

Computing cardinalities of Q -curve reductions over finite fields

François Morain, Charlotte Scribot, Benjamin Smith

► **To cite this version:**

François Morain, Charlotte Scribot, Benjamin Smith. Computing cardinalities of Q -curve reductions over finite fields. *LMS Journal of Computation and Mathematics*, London Mathematical Society, 2016, 19 (A), pp.15. <10.1112/S1461157016000267>. <hal-01320388v3>

HAL Id: hal-01320388

<https://hal.inria.fr/hal-01320388v3>

Submitted on 17 Jun 2016

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.

Copyright

Computing cardinalities of \mathbb{Q} -curve reductions over finite fields

François Morain, Charlotte Scribot, and Benjamin Smith

ABSTRACT

We present a specialized point-counting algorithm for a class of elliptic curves over \mathbb{F}_{p^2} that includes reductions of quadratic \mathbb{Q} -curves modulo inert primes and, more generally, any elliptic curve over \mathbb{F}_{p^2} with a low-degree isogeny to its Galois conjugate curve. These curves have interesting cryptographic applications. Our algorithm is a variant of the Schoof–Elkies–Atkin (SEA) algorithm, but with a new, lower-degree endomorphism in place of Frobenius. While it has the same asymptotic complexity as SEA, our algorithm is much faster in practice.

1. Introduction

Computing the cardinalities of the groups of rational points on elliptic curves over finite fields is a fundamental algorithmic challenge in computational number theory, and an essential tool in elliptic curve cryptography. Over finite fields of large characteristic, the best known algorithm is the Schoof–Elkies–Atkin (SEA) algorithm [20]. A lot of work has been put into optimizing the computations for prime fields of large characteristic (see [26] for the most recent record). Many of these improvements carry over to the case of more general finite fields. In this article we define a specialized, faster SEA algorithm for a class of elliptic curves over \mathbb{F}_{p^2} that have useful cryptographic applications. These curves have low-degree inseparable endomorphisms that can be used to accelerate scalar multiplication in elliptic curve cryptosystems [23, 24]; here, we use those endomorphisms to accelerate point counting. Going beyond cryptography, this class of curves also includes reductions of quadratic \mathbb{Q} -curves modulo inert primes, so our algorithm may be useful for studying these curves.

Let q be a power of a prime $p > 3$ (in our applications, $q = p^2$ and p is large). Let

$$\sigma(\cdot) : x \mapsto x^p$$

be the Frobenius automorphism on \mathbb{F}_q . We extend the action of Frobenius to polynomials over \mathbb{F}_q by acting on coefficients, and thus to curves over \mathbb{F}_q by acting on their defining equations: for example, an elliptic curve \mathcal{E}/\mathbb{F}_q and its Galois conjugate curve ${}^\sigma\mathcal{E}/\mathbb{F}_q$ would be defined by

$$\mathcal{E} : y^2 = x^3 + Ax + B \quad \text{and} \quad {}^\sigma\mathcal{E} : y^2 = x^3 + A^p x + B^p .$$

If \mathcal{E}/\mathbb{F}_q is an elliptic curve, then there is a p -isogeny $\pi_p : \mathcal{E} \rightarrow {}^\sigma\mathcal{E}$ defined by $\pi_p : (x, y) \mapsto (x^p, y^p)$. If $q = p^n$, then composing $\pi_p, {}^\sigma\pi_p, \dots, {}^{\sigma^{n-1}}\pi_p$ yields the Frobenius endomorphism $\pi_q : (x, y) \mapsto (x^q, y^q)$ of \mathcal{E} . Being an endomorphism, π_q has a characteristic polynomial

$$\chi_{\pi_q}(T) = T^2 - t_{\mathcal{E}}T + q$$

such that $\chi_{\pi_q}(\pi_q) = [0]$ in $\text{End}(\mathcal{E})$; the trace $t_{\mathcal{E}}$ satisfies the Hasse bound

$$|t_{\mathcal{E}}| \leq 2\sqrt{q} .$$

Knowing the cardinality of $\mathcal{E}(\mathbb{F}_q)$ is equivalent to knowing the trace $t_{\mathcal{E}}$, since

$$\#\mathcal{E}(\mathbb{F}_q) = \chi_{\pi_q}(1) = q + 1 - t_{\mathcal{E}} .$$

Schoof’s point counting algorithm [18] determines $t_\ell := t_\mathcal{E} \pmod{\ell}$ for small primes $\ell \neq p$ by examining the action of π_q on $\mathcal{E}[\ell]$, the ℓ -torsion subscheme of \mathcal{E} : we have

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

If we construct a general P as detailed in §2, then finding t_ℓ boils down to a series of polynomial operations modulo the ℓ -th division polynomial Ψ_ℓ . Schoof’s algorithm tests these relations until $\prod \ell > 4\sqrt{q}$, and then deduces $t_\mathcal{E}$ from the t_ℓ using the Chinese Remainder Theorem (CRT). To completely determine $t_\mathcal{E}$ we need to compute t_ℓ for $O(\log q)$ primes ℓ , the largest of which is in $O(\log q)$; fast polynomial evaluations add some more $O(\log q)$ factors, and the final cost is $O(\log^8 q)$ with classical arithmetic (or $O(\log^6 q)$ with fast arithmetic). This basic algorithm was subsequently improved by Atkin and Elkies; the resulting SEA algorithm (see §4) is now the standard point-counting algorithm for elliptic curves over large characteristic fields.

In this article, we present an algorithm that was designed to compute $\#\mathcal{E}(\mathbb{F}_{p^2})$ when \mathcal{E} is the reduction of a low-degree quadratic \mathbb{Q} -curve modulo an inert prime. In fact, our algorithm applies to a larger class of curves over finite fields, which we will call *admissible curves*.

First, recall that every d -isogeny $\vartheta : \mathcal{E} \rightarrow \mathcal{E}'$ has a *dual* d -isogeny $\vartheta^\dagger : \mathcal{E}' \rightarrow \mathcal{E}$ such that $\vartheta^\dagger \vartheta = [d]_\mathcal{E}$ and $\vartheta \vartheta^\dagger = [d]_{\mathcal{E}'}$. Also, σ acts on isogenies by p -th powering the coefficients of their defining polynomials; so every isogeny $\vartheta : \mathcal{E} \rightarrow \mathcal{E}'$ has a Galois conjugate isogeny $\sigma \vartheta : \sigma \mathcal{E} \rightarrow \sigma \mathcal{E}'$.

DEFINITION 1. Let d be a squarefree integer with $p \nmid d$. An elliptic curve $\mathcal{E}/\mathbb{F}_{p^2}$ is *d -admissible* if it is equipped with a d -isogeny

$$\phi : \mathcal{E} \longrightarrow \sigma \mathcal{E} \quad \text{such that} \quad \sigma \phi = \epsilon \phi^\dagger \quad \text{where} \quad \epsilon = \pm 1.$$

Composing $\pi_p : \mathcal{E} \rightarrow \sigma \mathcal{E}$ with $\sigma \phi : \sigma \mathcal{E} \rightarrow \mathcal{E}$, we obtain the *associated endomorphism*

$$\psi := \sigma \phi \circ \pi_p \in \text{End}(\mathcal{E})$$

of degree dp . Note that the requirement $p \nmid d$ implies that both ϕ and $\sigma \phi$ are separable.

We are particularly interested in curves that are d -admissible for small values of d . When d is extremely small the associated endomorphism can be evaluated very efficiently, and thus used to accelerate scalar multiplication on \mathcal{E} for more efficient implementations of elliptic curve cryptosystems (as in [23], [10], [3], [24], and [4]). Constructing cryptographically secure curves equipped with efficient endomorphisms is one major motivation for our algorithm; the other is the principle that the presence of special structures demands the use of a specialized algorithm.

From a practical point of view, suitable modifications of the SEA algorithm gives us a very fast probabilistic solution to the point counting problem for admissible curves. The essential idea is to use SEA with the associated endomorphism ψ in place of π_q . While the asymptotic complexity of our algorithm is the same as for the unmodified SEA algorithm when d is fixed, there are some important improvements in the big- O constants. Asymptotically, when d is small, our algorithm runs four times faster than SEA (and even faster for smaller p).

It is not hard to see that of the p^2 isomorphism classes of elliptic curves over \mathbb{F}_{p^2} , only $O(p)$ classes correspond to d -admissible curves for any fixed d . But while d -admissible curves with small d may be relatively rare, they appear naturally “in the wild” as reductions of quadratic \mathbb{Q} -curves of degree d (elliptic curves over quadratic number fields that are d -isogenous to their Galois conjugates) modulo inert primes. For some small d , these \mathbb{Q} -curves occur in one-parameter families; so our algorithm allows the reductions of these families modulo suitable primes to be rapidly searched for cryptographic curves. We explain this further in §9.

2. Computing with isogenies

We begin by recalling some standard results on isogenies, fixing notation and basic complexities in the process. A classical reference for all this is [7].

First, let $M(n)$ denote the cost in \mathbb{F}_q -operations (multiplications) of multiplying two polynomials of degree n . Traditional multiplication gives $M(n) = O(n^2)$; fast multiplication gives $\tilde{O}(n)$. Dividing a degree- $2n$ polynomial by a degree- n polynomial costs $O(M(n))$ \mathbb{F}_q -operations; the extended GCD of two degree- n polynomials can be computed in $O(M(n) \log n)$ \mathbb{F}_q -operations. The number of roots in \mathbb{F}_q of a degree- n polynomial F over \mathbb{F}_q is equal to $\deg \text{GCD}(x^q - x, F(x))$, which we can compute in $O((\log q)M(n))$ \mathbb{F}_q -operations if $n \ll q$ (this is dominated by the cost of computing $x^q \bmod F$; see Appendix A).

We will make extensive use of *modular composition*: if F , G , and H are polynomials over \mathbb{F}_q with $\deg F = n$, $\deg G < n$, and $\deg H < n$, then we can compute $(G \circ H) \bmod F$ in $O(n^{1/2}M(n) + n^{(\omega+1)/2})$ \mathbb{F}_q -operations, where $2 \leq \omega \leq 3$ is the constant for linear algebra. Using the method of [13], the cost in \mathbb{F}_q -operations of performing r modular compositions with the same H and F is

$$C_r(n) := O(r^{1/2}n^{1/2}M(n) + r^{(\omega-1)/2}n^{(\omega+1)/2}) .$$

We will always work with elliptic curves \mathcal{E}/\mathbb{F}_q using their Weierstrass models,

$$\mathcal{E} : y^2 = f_{\mathcal{E}}(x), \text{ where } f_{\mathcal{E}} \text{ is a monic cubic over } \mathbb{F}_q .$$

For $m > 0$, the m -th *division polynomial* $\Psi_m(x)$ is the polynomial in $\mathbb{F}_q[x]$ whose roots are precisely the x -coordinates of the points in $\mathcal{E}[m](\overline{\mathbb{F}_q})$.

If ℓ is a prime, then the level- ℓ modular polynomial $\Phi_{\ell}(J_1, J_2)$ has degree $\ell + 1$ in both J_1 and J_2 , and is defined over \mathbb{Z} . If $\Phi_{\ell}(j_1, j_2) = 0$ for some j_1 and j_2 in \mathbb{F}_q , then there is an \mathbb{F}_q -rational ℓ -isogeny between the curves with j -invariants j_1 and j_2 (possibly after a twist). In particular, if we fix an elliptic curve \mathcal{E}/\mathbb{F}_q , then the roots of $\Phi_{\ell}(j(\mathcal{E}), x)$ in \mathbb{F}_q correspond to (the isomorphism classes of) the curves that are ℓ -isogenous to \mathcal{E} over \mathbb{F}_q .

We will need explicit forms for d -isogenies where d is squarefree and prime to p . Every such isogeny can be expressed as a composition of at most one 2-isogeny with at most one odd-degree cyclic isogeny over \mathbb{F}_q . If ϑ is a 2-isogeny, then it is defined by a rational map

$$\vartheta : (x, y) \mapsto \left(\frac{N(x)}{D(x)}, y \frac{M(x)}{D^2(x)} \right) \tag{2.1}$$

where N , M , and D are polynomials over \mathbb{F}_q with $\deg N = \deg M = 2$ and $D = x - x_0$ where x_0 is the abscissa of a 2-torsion point. If ϑ is a d -isogeny where d is odd, squarefree, and prime to p , then ϑ is defined by a rational map

$$\vartheta : (x, y) \mapsto \left(\frac{N(x)}{D^2(x)}, y \frac{M(x)}{D^3(x)} \right) \tag{2.2}$$

where N , M , and D are polynomials over \mathbb{F}_q with $\deg N = \deg M = d$ and $\deg D = (d - 1)/2$.

In both cases, the polynomial $D(x)$ cuts out the kernel of ϑ , in the sense that $D(x(P)) = 0$ if and only if P is a nontrivial element of $\ker \vartheta$; we call D the *kernel polynomial* of ϑ . We suppose we have a subroutine `KERNELPOLYNOMIAL`(ℓ, \mathcal{E}, j_1) which, given \mathcal{E} and $j_1 = j(\mathcal{E}_1)$ such that there exists an ℓ -isogeny $\vartheta : \mathcal{E} \rightarrow \mathcal{E}_1$ over \mathbb{F}_q , computes the kernel polynomial D of ϑ and the isogenous curve \mathcal{E}_1 in $O(\ell^2)$ \mathbb{F}_q -operations (using the fast algorithms in [1]).

The algorithms in this article examine the actions of endomorphisms on $\ker \vartheta$, where ϑ is either $[\ell]$ or an ℓ -isogeny, for a series of small primes ℓ . The key is to define a symbolic element of $\ker \vartheta$. First, we compute the kernel polynomial D of ϑ (note that $D = \Psi_{\ell}$ if $\vartheta = [\ell]$); then, we can construct a symbolic point P of $\ker \vartheta$ as

$$P := (X, Y) \in \mathcal{E} (\mathbb{F}_q[X, Y]/(Y^2 - f_{\mathcal{E}}(X), D(X))) .$$

We reduce the coordinates of points in $\langle P \rangle$ modulo $D(X)$ and $Y^2 - f_{\mathcal{E}}(X)$ after each operation, so elements of $\langle P \rangle$ have a canonical form $Q = (Q_x(X), YQ_y(X))$ with $\deg Q_x, \deg Q_y < \deg D$.

Let $e = \deg D$; then we can compute $Q_1 + Q_2$ for any Q_1 and Q_2 in $\langle P \rangle$ in $O(M(e) \log e)$ \mathbb{F}_q -operations, using the standard affine Weierstrass addition formulæ. We can therefore compute $[m]Q$ for any m in \mathbb{Z} and Q in $\langle P \rangle$ in $O((\log m)M(e) \log e)$ \mathbb{F}_q -operations, using a binary method. We let $\text{DISCRETELOGARITHM}(Q_1, Q_2)$ be a subroutine which returns the discrete logarithm of Q_2 to the base Q_1 , where both points are in $\langle P \rangle$, in $O(\sqrt{e}M(e))$ \mathbb{F}_q -operations (using the approach in [8]; in some cases we can do better [15]).

LEMMA 1. *Let $P = (X, Y)$ in $\mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_{\mathcal{E}}(X), D(X)))$, and let $e = \deg D$. Then for any Q in $\langle P \rangle$, we can*

- (i) *compute $\pi_p(P) = (X^p, Y^p)$ in $O((\log p)M(e))$ \mathbb{F}_q -operations;*
- (ii) *compute $\pi_p(Q)$, given $\pi_p(P)$, in $O((\log p)M(e))$ \mathbb{F}_q -operations;*
- (iii) *compute $\phi(Q)$, where ϕ is a 2-isogeny (as in (2.1)) in $O(M(e))$ \mathbb{F}_q -operations;*
- (iv) *compute $\phi(Q)$, where ϕ is a d -isogeny with d odd, squarefree, and prime to p (as in (2.2)) in $O(M(d) + \mathcal{C}_3(e))$ \mathbb{F}_q -operations.*

Proof. See Appendix A. □

3. Atkin, Elkies, and volcanic primes

Given an elliptic curve \mathcal{E}/\mathbb{F}_q , we split the primes $\ell \neq p$ into three classes: Elkies, Atkin, and volcanic. The volcanic primes fall in two sub-classes: floor-volcanic and upper-volcanic. This classification reflects the structure of the ℓ -isogeny graph near \mathcal{E} , which reflects the factorization of $\Phi_{\ell}(j(\mathcal{E}), x)$. The facts stated below without proof all follow immediately from well-known observations of Atkin for general ordinary elliptic curves over \mathbb{F}_q (cf. [20, Prop. 6.2]).

Recall that the discriminant of χ_{π_q} is $\Delta_{\pi_q} := t_{\mathcal{E}}^2 - 4q < 0$. We say that ℓ is **volcanic** if ℓ divides Δ_{π_q} . A volcanic prime ℓ is **floor-volcanic** if

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

where h is an \mathbb{F}_q -irreducible polynomial of degree ℓ , or **upper-volcanic** if

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

with each j_i in \mathbb{F}_q . In each case, the roots j_i are the j -invariants of the elliptic curves \mathcal{E}_i that are ℓ -isogenous to \mathcal{E} over \mathbb{F}_q (up to isomorphism).

We say that ℓ is **Elkies** if Δ_{π_q} is a nonzero square modulo ℓ . Equivalently, ℓ is Elkies if

$$\Phi_{\ell}(x, j(\mathcal{E})) = (x - j_1)(x - j_2) \prod_{i=1}^{(\ell-1)/e} h_i(x) ,$$

where j_1 and j_2 are in \mathbb{F}_q and the h_i are \mathbb{F}_q -irreducible polynomials, all of the same degree $e > 1$, with $e \mid (\ell - 1)$. In this case, there exist \mathbb{F}_q -rational ℓ -isogenies $\vartheta_1 : \mathcal{E} \rightarrow \mathcal{E}_1$ and $\vartheta_2 : \mathcal{E} \rightarrow \mathcal{E}_2$ such that $j(\mathcal{E}_i) = j_i$, and the ℓ -torsion decomposes as $\mathcal{E}[\ell] = \ker \vartheta_1 \oplus \ker \vartheta_2$.

We say that ℓ is **Atkin** if Δ_{π_q} is not a square modulo ℓ . Equivalently, ℓ is Atkin if

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

where the h_i are all irreducible polynomials of the same degree $e > 1$, with $e \mid (\ell + 1)$. Since $\Phi_{\ell}(x, j(\mathcal{E}))$ has no roots in \mathbb{F}_q , there are no elliptic curves ℓ -isogenous to \mathcal{E} over \mathbb{F}_q .

We can determine the class of a prime ℓ by finding out how many roots $\Phi_\ell(j(\mathcal{E}), x)$ has in \mathbb{F}_q . We define a subroutine `EVALUATEDMODULARPOLYNOMIAL`(ℓ, \mathcal{E}), which computes $\Phi_\ell(j(\mathcal{E}), x)$ in $O(\ell^3(\log \ell)^3 \log \log \ell)$ bit operations (under the GRH) using the method of [26], assuming $\log q = \Theta(\ell)$. (Note that in practice, one generally uses precomputed modular polynomials over \mathbb{Z} .)

The number of roots is the degree of $J = \text{GCD}(x^q - x, \Phi_\ell(j(\mathcal{E}), x))$, which we compute at a further cost of $O((\log q)\mathbf{M}(\ell))$ \mathbb{F}_q -operations. We may then want one of these roots, if any exist; we therefore define a subroutine `ONEROOT`(J) which finds a single root of J . At worst, in the upper-volcanic case, this requires $O((\log q)\mathbf{M}(\deg J) \log \deg J) = O((\log q)\mathbf{M}(\ell) \log \ell)$ \mathbb{F}_q -operations; at best, in the lower-volcanic and Elkies cases (where J is linear and quadratic, respectively), `ONEROOT`(J) costs $O(1)$ \mathbb{F}_q -operations.

4. The SEA algorithm

Algorithm 1 presents a basic version of the SEA algorithm. The main loop computes $t_\ell := t_\mathcal{E} \pmod{\ell}$ for a series of small primes ℓ ; then we recover $t_\mathcal{E}$ from the t_ℓ via the CRT.

The complexity of Algorithm 1 (and Algorithm 2 below) depends on the number of non-Atkin primes less than a given bound. The standard (and naïve) heuristic on prime classes is to suppose that the number of Atkin and non-Atkin primes ℓ less than B for a given \mathcal{E}/\mathbb{F}_q is approximately equal when $B \sim \log q$, as $q \rightarrow \infty$. In particular, this means that $O(\log q)$ non-Atkin ℓ suffice to determine $t_\mathcal{E}$, and the largest such ℓ is in $O(\log q)$. While the standard heuristic holds on the average, it is known to fail for some \mathcal{E} ; Galbraith and Satoh have shown (under the GRH) that for some \mathcal{E}/\mathbb{F}_p we may need to use non-Atkin ℓ as large as $O(\log^{2+\epsilon} p)$ (see [17, App. A]). We refer the reader to [21] and [22] for further details and discussion.

PROPOSITION 2. *If \mathcal{E}/\mathbb{F}_q is an elliptic curve, then under the standard heuristic on prime classes, Algorithm 1 computes $t_\mathcal{E}$ in $\tilde{O}(\log^3 q)$ expected \mathbb{F}_q -operations (that is $\tilde{O}(\log^4 q)$ expected bit operations, using fast arithmetic).*

Proof. The main loop computes a set \mathcal{T} of pairs $(t_\ell := t_\mathcal{E} \pmod{\ell}, \ell)$ with $\prod_{\mathcal{T}} \ell > 4\sqrt{q}$. We then recover $t_\mathcal{E}$ from \mathcal{T} via an explicit CRT. Our procedure for computing t_ℓ depends on the class of ℓ , which we determine using the method at the end of §3 (Lines 6, 7, and 15).

If ℓ is volcanic (Lines 9 to 14), then $\ell \mid \Delta_{\pi_q}$, so $t_\ell = 0$ or $t_\ell \equiv \pm 2\sqrt{q} \pmod{\ell}$. We distinguish between the three cases by comparing $\pi_q(P)$ with $\pm[\sqrt{q} \pmod{\ell}]P$ for a generic element P of the kernel of the rational ℓ -isogeny corresponding to one of the roots of $\Phi_\ell(j(\mathcal{E}), x)$.

If ℓ is Elkies (Lines 16 to 20), then $\mathcal{E}[\ell]$ decomposes as a direct sum $(\ker \vartheta_1) \oplus (\ker \vartheta_2)$ of ℓ -isogeny kernels; π_q acts as multiplication by eigenvalues λ_1 and λ_2 on $\ker \vartheta_1$ and $\ker \vartheta_2$, respectively, with $\lambda_1 \lambda_2 \equiv q \pmod{\ell}$, so $t_\ell \equiv \lambda_1 + q/\lambda_1 \pmod{\ell}$; and we can determine λ_1 by solving the discrete logarithm problem $\pi_q(P) = [\lambda_1]P$ for a symbolic point P of $\ker \vartheta_1$.

If ℓ is Atkin, then we skip it completely and do not compute t_ℓ (see the discussion in §7).

In terms of \mathbb{F}_q -operations, determining the class of ℓ costs $O(\ell^2(\log \ell)^3 \log \log \ell + \log q \mathbf{M}(\ell))$; computing t_ℓ then costs $O((\log q + \log \ell)\mathbf{M}(\ell) \log \ell + \ell^{(\omega+1)/2})$ for volcanic ℓ , and $O((\log p + \ell^{1/2})\mathbf{M}(\ell) + \ell^{(\omega+1)/2})$ for Elkies ℓ . The standard heuristic on prime classes tells us that we will try $O(\log q)$ primes ℓ , and that the largest ℓ are in $O(\log q)$; so the total cost of the algorithm is $\tilde{O}(\log^3 q)$, as claimed. \square

Algorithm 1: SEATRACE

Input: An elliptic curve \mathcal{E}/\mathbb{F}_q , where $q = p^n$ with p large
Output: The trace of Frobenius of \mathcal{E}

```

1  $\mathcal{T} \leftarrow \{\}$  ; //  $\mathcal{T}$  will contain the pairs  $(t_{\mathcal{E}} \pmod{\ell}, \ell)$ 
2  $M \leftarrow 1$  ; // After each iteration,  $t_{\mathcal{E}}$  is known modulo  $M$ 
3  $\ell \leftarrow 1$  ;
4 while  $M \leq 4\sqrt{q}$  do
5    $\ell \leftarrow \text{NEXTPRIME}(\ell)$  ;
6    $J \leftarrow \text{GCD}(x^q - x, \text{EVALUATEDMODULARPOLYNOMIAL}(\ell, \mathcal{E}))$  ;
7   if  $\deg J = 1$  or  $\ell + 1$  then //  $\ell$  is volcanic
8     if  $q$  has a square root  $s$  modulo  $\ell$  then //  $s = p$  for  $q = p^2$ 
9        $F \leftarrow \text{KERNELPOLYNOMIAL}(\ell, \mathcal{E}, \text{ONEROOT}(J))$  ;
10       $P \leftarrow (X, Y)$  in  $\mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_{\mathcal{E}}(X), F(X)))$  ;
11       $Q_1 \leftarrow \pi_p(P)$  ;  $Q_2 \leftarrow \pi_p(Q_1)$  ;  $Q_3 \leftarrow [s]P$  ;
12       $t_{\ell} \leftarrow \begin{cases} -2s & \text{if } Q_2 = Q_3 ; \\ 2s & \text{if } Q_2 = -Q_3 ; \\ 0 & \text{otherwise ;} \end{cases}$ 
13    else  $t_{\ell} \leftarrow 0$  ;
14     $\mathcal{T} \leftarrow \mathcal{T} \cup \{(t_{\ell}, \ell)\}$  ;  $M \leftarrow \ell M$  ;
15  else if  $\deg J = 2$  then //  $\ell$  is Elkies
16     $F \leftarrow \text{KERNELPOLYNOMIAL}(\ell, \mathcal{E}, \text{ONEROOT}(J))$  ;
17     $P \leftarrow (X, Y)$  in  $\mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_{\mathcal{E}}(X), F(X)))$  ;
18     $Q_1 \leftarrow \pi_p(P)$  ;  $Q_2 \leftarrow \pi_p(Q_1)$  ;
19     $t_{\ell} \leftarrow \lambda + q/\lambda \pmod{\ell}$  where  $\lambda = \text{DISCRETELOGARITHM}(P, Q_2)$  ;
20     $\mathcal{T} \leftarrow \mathcal{T} \cup \{(t_{\ell}, \ell)\}$  ;  $M \leftarrow \ell M$  ;
21 return  $\text{CHINESEREMAINDERTHEOREM}(\mathcal{T})$  ;
```

5. Admissible curves

From now on, $q = p^2$.

Recalling Definition 1: let \mathcal{E} be a d -admissible curve over \mathbb{F}_{p^2} , with separable d -isogeny $\phi : \mathcal{E} \rightarrow {}^{\sigma}\mathcal{E}$ (satisfying $\sigma\phi = \epsilon\phi^{\dagger}$ with $\epsilon = \pm 1$), and associated endomorphism $\psi = \sigma\phi \circ \pi_p$.

PROPOSITION 3. *The associated and Frobenius endomorphisms of \mathcal{E} are related by*

$$\psi^2 = [\epsilon d]\pi_{p^2} . \quad (5.1)$$

The characteristic polynomial of ψ is

$$\chi_{\psi}(T) = T^2 - rT + dp , \quad (5.2)$$

where r is an integer satisfying

$$dr^2 = 2p + \epsilon t_{\mathcal{E}} . \quad (5.3)$$

In particular,

$$r\psi = p + \epsilon\pi_{p^2} \quad \text{in } \text{End}(\mathcal{E}) . \quad (5.4)$$

Proof. Equation (5.1) holds because $\psi^2 = (\sigma\phi\pi_p)(\sigma\phi\pi_p) = (\epsilon\phi^{\dagger}\phi)(\sigma\pi_p\pi_p) = [\epsilon d]\pi_q$. The degree of ψ is dp , so ψ has characteristic polynomial $\chi_{\psi}(T) = T^2 - xT + dp$ for some integer x .

On the other hand, $\epsilon d\pi_q$ has characteristic polynomial $T^2 - \epsilon dt_{\mathcal{E}}T + d^2p^2$; but $\psi^2 = x\psi - dp$ is a root, so $x = rd$ where r satisfies (5.3). We then have $\epsilon dr^2\pi_q = (\epsilon\pi_q + p)^2$ in $\mathbb{Z}[\pi_q]$. Comparing with (5.1), we find $r\psi = \pm(p + \epsilon\pi_q)$; but then $\chi_{\psi}(\psi) = 0$ implies (5.4). \square

Equation (5.3) has a number of interesting corollaries. First, $t_{\mathcal{E}} \equiv -\epsilon 2p \pmod{d}$, so we obtain some information on $t_{\mathcal{E}}$ for free. Second, r determines $t_{\mathcal{E}}$, and hence $\#\mathcal{E}(\mathbb{F}_{p^2})$. Third, we have a much smaller bound on r than on $t_{\mathcal{E}}$: for d -admissible curves the Hasse–Weil bound becomes

$$|r| \leq 2\sqrt{p/d}. \tag{5.5}$$

This suggests our point-counting strategy, which is to modify the SEA algorithm to compute r instead of $t_{\mathcal{E}}$, by considering the action on $\mathcal{E}[\ell]$ of ψ instead of π_q and using fewer primes ℓ .

We simplify the task by quickly disposing of the supersingular case, which can be efficiently detected using Sutherland’s algorithm [25], or slightly faster using a probabilistic algorithm.

PROPOSITION 4. *If $\mathcal{E}/\mathbb{F}_{p^2}$ is d -admissible, then it is supersingular if and only if $r = 0$, in which case $t_{\mathcal{E}} = -2\epsilon p$ and $\mathcal{E}(\mathbb{F}_{p^2}) \cong (\mathbb{Z}/(p + \epsilon)\mathbb{Z})^2$.*

Proof. The curve \mathcal{E} is supersingular if and only if $p \mid t_{\mathcal{E}}$, if and only if $p \mid r$ (by (5.3) mod p and $p \nmid d$), if and only if $r = 0$ (by (5.5)). The group structure follows from [27, Th. 1.1]. \square

From now on, we will **assume \mathcal{E} is ordinary**; so $\text{End}(\mathcal{E})$ is an order in the quadratic imaginary field $\mathbb{Q}(\pi_q)$, and $\mathbb{Z}[\pi_q]$ and $\mathbb{Z}[\psi]$ are orders contained in $\text{End}(\mathcal{E})$. Looking at (5.2), we see that the discriminants of $\mathbb{Z}[\psi]$ and $\mathbb{Z}[\pi_q]$ are related by

$$\Delta_{\psi} = d(dr^2 - 4p) \quad \text{and} \quad \Delta_{\pi_q} = t_{\mathcal{E}}^2 - 4p^2 = r^2\Delta_{\psi},$$

so $|r|$ is the conductor of $\mathbb{Z}[\pi_q]$ in $\mathbb{Z}[\psi]$: that is,

$$\mathbb{Z}[\pi_q] \subset \mathbb{Z}[\psi] \subseteq \text{End}(\mathcal{E}) \quad \text{with} \quad [\mathbb{Z}[\psi] : \mathbb{Z}[\pi_q]] = |r|.$$

Indeed, since \mathcal{E} is ordinary, we have $r \neq 0$; so we can rewrite (5.4) as

$$\psi = \frac{p + \epsilon\pi_q}{r} \quad \text{in } \text{End}(\mathcal{E}). \tag{5.6}$$

Deuring’s theorem on isogeny classes and class groups (cf. [19, §4]) can be used to show that the number of \mathbb{F}_q -isomorphism classes of ordinary d -admissible curves with a given r is $H(\Delta_{\psi})$, where H is the Kronecker class number. In particular, every r in the interval of (5.5) occurs for some d -admissible \mathcal{E}/\mathbb{F}_q .

In the language of isogeny volcanoes [6]: if ℓ is a prime dividing r , then \mathcal{E} is somewhere *strictly above* the floor of the volcano for ℓ ; that is, all $\ell \mid r$ are upper-volcanic.

6. Computing the cardinality of admissible curves

Let \mathcal{E}/\mathbb{F}_q be an ordinary d -admissible curve, with associated endomorphism ψ ; we want to compute $\#\mathcal{E}(\mathbb{F}_q)$. Many of the techniques used in the conventional SEA algorithm can be transposed to working with ψ instead of π_q . Equations (5.3) and (5.5) show that $t_{\mathcal{E}}$ is completely determined by $|r|$, which is bounded by $2\sqrt{p/d}$; so we can compute $t_{\mathcal{E}}$ by computing

$$r_{\ell} := r \pmod{\ell}$$

for ℓ in a collection of small primes \mathcal{L} such that

$$\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{p/d},$$

then recovering r from the r_ℓ using the CRT. As a quick comparison, using SEA with π_q to compute $t_\mathcal{E}$ directly would require $\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q} = 4p$.

PROPOSITION 5. *If $\mathcal{E}/\mathbb{F}_{p^2}$ is d -admissible, then under the standard heuristic on prime classes, Algorithm 2 computes $t_\mathcal{E}$ in $\tilde{O}(\log^3 p)$ expected \mathbb{F}_q -operations (that is $\tilde{O}(\log^4 p)$ expected bit operations, using fast arithmetic).*

Proof. We compute $t_\mathcal{E}$ from r , which we recover exactly using the CRT from the pairs (r_ℓ, ℓ) in \mathcal{R} , since $\prod_{(r_\ell, \ell) \in \mathcal{R}} \ell > 4\sqrt{p/d}$. Our approach for computing r_ℓ depends on which class ℓ falls into; we determine the class of ℓ in Lines 6, 7, and 15 (exactly as in Algorithm 1).

If ℓ is volcanic (Lines 9 to 14), then combining $\ell \mid \Delta_{\pi_q}$ with (5.3) yields $r \equiv 0$ or $\pm 2\sqrt{p/d} \pmod{\ell}$; in particular, if ℓ is volcanic and dp is a nonsquare modulo ℓ , then $r_\ell = 0$.

If ℓ is Elkies (Lines 16 to 20), then let $\mathcal{E}[\ell] = (\ker \vartheta_1) \oplus (\ker \vartheta_2)$ be the decomposition of the ℓ -torsion into eigenspaces for π_q . Since ℓ is not volcanic we have $r \not\equiv 0 \pmod{\ell}$, so (5.6) shows that the $\ker \vartheta_i$ are also eigenspaces for ψ . So let λ_π and λ_ψ be the eigenvalues of π_q and ψ on $\ker \vartheta_1$ (say); then (5.6) yields $\lambda_\psi \equiv (p + \epsilon\lambda_\pi)/r \pmod{\ell}$, and then $\chi_\psi(\lambda_\psi) \equiv 0 \pmod{\ell}$ implies $r_\ell \equiv \frac{\lambda_\psi}{d} + \frac{p}{\lambda_\psi} \pmod{\ell}$. We can therefore compute r_ℓ by computing λ_ψ , which is the discrete logarithm of $\psi(P)$ to the base P for a symbolic point P in $\ker \vartheta_1$.

If ℓ is Atkin then we skip it completely, as in Algorithm 1 (but see §7).

In terms of \mathbb{F}_q -operations, determining the class of ℓ costs $O(\ell^2(\log \ell)^3 \log \log \ell + \log q M(\ell))$, while computing r_ℓ costs $O((\log p + \log \ell)M(\ell) \log \ell + \ell^{(\omega+1)/2})$ if ℓ is volcanic, and $O((\log p + \ell^{1/2})M(\ell) + \ell^{(\omega+1)/2})$ if ℓ is Elkies. The standard heuristic on prime classes tell us that we will try $O(\log p)$ primes ℓ , the largest of which are in $O(\log p)$; so the total complexity is $\tilde{O}(\log^3 p)$ \mathbb{F}_q -operations, as claimed. \square

REMARK 1. Suppose $\ell \mid d$ and $\ell \neq 2$. Equation (5.3) tells us that $t_\mathcal{E} \equiv 2\epsilon p \pmod{\ell}$; so $\ell \mid \Delta_{\pi_q}$, and ℓ is volcanic. Moreover, since $\Delta_\psi = d(dr^2 - 4p)$, we can deduce that $\ell \parallel \Delta_\psi$. Note also that $\text{End}(\mathcal{E}) \cong \text{End}(\sigma \mathcal{E})$, so the ℓ -isogeny factoring ϕ is horizontal; this implies that $\text{End}(\mathcal{E})$ is ℓ -maximal. Combined with the above, we see that $\mathbb{Z}[\psi]$ is ℓ -maximal in $\mathbb{Q}(\pi_q)$. In particular, if ℓ is upper-volcanic then $\ell \mid r$ (and $(0, \ell)$ can be added to \mathcal{R} in Algorithm 2).

7. Complements

Schoof's original algorithm may be generalized from prime ℓ to small prime powers in a very simple way. Going further, we may use isogeny cycles to compute eigenspaces of π_q and ψ on $\mathcal{E}[\ell^n]$ for Elkies ℓ : the methods developed for π_q in [5] and [8] generalize to ψ without any difficulty. Once we have recovered

$$\psi(P) = [k_n]P \quad \text{for } P = (X, Y) \in \mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_\mathcal{E}(X), F_{\ell^n}(X))),$$

we have $k_{n+1} = k_n + \tau \ell^n$ for $0 \leq \tau < \ell$, and we need to test

$$\psi(P) - [k_n]P = [\tau]([\ell^n]P) \quad \text{in } \mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_\mathcal{E}(X), F_{\ell^{n+1}}(X)))$$

(here F_{ℓ^n} and $F_{\ell^{n+1}}$ are factors of Ψ_{ℓ^n} and $\Psi_{\ell^{n+1}}$ that are minimal polynomials for ℓ^n and ℓ^{n-1} -torsion points).

We may extend Algorithms 1 and 2 to use Atkin primes. If ℓ is Atkin, then π_q and ψ have no rational eigenspaces in $\mathcal{E}[\ell]$; but we may still compute t_ℓ and r_ℓ by working on the full ℓ -torsion, as in Schoof's original algorithm. If P is a symbolic point of $\mathcal{E}[\ell]$ then $(\chi_{\pi_q} \pmod{\ell})(P) = 0$, so in

Algorithm 2: ADMISSIBLETRACE

Input : A d -admissible curve $\mathcal{E}/\mathbb{F}_{p^2}$, where p is large
Output: The trace of Frobenius of \mathcal{E}

```

1  $\mathcal{R} \leftarrow \{\}$ ; //  $\mathcal{R}$  will contain the pairs  $(r \pmod{\ell}, \ell)$ 
2  $M \leftarrow 1$ ; // After each iteration,  $r$  is known modulo  $M$ 
3  $\ell \leftarrow 1$ ;
4 while  $M \leq 4\sqrt{p/d}$  do
5   repeat  $\ell \leftarrow \text{NEXTPRIME}(\ell)$  until  $\ell \nmid d$ ;
6    $J \leftarrow \text{GCD}(x^{p^2} - x, \text{EVALUATEDMODULARPOLYNOMIAL}(\ell, \mathcal{E}))$ ;
7   if  $\deg J = 1$  or  $\ell + 1$  then //  $\ell$  is volcanic
8     if  $dp$  has a square root  $s$  modulo  $\ell$  then
9        $F \leftarrow \text{KERNELPOLYNOMIAL}(\ell, \mathcal{E}, \text{ONEROOT}(J))$ ;
10       $P \leftarrow (X, Y) \in \mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_{\mathcal{E}}(X), F(X)))$ ;
11       $Q_1 \leftarrow \pi_p(P)$ ;  $Q_2 \leftarrow {}^\sigma\phi(Q_1)$ ;  $Q_3 \leftarrow [s]P$ ;
12       $r_\ell \leftarrow \begin{cases} 2s/d \pmod{\ell} & \text{if } Q_3 = Q_2; \\ -2s/d \pmod{\ell} & \text{if } Q_3 = -Q_2; \\ 0 & \text{otherwise;} \end{cases}$ 
13    else  $r_\ell \leftarrow 0$ ;
14     $\mathcal{R} \leftarrow \mathcal{R} \cup \{(r_\ell, \ell)\}$ ;  $M \leftarrow \ell M$ ;
15  else if  $\deg J = 2$  then //  $\ell$  is Elkies
16     $F \leftarrow \text{KERNELPOLYNOMIAL}(\ell, \mathcal{E}, \text{ONEROOT}(J))$ ;
17     $P \leftarrow (X, Y) \in \mathcal{E}(\mathbb{F}_q[X, Y]/(Y^2 - f_{\mathcal{E}}(X), F(X)))$ ;
18     $Q_1 \leftarrow \pi_p(P)$ ;  $Q_2 \leftarrow {}^\sigma\phi(Q_1)$ ;
19     $r_\ell \leftarrow \lambda/d + p/\lambda \pmod{\ell}$  where  $\lambda = \text{DISCRETELOGARITHM}(P, Q_2)$ ;
20     $\mathcal{R} \leftarrow \mathcal{R} \cup \{(r_\ell, \ell)\}$ ;  $M \leftarrow \ell M$ ;
21 return  $\epsilon(dr^2 - 2p)$  where  $r = \text{CHINESEREMAINDERTHEOREM}(\mathcal{R})$ ;

```

Algorithm 1, t_ℓ is the discrete logarithm of $\pi_q(\pi_q(P)) + [q \bmod \ell]P$ to the base $\pi_q(P)$; similarly, in Algorithm 2, r_ℓ is the discrete logarithm of $\epsilon\pi_q(P) + [p \bmod \ell](P)$ to the base $\psi(P)$ (here we use $\psi^2 - dr\psi + [dp] = d(\epsilon\pi_q - r\psi + [p]) = 0$ and $\ell \nmid d$). The kernel polynomial defining $\mathcal{E}[\ell]$ is Ψ_ℓ , which we can compute using standard recurrences involving the coefficients of $f_{\mathcal{E}}$ (using the method of [2], for example) in $O(M(\ell^2) \log \ell)$ \mathbb{F}_q -operations. But Ψ_ℓ has degree $(\ell^2 - 1)/2$, so computing t_ℓ resp. r_ℓ costs $O((\log q)M(\ell^2))$ resp. $O((\log p)M(\ell^2))$ \mathbb{F}_q -operations; for that cost, we would gain much more information by using a larger Elkies prime instead. Alternatively, we can use Atkin's initial ideas using the splitting degree of $\Phi_\ell(X, j(\mathcal{E}))$ to determine a list of potential t_ℓ to be used in a tricky match and sort algorithm, or the more advanced algorithm of [12]. In our setting, we could use (5.3) to transform the list of t_ℓ 's to build a list of r_ℓ 's (on average, this does not increase the size of the lists too much).

Finally, we mention the use of the baby-step giant-step approach to speed up the final computations. If $P \in \mathcal{E}(\mathbb{F}_q)$, then $\chi_\psi(P) = 0$ becomes $[\epsilon d + dp]P = [rd]\psi(P)$, so $[p + \epsilon](Q) = [r]\psi(Q)$ with $Q = [d]P$ (if $Q = O_{\mathcal{E}}$, then another P should be used). Suppose we stop the loop of Algorithm 2 early; then r is known modulo M . Writing $r = r_0 + sM$ with $|s| \leq 2\sqrt{p/d}/M$, we can find s by solving $[p + \epsilon - r_0]Q = [s]([M]\psi(Q))$ for a sufficiently general choice of Q in $\mathcal{E}(\mathbb{F}_q)$; this is a classical discrete logarithm problem with \mathbb{F}_q -points, but in a smaller search space than the whole of $\mathcal{E}(\mathbb{F}_q)$. The optimal threshold for M is best determined through experiments.

8. Comparison of Algorithms 1 and 2

Let us compare the cost of computing $t_{\mathcal{E}}$ with Algorithms 1 and 2 when \mathcal{E} is d -admissible. For simplicity, we will suppose that Algorithm 1 also avoids the primes dividing d (these are very few and very small, so they do not contribute asymptotically or practically to the comparison).

The first clear difference between the algorithms is the number and size of primes ℓ used: Algorithm 2 essentially uses the smaller half of the set of primes used by Algorithm 1. The largest primes in each set still have roughly the same size, $O(\log p)$, so asymptotically this makes no difference—but using half the number of primes, and the smaller half at that, represents an important improvement in practice.

Now consider the cost of computing t_{ℓ} (as in Algorithm 1) or r_{ℓ} (as in Algorithm 2) for the same ℓ . The costs of determining the class of ℓ and the calls to `KERNELPOLYNOMIAL` are identical, and the calls to `DISCRETELOGARITHM` are equivalent. The only real difference is in how each algorithm computes the relations used to determine t_{ℓ} and r_{ℓ} .

- if ℓ is Elkies, then Algorithm 1 uses $2 \times \pi_p$ while Algorithm 2 uses $1 \times \pi_p + 1 \times \sigma\phi$.
- if ℓ is volcanic, then (in the worst cases) Algorithm 1 uses $2 \times \pi_p + 1 \times [s \bmod \ell]$, while Algorithm 2 uses $1 \times \pi_p + 1 \times \sigma\phi + 1 \times [s \bmod \ell]$.

In each case, the asymptotic costs are the same; but if $d \ll \log p$, then the costs are dominated by computations of π_p on $\langle P \rangle$ (for the same P). The crucial practical difference is that for each class of prime, Algorithm 2 exchanges half of the computations of π_p required by Algorithm 1 for one computation of $\sigma\phi$, which has a very small cost when $d \ll \log p$. Hence, for any given prime ℓ , Algorithm 2 should compute r_{ℓ} twice as quickly as Algorithm 1 computes t_{ℓ} .

By our complexity analysis, we see that the largest ℓ is $O(\log p)$ instead of $O(\log q)$, and we use the smaller half of them, we expect a real speedup of a factor of four. This is confirmed by our experimental results in §11 below.

9. \mathbb{Q} -curves and other sources of admissible curves

Admissible curves appear naturally as reductions of quadratic \mathbb{Q} -curves modulo inert primes (cf. [24, §3]). As such, we can construct parametrized families of admissible curves over any \mathbb{F}_{p^2} .

DEFINITION 2. A *quadratic \mathbb{Q} -curve of degree d* is an elliptic curve $\tilde{\mathcal{E}}$ without complex multiplication, defined over a quadratic field $\mathbb{Q}(\sqrt{\Delta})$, such that there exists an isogeny of degree d from $\tilde{\mathcal{E}}$ to its Galois conjugate $\tau\tilde{\mathcal{E}}$, where τ is the conjugation of $\mathbb{Q}(\sqrt{\Delta})$ over \mathbb{Q} .

PROPOSITION 6. Let $\tilde{\mathcal{E}}/\mathbb{Q}(\sqrt{\Delta})$ be a quadratic \mathbb{Q} -curve of degree d . If $p \nmid d$ is a prime of good reduction for $\tilde{\mathcal{E}}$ that is inert in $\mathbb{Q}(\sqrt{\Delta})$, then the reduction of $\tilde{\mathcal{E}}$ modulo p is d -admissible.

Proof. González shows that a d -isogeny $\tilde{\phi} : \tilde{\mathcal{E}} \rightarrow \tau\tilde{\mathcal{E}}$ must be defined over $\mathbb{Q}(\sqrt{\Delta}, \sqrt{\pm d})$ (see [9, §3]); so if we extend τ to the involution of $\mathbb{Q}(\sqrt{\Delta}, \sqrt{\pm d})$ that acts trivially on $\mathbb{Q}(\sqrt{\pm d})$ if and only if $\sqrt{\pm d}$ is in \mathbb{F}_p , then $\tilde{\phi}$ reduces modulo p to a d -isogeny $\phi : \mathcal{E} \rightarrow \sigma\mathcal{E}$ over \mathbb{F}_{p^2} , and $\tau\tilde{\phi}$ reduces to $\sigma\phi$. Observe that $\tau\tilde{\phi}\tilde{\phi}$ is an endomorphism of $\tilde{\mathcal{E}}$ of degree d^2 . Since $\tilde{\mathcal{E}}$ does not have complex multiplication, its only endomorphisms of degree d^2 are $[\pm d]$; hence $\tau\tilde{\phi} = \epsilon\tilde{\phi}^\dagger$ with $\epsilon = \pm 1$. Reducing modulo p we have $\sigma\phi = \epsilon\phi^\dagger$, so \mathcal{E} is d -admissible. \square

We emphasize that if a d -admissible curve \mathcal{E} is the reduction of a quadratic \mathbb{Q} -curve $\tilde{\mathcal{E}}$, then the associated endomorphism on \mathcal{E} is *not* the reduction of any endomorphism on $\tilde{\mathcal{E}}$. Indeed, $\tilde{\mathcal{E}}$ has no non-integer endomorphisms by definition.

EXAMPLE 1. Fix any prime $p > 3$; the following construction (carried much further in [23] and [24]) yields a 1-parameter family of 2-admissible curves over \mathbb{F}_{p^2} . Let Δ be a squarefree integer that is not a square modulo p (so p is inert in $\mathbb{Q}(\sqrt{\Delta})$), let τ be the involution of $\mathbb{Q}(\sqrt{\Delta}, \sqrt{-2})$ that restricts to σ modulo p , and let s be a free parameter taking values in \mathbb{Q} . The family of curves over $\mathbb{Q}(\sqrt{\Delta})$ defined by $\tilde{\mathcal{E}} : y^2 = x^3 - 6(5 - 3s\sqrt{\Delta})x + 8(7 - 9s\sqrt{\Delta})$ is equipped with a 2-isogeny $\tilde{\phi} : \tilde{\mathcal{E}} \rightarrow {}^\tau\tilde{\mathcal{E}}$ over $\mathbb{Q}(\sqrt{\Delta}, \sqrt{-2})$ with kernel polynomial $D(x) = x - 4$ (see [11, Prop 3.3]). Computing $\tilde{\phi}^\dagger$ and ${}^\tau\tilde{\phi}$, we find that ${}^\tau\tilde{\phi} = \epsilon\tilde{\phi}^\dagger$, where $\epsilon = 1$ if $p \equiv 5, 7 \pmod{8}$ and $\epsilon = -1$ if $p \equiv 1, 3 \pmod{8}$. Reducing everything modulo p , as in the proof of Prop. 6, we obtain a family of curves

$$\mathcal{E} : y^2 = x^3 - 6(5 - 3s\sqrt{\Delta})x + 8(7 - 9s\sqrt{\Delta}) \quad \text{over } \mathbb{F}_{p^2} = \mathbb{F}_p(\sqrt{\Delta})$$

with the parameter s taking values in \mathbb{F}_p , equipped with a 2-isogeny $\phi : \mathcal{E} \rightarrow {}^\sigma\mathcal{E}$ over \mathbb{F}_{p^2} . Composing π_p with ${}^\sigma\phi$ yields the associated endomorphism ψ of \mathcal{E} , defined by

$$\psi : (x, y) \mapsto \left(\frac{x^p(x^p - 4) + 18(1 - s\sqrt{\Delta})}{-2(x^p - 4)}, \frac{y^p}{\sqrt{-2}^p} \left(\frac{(x^p - 4)^2 - 18(1 - s\sqrt{\Delta})}{-2(x^p - 4)^2} \right) \right).$$

Since the definition of admissible curves involves only isogenies over \mathbb{F}_{p^2} , we would expect a characterization of admissible curves over a given \mathbb{F}_{p^2} in terms of modular polynomials.

PROPOSITION 7. *If \mathcal{E} is an ordinary elliptic curve over $\mathbb{F}_q = \mathbb{F}_{p^2}$ such that $j(\mathcal{E})$ is a simple root of $\Phi_d(x, x^p)$ in $\mathbb{F}_q \setminus \{0, 1728\}$ (so in particular, $\text{Aut}_{\mathbb{F}_q}(\mathcal{E}) = \{[\pm 1]\}$), then \mathcal{E} is d -admissible.*

Proof. If $j(\mathcal{E})$ is a simple root of $\Phi_d(x, x^p)$ in \mathbb{F}_q , then up to $\overline{\mathbb{F}_q}$ -isomorphism there is a unique d -isogeny $\phi : \mathcal{E} \rightarrow {}^\sigma\mathcal{E}$. If ϕ were not defined over \mathbb{F}_q , then the endomorphism ${}^\sigma\pi_p\phi$ would not be defined over \mathbb{F}_q , hence not commute with π_q , contradicting non-supersingularity. For d -admissibility, it remains to show that ${}^\sigma\phi = \epsilon\phi^\dagger$ with $\epsilon = \pm 1$. But if this were not the case, then $({}^\sigma\phi)^\dagger$ would be a second d -isogeny $\mathcal{E} \rightarrow {}^\sigma\mathcal{E}$, not isomorphic to ϕ (since $\text{Aut}_{\mathbb{F}_q}(\mathcal{E}) = \{[\pm 1]\}$); that is, $j(\mathcal{E})$ would be (at least) a double root of $\Phi_d(x, x^p)$. \square

EXAMPLE 2. Multiple roots of $\Phi_d(x, x^p)$ may not yield d -admissible curves. Consider the ordinary curve $\mathcal{E} : y^2 = x^3 + (38 + 53i)x + 27 - 3i$ over $\mathbb{F}_q = \mathbb{F}_{103}(i)$ where $i^2 = -1$: then $j(\mathcal{E}) = 35 + 5i$ is a double root of $\Phi_3(x, x^{103})$. Indeed, we have a pair of non-isomorphic 3-isogenies $\phi_1 : \mathcal{E} \rightarrow {}^\sigma\mathcal{E}$ and $\phi_2 : \mathcal{E} \rightarrow {}^\sigma\mathcal{E}$, with kernel polynomials $x + 1 + 39i$ and $x - 4 + 32i$, respectively; but ${}^\sigma\phi_1 = \pm\phi_2^\dagger$ and ${}^\sigma\phi_2 = \pm\phi_1^\dagger$, so \mathcal{E} is not 3-admissible.

10. Generating cryptographically strong curves

One of the important motivations for developing our algorithm was the generation of cryptographically strong curves. Indeed, the curves proposed for cryptographic applications in [23] and [24], and which were subsequently used in fast, compact Diffie–Hellman key exchange software [3], are admissible. These curves were designed to offer accelerated scalar multiplication (using the associated endomorphism) over fast finite fields, without obstructing twist-security; but when generating twist-secure curves at and above the 128-bit security level, we can expect to try hundreds of thousands of curves before finding a suitable one. In this context of counting many curves, practical speedups become very important.

For cryptographic applications based on the hardness of the discrete logarithm problem, the minimum requirement for a “secure” curve $\mathcal{E}/\mathbb{F}_{p^2}$ is that $\#\mathcal{E}(\mathbb{F}_{p^2}) = c \cdot n$, where n is prime and c is tiny (traditionally, we want $c = 1$; more modern software using Montgomery and Edwards

models requires $c = 2$ or 4). For some applications we further require “twist-security”: that is, the quadratic twist \mathcal{E}' should satisfy $\#\mathcal{E}'(\mathbb{F}_{p^2}) = c' \cdot n'$, where n' is prime and c' is tiny.

To find a secure or twist-secure curve over \mathbb{F}_{p^2} we typically fix a prime p of bitlength around the required security parameter, then test a series of curves over \mathbb{F}_{p^2} , computing their orders until we find a curve with the right structure. Equation (5.3) implies

$$\#\mathcal{E}(\mathbb{F}_{p^2}) = (p + \epsilon)^2 - \epsilon dr^2 \quad \text{and} \quad \#\mathcal{E}'(\mathbb{F}_{p^2}) = (p - \epsilon)^2 + \epsilon dr^2 .$$

This places some immediate constraints on the combinations of d , p , and ϵ that can yield suitable curves. For example, $\#\mathcal{E}(\mathbb{F}_{p^2}) \equiv (p + \epsilon)^2 \pmod{d}$, so $d \mid \#\mathcal{E}(\mathbb{F}_{p^2})$ if and only if $p \equiv -\epsilon \pmod{d}$; such p should be avoided unless we can accept $d \mid c$. Similarly, if twist-security prohibits $d \mid c'$ then we should must avoid $p \equiv \epsilon \pmod{d}$. Clearly if \mathcal{E} is 2-admissible, then it must have a rational point of order 2, so we cannot do better than having $c = c' = 2$. Similarly, 3-admissible curves must have either $3 \mid c$ or $3 \mid c'$.

Extensive computations done for $d = 2$ and 3 over a range of primes revealed densities of twist-secure d -admissible curves (modulo the constraints above) similar to the densities of twist-secure general elliptic curves over the same fields.

With Algorithm 1, we can speed up the search for secure curves by checking whether $t_\ell \equiv p^2 + 1 \pmod{\ell}$ for each ℓ ; if so, then $\ell \mid \#\mathcal{E}(\mathbb{F}_{p^2})$, so we can abort the computation and move on to the next candidate curve [14]. Similarly, if $t_\ell \equiv -(p^2 + 1) \pmod{\ell}$ then $\ell \mid \#\mathcal{E}'(\mathbb{F}_{p^2})$.

With Algorithm 2, if ℓ divides $\#\mathcal{E}(\mathbb{F}_{p^2})$ then $(p + \epsilon)^2 \equiv \epsilon dr^2 \pmod{\ell}$, so ℓ cannot divide $\#\mathcal{E}(\mathbb{F}_{p^2})$ unless ϵd is a square mod ℓ ; and if ϵd is a square mod ℓ , then we should abort if $r_\ell \equiv \pm(p + \epsilon)/\sqrt{\epsilon d} \pmod{\ell}$. In fact, if $r_\ell \equiv 0$ and $p + \epsilon \equiv 0 \pmod{\ell}$, then the nondegeneracy of the ℓ -Weil pairing implies that $\mathcal{E}[\ell](\mathbb{F}_{p^2}) \cong (\mathbb{Z}/\ell\mathbb{Z})^2$. Replacing ϵ with $-\epsilon$ yields analogous results for the twist \mathcal{E}' .

We note also that there may be an advantage in generating curves using the parameter r and not $t_\mathcal{E}$. We could force some value of ℓ to divide r by rejecting curves \mathcal{E} for which $\Phi_\ell(X, j(\mathcal{E}))$ does not have 1 or $\ell + 1$ roots. This has no impact on $t_\mathcal{E}$, and we already know $r \pmod{\ell}$. We just need to hope that such curves are as secure as general d -admissible curves.

11. Implementation and experiments

We implemented the new algorithm on top of our implementation of SEA, realized in C++ using NTL 9.6.4 (with gcc 4.9.2). The timings below (in seconds) are for an Intel Xeon platform (E5520 CPU at 2.27GHz). We define two primes (of 128 and 255 bits), derived from the decimal expansion of π :

$$p_{128} := 314159265358979323846264338327950288459 ,$$

$$p_{255} := 31415926535897932384626433832795028841971693993751058209749445923078164062963 .$$

First, we compare the straightforward computation of $X^q \pmod{\Phi_\ell}$ to a modular composition over \mathbb{F}_{p^2} with $p = p_{128}$ and p_{255} , for two choices of ℓ :

p_{128}				p_{255}			
ℓ	$X^p \pmod{\Phi_\ell}$	$X^p \circ X^p$	X^q	ℓ	$X^p \pmod{\Phi_\ell}$	$X^p \circ X^p$	X^q
101	0.23	0.04	0.47	101	0.69	0.07	1.40
173	0.43	0.11	0.88	173	1.38	0.18	2.80

Then we ran our program on curves from the family of Example 1, for each $1 \leq s \leq 100$. This gave the following average values:

p_{128}			p_{255}		
	Algorithm 1	Algorithm 2		Algorithm 1	Algorithm 2
max. ℓ	164	62	max. ℓ	352	160.76
X^q time	9.11	2.62	X^q time	89.73	22.55
Total time	20.11	4.1	Total time	171.95	39.16

Finally, we searched for twist-secure curves with small values of the parameter s . For instance, with $p = p_{128}$ and $s = 113$, we get a curve of cardinality $2p'$, whose twist has cardinality $6p''$; with $p = p_{255}$, taking $s = 269$ yields a pair of curves each with cardinality two times a prime.

Acknowledgments. We thank A. Sutherland for pointing out an error in the complexity analysis of the SEA algorithm.

Appendix A. Detailed complexity of basic computations

Let $F(X)$ be a degree e polynomial with coefficients in $\mathbb{F}_q[X]$. We define G and H to be the polynomials of degree $< e$ such that $H \equiv X^p \pmod{F}$ and

$$Y^p \equiv YG(X) \text{ with } G(X) \equiv f_{\mathcal{E}}^{(p-1)/2}(X) \pmod{F(X)}. \quad (\text{A.1})$$

A.1. Computing $X^q \pmod{F}$

The first step in factoring F is to compute $X^q \pmod{F}$. When $q = p^n$ for some prime p , we may start by computing H and then proceed with modular composition.

If $R(X) = \sum_{i=0}^{e-1} r_i X^i$ with $r_i \in \mathbb{F}_q$, then $\sigma R(X) = \sum_{i=0}^{e-1} r_i^p X^i$ satisfies $R^p \pmod{F} = \sigma R \circ X^p \pmod{F}$. We assume that the cost of computing all the r_i^p is negligible (as it is with a suitable choice of basis for $\mathbb{F}_q/\mathbb{F}_p$: if $\mathbb{F}_{p^2} = \mathbb{F}_p(\sqrt{\Delta})$, then $(a + b\sqrt{\Delta})^p = a - b\sqrt{\Delta}$ for all a and b in \mathbb{F}_p). For our purposes, the computation of X^{p^2} computes $H(X)$ and $X^{p^2} = \sigma H \circ H \pmod{F}$, which costs $O((\log p)\mathcal{M}(e) + \mathcal{C}(e))$ instead of $O((\log q)\mathcal{M}(e))$, which is larger provided that $2e \leq (\log p)^2$. When $q = p^n$ with $n > 2$, similar savings can be obtained.

A.2. Proof of Lemma 1

Let $F = D$, or any factor of D (as in the extensions of the algorithm mentioned in §7).

For (i), the obvious way is to compute H , then G , in $O((\log p)\mathcal{M}(e))$ \mathbb{F}_q -operations. Alternatively, we can adapt the methods of [8]: first compute G in $O((\log p)\mathcal{M}(e))$ operations. Consider the polynomial $P(W) = W^3 + A^p W + B^p - (X^3 + AX + B)G(X)^2$. Then $X^p \pmod{F}$ is a root of both $\sigma F(W)$ and $P(W)$ in $\mathbb{F}_q[X]/(F(X))$, so $W - H(X) \mid g = \gcd(P(W), \sigma F(W))$. Very generally, $g = W - H(X)$. The main cost is that of reducing $\sigma F(W)$ modulo $P(W)$, which is $O(e\mathcal{M}(e))$. This can be reduced to $\mathcal{C}_3(e)$ or even $O((\log \ell)\mathcal{M}(e))$ if F divides Ψ_ℓ .

For (ii): we can compute $\pi_p(Q) = (Q_x^p, Y^p Q_y^p) = (\sigma Q_x \circ H \pmod{F}, YG(\sigma Q_y \circ H) \pmod{F})$ in $\mathcal{C}_2(e)$ \mathbb{F}_q -operations. This also applies for computing $\pi_q(P) = (X^{p^2}, Y^{p^2}) = \pi_p(H, YG)$.

For (iv): suppose $\phi = (N/D, M/D^2)$ with $\deg N = \deg M = 2$ and $\deg D = 1$. We compute $N \circ Q_x \pmod{F}$, $M \circ Q_x \pmod{F}$, and $D \circ Q_x \pmod{F}$ followed by some multiplications, keeping numerators and denominators. We only need a few modular multiplications, for a cost of $O(\mathcal{M}(e))$.

For (v), we have $\phi = (N/D^2, M/D^3)$ with $\deg(N) = \deg(M) = d$, and $\deg(D) = (d-1)/2$. First we reduce N , M , and D modulo F (if necessary), at a cost of $O(\mathcal{M}(d))$. We then compute $N \circ Q_x \pmod{F}$, $M \circ Q_x \pmod{F}$, and $D \circ Q_x \pmod{F}$ followed by some multiplications, keeping numerators and denominators. The dominating cost is bounded by $O(\mathcal{M}(d) + \mathcal{C}_3(e))$.

References

1. Alin Bostan, François Morain, Bruno Salvy, and Éric Schost. Fast algorithms for computing isogenies between elliptic curves. *Math. Comput.*, 77(263):1755–1778, 2008.
2. Qi Cheng. Straight-line programs and torsion points on elliptic curves. *Comput. Complexity*, 12:150–161, 2003.
3. Craig Costello, Hüseyin Hisil, and Benjamin Smith. Faster compact Diffie-Hellman: Endomorphisms on the x -line. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Comput. Sci.*, pages 183–200. Springer, 2014.
4. Craig Costello and Patrick Longa. Four \mathbb{Q} : Four-dimensional decompositions on a \mathbb{Q} -curve over the Mersenne prime. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, volume 9452 of *Lecture Notes in Comput. Sci.*, pages 214–235. Springer, 2015.
5. Jean-Marc Couveignes and François Morain. Schoof’s algorithm and isogeny cycles. In L. Adleman and M.-D. Huang, editors, *Algorithmic Number Theory*, volume 877 of *Lecture Notes in Comput. Sci.*, pages 43–58. Springer-Verlag, 1994. 1st Algorithmic Number Theory Symposium - Cornell University, May 6-9, 1994.
6. Mireille Fouquet and François Morain. Isogeny volcanoes and the SEA algorithm. In C. Fieker and D. R. Kohel, editors, *Algorithmic Number Theory*, volume 2369 of *Lecture Notes in Comput. Sci.*, pages 276–291. Springer-Verlag, 2002. 5th International Symposium, ANTS-V, Sydney, Australia, July 2002, Proceedings.
7. Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
8. Pierrick Gaudry and François Morain. Fast algorithms for computing the eigenvalue in the Schoof-Elkies-Atkin algorithm. In *ISSAC ’06: Proceedings of the 2006 international symposium on Symbolic and algebraic computation*, pages 109–115, New York, NY, USA, 2006. ACM Press.
9. Josep González. Isogenies of polyquadratic \mathbb{Q} -curves to their Galois conjugates. *Archiv der Mathematik*, 77:383–390, 2001.
10. Aurore Guillevic and Sorina Ionica. Four-dimensional GLV via the Weil restriction. In Sako and Sarkar [16], pages 79–96.
11. Yuji Hasegawa. \mathbb{Q} -curves over quadratic fields. *Manuscripta Math.*, 94(1):347–364, 1997.
12. Antoine Joux and Reynald Lercier. ”Chinese & Match”, an alternative to Atkin’s ”Match and Sort” method used in the SEA algorithm. *Math. Comp.*, 70(234):827–836, April 2001.
13. Erich Kaltofen and Victor Shoup. Subquadratic-time factoring of polynomials over finite fields. *Math. Comp.*, 67(223):1179–1197, 1998.
14. Reynald Lercier. Finding good random elliptic curves for cryptosystems defined over \mathbb{F}_{2^n} . In W. Fumy, editor, *Advances in Cryptology - EUROCRYPT ’97*, volume 1233 of *Lecture Notes in Comput. Sci.*, pages 379–392. Springer-Verlag, 1997.
15. Preda Mihăilescu, François Morain, and Éric Schost. Computing the eigenvalue in the Schoof-Elkies-Atkin algorithm using Abelian lifts. In *ISSAC ’07: Proceedings of the 2007 international symposium on Symbolic and algebraic computation*, pages 285–292, New York, NY, USA, 2007. ACM Press.
16. Kazue Sako and Palash Sarkar, editors. *Advances in Cryptology - ASIACRYPT 2013*, volume 8269 of *Lecture Notes in Comput. Sci.* Springer, 2013.
17. Takakazu Satoh. On p -adic point counting algorithms for elliptic curves over finite fields. In Claus Fieker and David R. Kohel, editors, *Algorithmic Number Theory, 5th International Symposium, ANTS-V, Sydney, Australia, July 7-12, 2002, Proceedings*, volume 2369 of *Lecture Notes in Comput. Sci.*, pages 43–66. Springer, 2002.
18. Rene Schoof. Elliptic curves over finite fields and the computation of square roots mod p . *Math. Comp.*, 44:483–494, 1985.
19. Rene Schoof. Nonsingular plane cubic curves over finite fields. *J. Combin. Theory Ser. A*, 46(2):183–211, 1987.
20. Rene Schoof. Counting points on elliptic curves over finite fields. *J. Théor. Nombres Bordeaux*, 7:219–254, 1995.
21. Igor E. Shparlinski and Andrew V. Sutherland. On the distribution of Atkin and Elkies primes. *Found. Comput. Math.*, 14(2):285–297, 2014.
22. Igor E. Shparlinski and Andrew V. Sutherland. On the distribution of Atkin and Elkies primes for reductions of elliptic curves on average. *LMS J. Comput. Math.*, 18:308–322, 1 2015.
23. Benjamin Smith. Families of fast elliptic curves from \mathbb{Q} -curves. In Sako and Sarkar [16], pages 61–78.
24. Benjamin Smith. The \mathbb{Q} -curve construction for endomorphism-accelerated elliptic curves. *J. Cryptology*, 2015.
25. Andrew V. Sutherland. Identifying supersingular elliptic curves. *LMS J. Comput. Math.*, 15:317–325, 2012.
26. Andrew V. Sutherland. On the evaluation of modular polynomials. In *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium*, volume 1 of *Open Book Ser.*, pages 531–555. Math. Sci. Publ., Berkeley, CA, 2013.

27. Hui June Zhu. Group structures of elementary supersingular abelian varieties over finite fields. *J. Number Theory*, 81:292–309, 2000.

F. Morain
École Polytechnique/LIX
and Centre national de la recherche
scientifique (CNRS)
and Institut national de recherche en
informatique et en automatique (INRIA)
France
morain@lix.polytechnique.fr

C. Scribot
Ministère de l'Éducation Nationale
France

B. Smith
Institut national de recherche en
informatique et en automatique (INRIA)
and École Polytechnique/LIX
and Centre national de la recherche
scientifique (CNRS)
France
smith@lix.polytechnique.fr