π acts as an endomorphism. The following result is a general statement on the relation between π and the factorization patterns of the modular equations; the precise form is given in Lemma 2 below.

Theorem 3. *Assume that Ξ_ℓ is squarefree. Then the endomorphism π of the $\mathbb{F}_\ell$-vector space $\mathbf{Jac}[\ell]$ determines the factorization pattern of Ξ_ℓ.*

The proof of Theorem 3 is derived from the following lemma.

Lemma 1. *Assume that Ξ_ℓ is squarefree and let D and E be nonzero ℓ-torsion divisors. Then*

$$t_\ell(D) = t_\ell(E) \iff E \in \langle D \rangle.$$

Proof. The direction $\Leftarrow$ follows from the definition of t_ℓ. For the converse implication, since Ξ_ℓ is squarefree, the number of values taken by t_ℓ is maximal; therefore two distinct subgroups cannot give the same value. □

We define the *modified order* of a polynomial P with coefficients in a finite field as

$$\mathrm{ord}^*(P) = \min\{k \in \mathbb{N}^*,\ \deg(X^k \bmod P(X)) = 0\}.$$

This notation is used in this section and in the Appendix.

Lemma 2. *Let D be a nonzero ℓ-torsion divisor of $\mathbf{Jac}[\ell]$. Let V_D be the $\mathbb{F}_\ell$-vector space generated by the Galoisian conjugates of D:*

$$V_D = \mathrm{Span}_{\mathbb{F}_\ell}\{\pi^n(D),\ n \in \mathbb{N}\}.$$

Let P be the characteristic polynomial of π restricted to V_D. If Ξ_ℓ is squarefree, then the degree of the extension of $\mathbb{F}_q$ where $t_\ell(D)$ is defined is $\mathrm{ord}^*(P)$.

Proof. Let k be the smallest positive integer such that $\pi^k(D) \in \langle D \rangle$. Since π fixes $\mathbb{F}_q$, we check as in Section 2 that $\pi\big(t_\ell(D)\big) = t_\ell\big(\pi(D)\big)$, and so for all $i \geq 0$, $\pi^i\big(t_\ell(D)\big) = t_\ell\big(\pi^i(D)\big)$. By definition of k, for $i < k$, $\pi^i(D)$ is not in $\langle D \rangle$, so $t_\ell\big(\big(\pi^i(D)\big) \neq t_\ell(D)$ by the previous lemma, i.e., $\pi^i\big(t_\ell(D)\big) \neq t_\ell(D)$. Since $t_\ell\big(\pi^k(D)\big) = t_\ell(D)$, we have proven that k is the degree of the extension of $\mathbb{F}_q$ where $t_\ell(D)$ is defined.

We now consider the characteristic polynomial P of π restricted to the space V_D generated by the conjugates of D. Due to the definition of V_D, P is the minimal polynomial of π restricted to V_D. Denoting by λ the element of $\mathbb{F}_\ell^*$ such that $\pi^k(D) = \lambda D$, we see that the endomorphism $\pi^k - \lambda \mathrm{Id}$ is trivial on D and on all its conjugates; therefore it is trivial on the whole of V_D. Hence we have

$$X^k - \lambda \equiv 0 \mod P(X).$$

Conversely, let k' and μ be such that $X^{k'} - \mu \equiv 0 \mod P(X)$; then $\pi^{k'}$ is equal to $\mu\mathrm{Id}$ on V_D and in particular on D: therefore k' must be larger than or equal to k. The integer k is then the minimal having this property, as announced. □

Proof of Theorem 3. Let M be the matrix of the action of π on a basis $\mathcal{B}$ of $\mathbf{Jac}[\ell]$. For all D in $\mathbf{Jac}[\ell]$ expressed as a vector according to the same basis, finding V_D and the associated polynomial $P(X)$ is a matter of elementary linear algebra. The extension where the root of Ξ_ℓ associated to D is defined follows easily. Hence knowing the matrix M gives all the extensions where the roots of Ξ_ℓ are defined, which gives the factorization pattern, because the field of definition is finite. □

5.2. **Genus 1.** The factorization patterns of the modular polynomial of an elliptic curve are well known [6, p. 119]. We restate them here for completeness; this gives a flavor of the genus 2 case that follows. Note that this result is usually given for the classical modular equations Φ_ℓ relating the j-invariants of ℓ-isogenous elliptic curves; they have the same factorization patterns as ours.

In the sequel, notations such as

$$(n_1)^{\alpha_1} \cdots (n_k)^{\alpha_k} \quad \text{or} \quad (\underbrace{n_1, \ldots, n_1}_{\beta_1}, \ldots, \underbrace{n_k, \ldots, n_k}_{\beta_k})$$

stand for a squarefree polynomial having α_i irreducible factors of degree n_i, for i in $1, \ldots, k$; then $\beta_i = n_i \alpha_i$.

Theorem 4. *Let E be an elliptic curve over $\mathbb{F}_q$, and let ℓ be coprime to q. Let $X^2 - cX + q$ be the characteristic polynomial of the Frobenius endomorphism π. Assume that the modular polynomial Ξ_ℓ is squarefree. Then the set of degrees of its irreducible factors is one of the following, where r is an integer greater than 1:*

(1) $(1, \ell)$ *or* $(1, 1, \ldots, 1)$ *if* $c^2 - 4q \equiv 0 \mod \ell$.
(2) $(1, 1, r, r, \ldots r)$ *if $c^2 - 4q$ is a nonzero square modulo ℓ.*
(3) $(r, r, \ldots, r)$ *if $c^2 - 4q$ is not a square modulo ℓ.*

Case 1 corresponds to a Frobenius endomorphism having a double eigenvalue. According to the presence of one or two blocks in the Jordan decomposition we get one or the other subcase. Case 2 is the case where there are two distinct eigenvalues. Case 3 is the nondiagonalizable case.

5.3. **Genus 2.** In genus 2, we study all possible reductions of the matrix of the Frobenius endomorphism and get the corresponding factorization patterns for Ξ_ℓ. The classification is done according to the possible factorization patterns of the minimal and characteristic polynomials of the Frobenius endomorphism π, from which we deduce a block-reduction of the matrix of π; the factorization pattern of the modular equation is then deduced from Lemma 2. The two lemmata below reduce the number of cases to consider.

Lemma 3. *If the characteristic polynomial of the Frobenius endomorphism has a triple root modulo ℓ, then the multiplicity is actually 4.*

Proof. Let $(X-A)^3(X-B)$ be the characteristic polynomial of π modulo ℓ. By Weil's theorem, if λ is a root of the characteristic polynomial of π, then q/λ is also a root of it. Therefore we have $A^2 \equiv q \mod \ell$ and $AB \equiv q \mod \ell$. But q is nonzero modulo ℓ, so that $A \equiv B \mod \ell$. □

Lemma 4. *Let $V \subset \mathbb{F}_\ell^4$ be a π-stable subspace of dimension 3. Then either V is a nontrivial direct sum of π-stable subspaces, or the minimal polynomial of π equals $(X-A)^4$, for some nonzero $A \in \mathbb{F}_\ell$.*

Proof. We suppose that V cannot be decomposed as a nontrivial direct sum of π-stable subspaces, and we prove that the minimal polynomial of π equals $(X-A)^4$, for some $A \in \mathbb{F}_\ell$. We first prove that the characteristic polynomial of π takes this form.

Let χ and χ_V be the characteristic polynomials of respectively π and its restriction to the π-stable subspace V. From our hypothesis on V, we deduce that χ_V admits only one irreducible factor. Let D be a supplementary vector of V: there exists $A \in \mathbb{F}_\ell$ such that $\pi(D) = AD + R$, with $R \in V$. We deduce that A is a root of χ; by Weil's theorem, A is nonzero, and q/A is also a root of χ. But then q/A is a root of χ_V, so $\chi_V = (T - q/A)^3$. By the above lemma, we deduce that $A = q/A$, and $\chi = (T-A)^4$.

The minimal polynomial of π is some power of $(T-A)$; we now prove that it precisely equals $(T-A)^4$. To this end, let us consider the Jordan decomposition of the restriction of π to V. Our assumption on V implies that there is only one Jordan block, with A on the diagonal. Thus, there exist D_1, D_2, D_3 in V such that in the basis (D_1, D_2, D_3, D) of $\mathbb{F}_\ell^4$, the matrix of π takes the form

$$\begin{bmatrix} A & 1 & 0 & Z \\ 0 & A & 1 & Y \\ 0 & 0 & A & X \\ 0 & 0 & 0 & A \end{bmatrix},$$

for some X, Y, Z in $\mathbb{F}_\ell$. Note that by Weil's theorem, we have $A^2 = q$ in $\mathbb{F}_\ell$.

Let $e_\ell(.,.)$ denote the Weil pairing on $\mathbf{Jac}[\ell]$. Recall that this is an alternate nondegenerate bilinear form with values in the ℓ-th roots of unity in an algebraic closure of $\mathbb{F}_q$ and such that $e_\ell(\pi(\Delta), \pi(\Delta')) = (e_\ell(\Delta, \Delta'))^q$ for all ℓ-torsion divisors Δ, Δ'. Using this pairing, we now prove that $X \neq 0$; this concludes the proof, since this implies that $(\pi - AI_d)^3$ is not zero.

Using the equality $A^2 = q$, we have

$$e_\ell(D_1, D_3)^q = e_\ell(AD_1, AD_3 + D_2) = e_\ell(D_1, D_3)^q e_\ell(D_1, D_2)^A.$$

Thus, since A is not zero modulo ℓ, we deduce that $e_\ell(D_1, D_2) = 1$. Similarly, writing

$$e_\ell(D_2, D_3)^q = e_\ell(AD_2 + D_1, AD_3 + D_2) = e_\ell(D_2, D_3)^q e_\ell(D_1, D_3)^A e_\ell(D_1, D_2),$$

we deduce that $e_\ell(D_1, D_3) = 1$. Let us finally consider the equality

$$\begin{aligned} e_\ell(D_2, D)^q &= e_\ell(AD_2 + D_1, AD + XD_3 + YD_2 + ZD_1) \\ &= e_\ell(D_2, D)^q e_\ell(D_2, D_3)^{AX} e_\ell(D_1, D)^A. \end{aligned}$$

If $X = 0$, then $e_\ell(D_1, D) = 1$, which contradicts the nondegeneracy of e_ℓ. Thus $X \neq 0$. □

Using Lemmata 3 and 4, we now proceed to enumerate all admissible reduction patterns of π and apply Lemma 2 to get the corresponding factorization pattern of Ξ_ℓ. We give below all details for one case and leave the others to the reader. The results are collected in the Appendix.

5.3.1. *A complete example.* Assume that the minimal (resp. characteristic) polynomial of π can be written PQ (resp. PQ^2), with P irreducible of degree 2 and Q of degree 1. Then there exists a basis (D_1, D_2, D_3, D_4) of $\mathbf{Jac}[\ell]$ in which the matrix of π has the form

$$M = \begin{array}{|c|cc|} \hline A_2 & \multicolumn{2}{c|}{0} \\ \hline 0 & b & 0 \\ & 0 & b \\ \hline \end{array},$$

where A_2 is a 2×2 matrix whose characteristic polynomial is irreducible modulo ℓ, and b is a scalar.

For all $0 \neq D \in \langle D_3, D_4 \rangle$, we have $\pi(D) = bD$, and by Lemma 2, the subgroups of order ℓ inside $\langle D_3, D_4 \rangle$ correspond to roots of Ξ_ℓ defined over $\mathbb{F}_\ell$. There are as many roots as lines in $\langle D_3, D_4 \rangle$, i.e., $\frac{\ell^2-1}{\ell-1} = \ell + 1$.

For all $0 \neq D \in \langle D_1, D_2 \rangle$, the vector space generated by the conjugates of D is $V_D = \langle D_1, D_2 \rangle$ and the characteristic polynomial of π restricted to V_D is the characteristic polynomial of the matrix A_2. By Lemma 2, the root of Ξ_ℓ associated to the subgroup generated by D is therefore in an extension of degree $\mathrm{ord}^*(\text{Characteristic polynomial of } A_2)$, which we denote by $\mathrm{ord}(A_2)$. Again there are $\ell + 1$ roots; thus $\mathrm{ord}(A_2)$ divides $\ell + 1$. Finally, we have $(\ell + 1)/\mathrm{ord}(A_2)$ factors of degree $\mathrm{ord}(A_2)$ in Ξ_ℓ.

For all $D = E + F$, where $0 \neq E \in \langle D_1, D_2 \rangle$ and $0 \neq F \in \langle D_3, D_4 \rangle$, the vector space generated by the conjugates of D is $V_D = \langle D_1, D_2, F \rangle$ and the characteristic polynomial of π on V_D is equal to the characteristic polynomial of A_2 multiplied by $(X - b)$. We denote by $\mathrm{ord}(A_2 b)$ the value of ord^* for this polynomial; this is the degree of the extension where we find the roots. The number of lines generated by such D's is $(\ell - 1)(\ell + 1)^2$.

To summarize, Ξ_ℓ contains

- $\ell + 1$ linear factors,
- $(\ell + 1)/\mathrm{ord}(A_2)$ factors of degree $\mathrm{ord}(A_2)$,
- $(\ell - 1)(\ell + 1)^2/\mathrm{ord}(A_2 b)$ factors of degree $\mathrm{ord}(A_2 b)$.

5.4. **Application to point-counting.** Let $\mathcal{C}$ be a hyperelliptic curve of genus g over a finite field $\mathbb{F}_q$ for which we want to compute the cardinality of the Jacobian. The characteristic polynomial of the Frobenius endomorphism is denoted by $\mathcal{F}$; then $\#\mathbf{Jac}(\mathcal{C}) = \mathcal{F}(1)$.

The principle of the high genus variants of Schoof's algorithm [26, 21, 18] is to compute $\mathcal{F}(1)$ modulo several small primes ℓ using the restriction of π to $\mathbf{Jac}[\ell]$. When this is done for sufficiently many primes ℓ, $\mathcal{F}(1)$ can be deduced using the Chinese Remainder Theorem.

In the 1980's, Atkin [5] proposed a modification of Schoof's algorithm for elliptic curves: for each ℓ, he only computes the factorization pattern of the modular polynomial and reduces the number of candidates for $\mathcal{F}(1)$ modulo ℓ. It then requires more primes ℓ before having enough information to conclude. This method has exponential complexity, but with further improvements by Elkies it led to the so-called Schoof–Elkies–Atkin algorithm [28], which is the fastest method to date for large prime q.

Our modular equations can be used in the same manner for a curve $\mathcal{C}$ of genus g. Let ℓ be a small prime coprime to q. Assume that we have computed Ξ_ℓ for this curve and that it is squarefree. Then a partial factorization can be computed in order to find its factorization pattern. Following the method developed in the previous section, it is possible to find all characteristic polynomials $\mathcal{F}$ that could yield this pattern. Then computing the group order modulo ℓ for each of these cases gives a list of candidates among which the actual one lies.

This yields the following algorithm "à la Atkin":

(1) While we do not have enough information, do
   (a) Choose a new small prime ℓ coprime to $2q$;
   (b) Compute the factorization pattern of Ξ_ℓ;
   (c) Compute the values of $\chi(1) \bmod \ell$ compatible with this pattern;

(2) Determine the only value of $\chi(1)$ compatible with all the modular information.

To illustrate this approach, we give a table of the possible patterns for $\ell = 3$ in genus 2 and the corresponding candidates for the cardinality of the Jacobian modulo 3. Recall that all genus 2 curves satisfy the genericity assumption for 3-torsion.

Theorem 5. *Let $\mathcal{C}$ be a curve of genus 2 over a finite field $\mathbb{F}_q$ of characteristic different from 2 and 3. Assume that $\Xi_3(\mathcal{C})$ is squarefree. Then the factorization pattern of $\Xi_3(\mathcal{C})$ implies the values of $\#\mathbf{Jac}(\mathcal{C}) \bmod 3$ according to Table 1.*

As an example, consider again the curve $\mathcal{C}$ defined in Section 4.1 over the field $\mathbb{F}_q$ with $q = 101009 \equiv 2 \bmod 3$. Then the factorization pattern of $\Xi_3(\mathcal{C})$ is $(1)^2(2)(4)(8)^4$, which implies that $\#\mathbf{Jac}(\mathcal{C}) \equiv 0 \mod 3$.

Similarly, the factorization pattern of $\Xi_5(\mathcal{C})$ is $(3)^{52}$. Referring to the table given in the Appendix, we deduce that only the first four cases of that table are to be considered. An exhaustive enumeration of all possible characteristic polynomials in these four cases reveals that this polynomial is either $(T^2+2T+4)^2$ or $(T^2+3T+4)^2$; they both give the same number of points $\#\mathbf{Jac}(\mathcal{C}) \equiv 4 \mod 5$.

Complexity considerations. We concentrate on the case of genus 2 curves, as the analysis is already complex in this case, and gets worse in higher genus. We use

TABLE 1

$q \equiv 2 \bmod 3$	
pattern	$\#\mathbf{Jac}(\mathcal{C}) \bmod 3$
$(10)^4$	2
$(2)^2(6)^6$	1
$(4)(12)^3$	1
$(2)^{20}$	1
$(1)^2(2)(4)(8)^4$	0
$(1)^2(2)(3)^2(6)^5$	0
$(1)^8(2)^{16}$	0
$(1)^5(2)^4(3)(6)^4$	0
$(4)^{10}$	1,2

$q \equiv 1 \bmod 3$	
pattern	$\#\mathbf{Jac}(\mathcal{C}) \bmod 3$
$(5)^8$	1,2
$(1)(2)^2(3)(4)^2(12)^2$	0,2
$(1)^4(2)^2(4)^8$	0,2
$(1)(3)^4(9)^3$	0,1
$(1)^4(3)^{12}$	0,1
$(1)^{40}$	0,1
$(4)^{10}$	2
$(2)^2(6)^6$	1
$(2)^{20}$	1
$(1)^2(2)(3)^2(6)^5$	0
$(1)^5(2)^4(3)(6)^4$	0
$(1)^8(2)^{16}$	0

again the $O^{\sim}$ notation, meaning that logarithmic factors are neglected. Furthermore, we assume that fast polynomial arithmetic algorithms are used.

We fix a curve $\mathcal{C}$ of genus 2 over a finite field $\mathbb{F}_q$, and we compare the two strategies for getting some information on cardinality modulo a prime ℓ coprime to q, namely the point-counting algorithm "à la Atkin" described above, which gives only partial information, and the "plain Schoof" algorithm, which gives the exact number of points modulo ℓ. We assume that we have computed the ℓ-torsion ideal I_ℓ and the modular equation Ξ_ℓ for the curve $\mathcal{C}$.

In Atkin's algorithm, we need to find the factorization pattern of Ξ_ℓ, which amounts to performing the so-called distinct-degree factorization of Ξ_ℓ. The degree of Ξ_ℓ is in $O(\ell^3)$; therefore, using classical algorithms [31], the cost of the determination of the factorization pattern is in $O^{\sim}(\ell^6 \log q)$ operations in $\mathbb{F}_q$.

This has to be compared to the plain Schoof algorithm. For that algorithm, we assume that the ideal I_ℓ has been put in a convenient form, as explained in subsection 4.1, so that any operation in the quotient algebra costs $O^{\sim}(\ell^4)$ operations in $\mathbb{F}_q$. There are ℓ^2 possible values for the polynomial χ modulo ℓ. Testing which of them vanish modulo I_ℓ can be done in $O(\ell + \log q)$ operations in the quotient algebra, that is, $O^{\sim}\big((\ell + \log q)\ell^4\big)$ operations in $\mathbb{F}_q$.

6. Conclusion and future work

As a conclusion of the complexity estimates of the previous section, in the present state, plain Schoof's algorithm should be more efficient that our extension of Atkin's algorithm. However, we can remark that even for elliptic curves, Atkin's algorithm is not really faster than Schoof's algorithm, and becomes of practical use only when combined with Elkies' improvements. Hence our construction must be viewed as a first step towards better point-counting algorithms. We now give more details on the improvements that we expect to be possible and that could lead to better complexities.

- For the moment, we do not know how to compute Ξ_ℓ without first computing the torsion ideal I_ℓ introduced in subsection 4.1 and a convenient basis for it. We believe that it is possible to improve on the complexity estimate $O^{\sim}(\ell^8)$ that we gave there.

- The large cost of Atkin's algorithm comes from the determination of the factorization pattern of Ξ_ℓ. In the analysis we took a quadratic complexity for factoring algorithms. There exist better algorithms, based on asymptotically fast linear algebra [20]. Even with their subquadratic complexity Atkin's algorithm remains slower than plain Schoof's algorithm, but further progress in that direction could turn the situation in favor of Atkin's method.
- In the case of elliptic curves, Atkin's algorithm is used in combination with Elkies' method to form the so-called SEA algorithm. The idea is to check if Ξ_ℓ has a linear factor, then, if this is the case, to compute the factor of the torsion ideal I_ℓ corresponding to that root, and work modulo this lower degree ideal to conclude. In higher genus, we wish to apply the same strategy, so as to design a hyperelliptic Elkies' algorithm:
  - Define *hyperelliptic Elkies primes* to be the primes ℓ for which Ξ_ℓ has a linear factor, or more generally a small degree factor, and study the probability for a prime to be of this type.
  - Design an algorithm to compute a component of the ℓ-torsion ideal corresponding to a given factor of Ξ_ℓ. Note that in the elliptic case, the classical approach for this task is based on modular forms; we expect that our purely algebraic formulation will enable us to use effective elimination algorithms instead.
- Finally, note that the degree of our modular equation Ξ_ℓ is $(\ell^{2g}-1)/(\ell-1)$, so it is always divisible by $\ell+1$. Finding an explicit Galois action on the roots of Ξ_ℓ having cycles of length $\ell+1$ would yield a modular polynomial of degree $(\ell^{2g}-1)/(\ell^2-1)$. For instance, in genus 2, this would give a polynomial of degree ℓ^2+1 whose rational roots could be used to construct factors of the torsion ideal.

Appendix A. Factorization patterns of Ξ_ℓ in genus 2

The lines in the following table are indexed by the possible decompositions of the Frobenius endomorphism π on $\mathbf{Jac}[\ell]$. These are given in the second column. The first column gives the corresponding factorization pattern of the characteristic polynomial of π, and the last gives the factorization pattern of the modular equation Ξ_ℓ.

The notation for the factorization patterns is defined in Section 5. Notation such as A_i or B_i denotes $i \times i$ matrices with irreducible characteristic polynomials; the notation $*$ denotes any nonzero matrix. Lower case letters stand for scalars, which may be thought as 1×1 matrices.

The modified order $\mathrm{ord}^*(P)$ of a polynomial P is defined in Section 5. Given two matrices A_i, B_j, the notation $\mathrm{ord}(A_iB_j)$ is defined as the modified order of the product of their characteristic polynomials. This notation extends to three or four blocks.

The characteristic polynomial of π is not necessarily squarefree, so we adopt another notation for its factorization pattern:

$$[n_1]^{\alpha_1} \cdots [n_k]^{\alpha_k}$$

stands for a product $P_1^{\alpha_1} \cdots P_k^{\alpha_k}$, where the P_i are distinct irreducible polynomials of degree n_i, for $i = 1, \ldots, k$.

TABLE A.1. Factorization patterns of Ξ_ℓ in genus 2, continued.

Char. pol.	Matrix	Pattern of modular polynomial
	One factor of degree 4	
$[4]$	$\boxed{A_4}$	$(\underbrace{\mathrm{ord}(A_4),\ldots,\mathrm{ord}(A_4)}_{\ell^3+\ell^2+\ell+1})$
	Two factors of degree 2	
$[2]^2$	$\begin{array}{\|c\|c\|}\hline A_2 & 0 \\ \hline 0 & A_2 \\ \hline\end{array}$	$(\underbrace{\mathrm{ord}(A_2),\ldots,\mathrm{ord}(A_2)}_{\ell^3+\ell^2+\ell+1})$
$[2]^2$	$\begin{array}{\|c\|c\|}\hline A_2 & * \\ \hline 0 & A_2 \\ \hline\end{array}$	$(\underbrace{\mathrm{ord}(A_2),\ldots,\mathrm{ord}(A_2)}_{\ell+1}, \underbrace{\mathrm{ord}(A_2^2),\ldots,\mathrm{ord}(A_2^2)}_{\ell^2(\ell+1)})$
$[2][2]$	$\begin{array}{\|c\|c\|}\hline A_2 & 0 \\ \hline 0 & B_2 \\ \hline\end{array}$	$(\underbrace{\mathrm{ord}(A_2),\ldots,\mathrm{ord}(A_2)}_{\ell+1}, \underbrace{\mathrm{ord}(B_2),\ldots,\mathrm{ord}(B_2)}_{\ell+1}, \underbrace{\mathrm{ord}(A_2B_2),\ldots,\mathrm{ord}(A_2B_2)}_{(\ell-1)(\ell+1)^2})$
	One factor of degree 2, two of degree 1	
$[2][1]^2$	$\begin{array}{\|c\|c\|c\|}\hline A_2 & \multicolumn{2}{c\|}{0} \\ \hline 0 & b & 0 \\ \cline{2-3} & 0 & b \\ \hline\end{array}$	$(\underbrace{1,\ldots,1}_{\ell+1}, \underbrace{\mathrm{ord}(A_2),\ldots,\mathrm{ord}(A_2)}_{\ell+1}, \underbrace{\mathrm{ord}(A_2b),\ldots,\mathrm{ord}(A_2b)}_{(\ell-1)(\ell+1)^2})$
$[2][1]^2$	$\begin{array}{\|c\|c\|c\|}\hline A_2 & \multicolumn{2}{c\|}{0} \\ \hline 0 & b & * \\ \cline{2-3} & 0 & b \\ \hline\end{array}$	$(1,\ \ell,\ \underbrace{\mathrm{ord}(A_2),\ldots,\mathrm{ord}(A_2)}_{\ell+1}, \underbrace{\mathrm{ord}(A_2b),\ldots,\mathrm{ord}(A_2b)}_{(\ell-1)(\ell+1)}, \underbrace{\mathrm{ord}(A_2b^2),\ldots,\mathrm{ord}(A_2b^2)}_{\ell(\ell-1)(\ell+1)})$
$[2][1][1]$	$\begin{array}{\|c\|c\|c\|}\hline A_2 & \multicolumn{2}{c\|}{0} \\ \hline 0 & b & 0 \\ \cline{2-3} & 0 & c \\ \hline\end{array}$	$(1,\ 1,\ \underbrace{\mathrm{ord}(A_2),\ldots,\mathrm{ord}(A_2)}_{\ell+1}, \underbrace{\mathrm{ord}(bc),\ldots,\mathrm{ord}(bc)}_{\ell-1}, \underbrace{\mathrm{ord}(A_2b),\ldots,\mathrm{ord}(A_2b)}_{(\ell+1)(\ell-1)},$ $\underbrace{\mathrm{ord}(A_2c),\ldots,\mathrm{ord}(A_2c)}_{(\ell+1)(\ell-1)}, \underbrace{\mathrm{ord}(A_2bc),\ldots,\mathrm{ord}(A_2bc)}_{(\ell-1)^2(\ell+1)})$

TABLE A.2. Factorization patterns of Ξ_ℓ in genus 2.

Char. pol.	Matrix	Pattern of modular polynomial
Four factors of degree 1		
$[1]^4$	$\begin{pmatrix} a & 0 & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & a & 0 \\ 0 & 0 & 0 & a \end{pmatrix}$	$(\underbrace{1, \ldots, 1}_{\ell^3+\ell^2+\ell+1})$
$[1]^4$	$\begin{pmatrix} a & * & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & a & 0 \\ 0 & 0 & 0 & a \end{pmatrix}$	$(\underbrace{1, \ldots, 1}_{\ell^2+\ell+1}, \underbrace{\ell, \ldots, \ell}_{\ell^2})$
$[1]^4$	$\begin{pmatrix} a & * & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & a & * \\ 0 & 0 & 0 & a \end{pmatrix}$	$(\underbrace{1, \ldots, 1}_{\ell+1}, \underbrace{\ell, \ldots, \ell}_{\ell^2+\ell})$
$[1]^4$	$\begin{pmatrix} a & * & * & * \\ 0 & a & * & * \\ 0 & 0 & a & * \\ 0 & 0 & 0 & a \end{pmatrix}$	$(1, \ell, \underbrace{\mathrm{ord}(a^3), \ldots, \mathrm{ord}(a^3)}_{\ell^2}\ \underbrace{\mathrm{ord}(a^4), \ldots, \mathrm{ord}(a^4)}_{\ell^3})$
$[1]^2[1]^2$	$\begin{pmatrix} a & 0 & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & b & 0 \\ 0 & 0 & 0 & b \end{pmatrix}$	$(\underbrace{1, \ldots, 1}_{2\ell+2}, \underbrace{\mathrm{ord}(ab), \ldots, \mathrm{ord}(ab)}_{(\ell-1)(\ell+1)^2})$

TABLE A.3. Factorization patterns of Ξ_ℓ in genus 2, continued.

Char. pol.	Matrix	Pattern of modular polynomial
Four factors of degree 1, continued		
$[1]^2[1]^2$	$\begin{pmatrix} a & * & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & b & 0 \\ 0 & 0 & 0 & b \end{pmatrix}$	$(\underbrace{1, \ldots, 1}_{\ell+2}, \ell, \underbrace{\mathrm{ord}(ab), \ldots, \mathrm{ord}(ab)}_{(\ell-1)(\ell+1)}, \underbrace{\mathrm{ord}(a^2b), \ldots, \mathrm{ord}(a^2b)}_{\ell(\ell-1)(\ell+1)})$
$[1]^2[1]^2$	$\begin{pmatrix} a & * & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & b & * \\ 0 & 0 & 0 & b \end{pmatrix}$	$(1, 1, \ell, \ell, \underbrace{r, \ldots, r}_{\ell-1}, \underbrace{s_1, \ldots, s_1}_{\ell(\ell-1)}, \underbrace{s_2, \ldots, s_2}_{\ell(\ell-1)}, \underbrace{t, \ldots, t}_{\ell^2(\ell-1)})$ where $r = \mathrm{ord}(ab)$, $s_1 = \mathrm{ord}(ab^2)$, $s_2 = \mathrm{ord}(a^2b)$, $t = \mathrm{ord}(a^2b^2)$
$[1]^2[1][1]$	$\begin{pmatrix} a & 0 & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & b & 0 \\ 0 & 0 & 0 & c \end{pmatrix}$	$(\underbrace{1, \ldots, 1}_{\ell+3}, \underbrace{r, \ldots, r}_{\ell-1}, \underbrace{s_1, \ldots, s_1}_{\ell^2-1}, \underbrace{s_2, \ldots, s_2}_{\ell^2-1}, \underbrace{t, \ldots, t}_{(\ell^2-1)(\ell-1)})$ where $r = \mathrm{ord}(bc)$, $s_1 = \mathrm{ord}(ab)$, $s_2 = \mathrm{ord}(ac)$, $t = \mathrm{ord}(abc)$
$[1]^2[1][1]$	$\begin{pmatrix} a & * & 0 & 0 \\ 0 & a & 0 & 0 \\ 0 & 0 & b & 0 \\ 0 & 0 & 0 & c \end{pmatrix}$	$\Big(1, 1, 1, \ell, \underbrace{r_1, \ldots, r_1}_{\ell-1}, \ldots, \underbrace{r_3, \ldots, r_3}_{\ell-1}, \underbrace{s_1, \ldots, s_1}_{\ell(\ell-1)}, \underbrace{s_2, \ldots, s_2}_{\ell(\ell-1)}, \underbrace{t, \ldots, t}_{(\ell-1)^2}, \underbrace{u, \ldots, u}_{\ell(\ell-1)^2}\Big)$ where $r_1 = \mathrm{ord}(bc)$, $r_2 = \mathrm{ord}(ab)$, $r_3 = \mathrm{ord}(ac)$, $s_1 = \mathrm{ord}(a^2b)$, $s_2 = \mathrm{ord}(a^2c)$, $t = \mathrm{ord}(abc)$, $u = \mathrm{ord}(a^2bc)$
$[1][1][1][1]$	$\begin{pmatrix} a & 0 & 0 & 0 \\ 0 & b & 0 & 0 \\ 0 & 0 & c & 0 \\ 0 & 0 & 0 & d \end{pmatrix}$	$\Big(1, 1, 1, 1, \underbrace{r_1, \ldots, r_1}_{\ell-1}, \ldots, \underbrace{r_6, \ldots, r_6}_{\ell-1}, \underbrace{s_1, \ldots, s_1}_{(\ell-1)^2}, \ldots, \underbrace{s_4, \ldots, s_4}_{(\ell-1)^2}, \underbrace{t, \ldots, t}_{(\ell-1)^3}\Big)$ where $r_1 = \mathrm{ord}(ab)$, $r_2 = \mathrm{ord}(ac)$, $r_3 = \mathrm{ord}(ad)$, $r_4 = \mathrm{ord}(bc)$, $r_5 = \mathrm{ord}(bd)$, $r_6 = \mathrm{ord}(cd)$, $s_1 = \mathrm{ord}(abc)$, $s_2 = \mathrm{ord}(abd)$, $s_3 = \mathrm{ord}(acd)$, $s_4 = \mathrm{ord}(bcd)$, $t = \mathrm{ord}(abcd)$

References

1. *Magma*, http://www.maths.usyd.edu.au:8000/u/magma/.
2. *Medicis*, http://www.medicis.polytechnique.fr/.
3. *Tera*, http://tera.medicis.polytechnique.fr/.
4. L. Adleman and M.-D. Huang, *Counting points on curves and abelian varieties over finite fields*, J. Symbolic Comput. **32** (2001), 171–189. MR 2002j:14027
5. A. O. L. Atkin, *The number of points on an elliptic curve modulo a prime*, Series of e-mails to the `NMBRTHRY` mailing list, 1992.
6. I. Blake, G. Seroussi, and N. Smart, *Elliptic curves in cryptography*, London Math. Soc. Lecture Note Ser., vol. 265, Cambridge University Press, 1999. MR 2001i:94048
7. J. Boxall and D. Grant, *Examples of torsion points on genus two curves*, Trans. Amer. Math. Soc. **352** (2000), 4533–4555. MR 2001b:11057
8. B. Buchberger, *Gröbner bases: An algorithmic method in polynomial ideal theory*, Chapter 6, in N. K. Bose et al., Multidimensional Systems Theory, Reidel, Dordrecht, 1985, pp. 374–383. MR 87m:93017
9. D. G. Cantor, *Computing in the Jacobian of a hyperelliptic curve*, Math. Comp. **48** (1987), 95–101. MR 88f:11118
10. ———, *On the analogue of the division polynomials for hyperelliptic curves*, J. Reine Angew. Math. **447** (1994), 91–145. MR 94m:11071
11. L. S. Charlap, R. Coley, and D. P. Robbins, *Enumeration of rational points on elliptic curves over finite fields*, Draft, 1991.
12. N. Elkies, *Elliptic and modular curves over finite fields and related computational issues*, Computational Perspectives on Number Theory (D.A. Buell and J.T. Teitelbaum, eds.), AMS/International Press, 1998, Proceedings of a Conference in Honor of A.O.L. Atkin, pp. 21–76. MR 99a:11078
13. P. Gaudry and R. Harley, *Counting points on hyperelliptic curves over finite fields*, ANTS-IV (W. Bosma, ed.), Lecture Notes in Comput. Sci., vol. 1838, Springer–Verlag, 2000, pp. 313–332. MR 2002f:11072
14. M. Giusti, J. Heintz, J. E. Morais, J. Morgenstern, and L. M. Pardo, *Straight-line programs in geometric elimination theory*, J. of Pure and App. Algebra **124** (1998), 101–146. MR 99d:68128
15. M. Giusti, G. Lecerf, and B. Salvy, *A Gröbner free alternative for polynomial system solving*, J. Complexity **17** (2001), 154–211. MR 2002b:68123
16. J. Heintz, T. Krick, S. Puddu, J. Sabia, and A. Waissbein, *Deformation techniques for efficient polynomial equation solving*, J. Complexity **16** (2000), 70–109. MR 2001c:65067
17. M. Hindry and J. Silverman, *Diophantine geometry. an introduction*, Graduate Texts in Mathematics, vol. 201, Springer–Verlag, 2000. MR 2001e:11058
18. M.-D. Huang and D. Ierardi, *Counting points on curves over finite fields*, J. Symbolic Comput. **25** (1998), 1–21. MR 98i:11040
19. C. Dicrescenzo J. Della Dora and D. Duval, *About a new method for computing in algebraic number fields*, Proceedings Eurocal '85 Vol. 2, Lecture Notes in Comput. Sci., vol. 204, 1985, pp. 289–290.
20. E. Kaltofen and V. Shoup, *Subquadratic-time factoring of polynomials over finite fields*, Math. Comp., **67** (1998), 1179–1197. MR 99m:68097
21. W. Kampkötter, *Explizite Gleichungen für Jacobische Varietäten hyperelliptischer Kurven*, Ph.D. thesis, Universität Gesamthochschule Essen, August 1991.
22. J. S. Milne, *Abelian varieties*, Arithmetic Geometry (G. Cornell and J. H. Silverman, eds.), Springer–Verlag, 1986, pp. 103–150. MR 89b:14029
23. ———, *Jacobian varieties*, Arithmetic Geometry (G. Cornell and J. H. Silverman, eds.), Springer–Verlag, 1986, pp. 167–212. MR 89b:14029
24. F. Morain, *Calcul du nombre de points sur une courbe elliptique dans un corps fini: aspects algorithmiques*, J. Théor. Nombres Bordeaux **7** (1995), 255–282. MR 97i:11069
25. D. Mumford, *Tata lectures on theta I*, Progr. Math., vol. 28, Birkhäuser, 1983. MR 85h:14026
26. J. Pila, *Frobenius maps of abelian varieties and finding roots of unity in finite fields*, Math. Comp. **55** (1990), 745–763. MR 91a:11071
27. R. Schoof, *Elliptic curves over finite fields and the computation of square roots mod p*, Math. Comp. **44** (1985), 483–494. MR 86e:11122

28. ———, *Counting points on elliptic curves over finite fields*, J. Théor. Nombres Bordeaux **7** (1995), 219–254. MR 97i:11070
29. É. Schost, *Complexity results for triangular sets*, J. Symbolic Comput. **36** (2003), 555–594.
30. J. H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics, vol. 106, Springer–Verlag, 1986. MR 87g:11070
31. J. von zur Gathen and J. Gerhard, *Modern computer algebra*, Cambridge University Press, 1999. MR 2000j:68205
32. D. Y. Y. Yun, *On square-free decompositions algorithms*, Proceedings of the 1976 ACM Symposium on Symbolic and Algebraic Computation, ACM Press, 1976, pp. 26–35.

Laboratoire LIX, École polytechnique, 91128 Palaiseau, France
E-mail address: gaudry@lix.polytechnique.fr

Laboratoire STIX, École polytechnique, 91128 Palaiseau, France
E-mail address: schost@stix.polytechnique.fr