

8 Schoof's algorithm

In the early 1980's, René Schoof [3, 4] introduced the first polynomial-time algorithm to compute $\#E(\mathbb{F}_q)$. Extensions of Schoof's algorithm remain the point-counting method of choice when the characteristic of \mathbb{F}_q is large (e.g., when q is a cryptographic size prime).¹

Schoof's basic strategy is simple: compute the trace of Frobenius t modulo many small primes ℓ and use the Chinese remainder theorem to uniquely determine t , which then determines $\#E(\mathbb{F}_q) = q + 1 - t$. Here is a high-level version of the algorithm.

Algorithm 8.1. Given an elliptic curve E over a finite field \mathbb{F}_q compute $\#E(\mathbb{F}_q)$ as follows:

1. Initialize $M \leftarrow 1$ and $t \leftarrow 0$.
2. While $M \leq 4\sqrt{q}$, for increasing primes $\ell = 2, 3, 5, \dots$ that do not divide q :
 - a. Compute $t_\ell = \text{tr } \pi \pmod{\ell}$.
 - b. Set $t \leftarrow (M(M^{-1} \pmod{\ell})t_\ell + \ell(\ell^{-1} \pmod{M})t) \pmod{\ell M}$ and then $M \leftarrow \ell M$.
3. If $t > M/2$ then set $t \leftarrow t - M$.
4. Output $q + 1 - t$.

Step 2b uses an iterative version of the Chinese remainder theorem to ensure that

$$t \equiv \text{tr } \pi_E \pmod{M}$$

holds throughout.² This invariant holds trivially after step 1, modulo $M = 1$, and is maintained in step 2b: note that the integer $M(M^{-1} \pmod{\ell})$ is congruent to $1 \pmod{\ell}$ and $0 \pmod{M}$, while the integer $\ell(\ell^{-1} \pmod{M})$ is congruent to $0 \pmod{\ell}$ and $1 \pmod{M}$.

Once M exceeds $4\sqrt{q}$, the value of $t \in \mathbb{Z}/M\mathbb{Z}$ uniquely determines $\text{tr } \pi_E \in \mathbb{Z}$: by Hasse's theorem, $|\text{tr } \pi_E| \leq 2\sqrt{q} < M/2$, and this allows us to determine the sign of $\text{tr } \pi_E$ in step 3. The key to the algorithm is the implementation of step 2a, which is described in the next section, but let us first consider the primes ℓ that the algorithm uses. Let ℓ_{max} be the largest prime ℓ for which the algorithm computes t_ℓ . The Prime Number Theorem implies³

$$\sum_{\text{primes } \ell \leq x} \log \ell \sim x,$$

so $\ell_{max} \approx \log 4\sqrt{q} \approx \frac{1}{2}n = O(n)$, and we need $O(\frac{n}{\log n})$ primes ℓ (as usual, $n = \log q$). The cost of updating t and M is bounded by $O(M(n) \log n)$, thus if we can compute t_ℓ in time bounded by a polynomial in n and ℓ , then the whole algorithm will run in polynomial time.

8.1 Computing the trace of Frobenius modulo 2.

We first consider the case $\ell = 2$. Assuming q is odd (which we do), $t = q + 1 - \#E(\mathbb{F}_q)$ is divisible by 2 if and only if $\#E(\mathbb{F}_q)$ is divisible by 2, equivalently, if and only if $E(\mathbb{F}_q)$ contains a point of order 2. If E has Weierstrass equation $y^2 = f(x)$, then the points of order 2 in $E(\mathbb{F}_q)$ are precisely those of the form $(x_0, 0)$, where $x_0 \in \mathbb{F}_q$ is a root $f(x)$. Recall

¹There are deterministic p -adic algorithms for computing $\#E(\mathbb{F}_q)$ that are faster than Schoof's algorithm when the characteristic p of \mathbb{F}_q is very small; see [2]. But their running times are exponential in $\log p$.

²There are faster ways to apply the Chinese remainder theorem; see [1, §10.3]. They are not relevant here because the complexity is overwhelmingly dominated by step 2a.

³In fact we only need Chebyshev's Theorem to get this.

from Lecture 4 that the distinct roots of f in \mathbb{F}_q are precisely the roots of $\gcd(x^q - x, f(x))$. We can thus compute $t_2 := \text{tr } \pi_E \bmod 2$ as

$$t_2 = \begin{cases} 0 & \text{if } \deg(\gcd(f(x), x^q - x)) > 0; \\ 1 & \text{otherwise.} \end{cases}$$

Note that is a deterministic computation (we need randomness to efficiently *find* the roots of $g(x)$, but not to *count* them), and it takes $O(nM(n))$ time.

Having addressed the case $\ell = 2$ we henceforth assume that ℓ is odd.

8.2 The characteristic equation of Frobenius modulo ℓ

Recall that for E/\mathbb{F}_q , the Frobenius endomorphism $\pi_E \in \text{End}(E)$ is defined by the rational map $(x : y : z) \mapsto (x^q : y^q : z^q)$. By Theorem 6.18, it satisfies the characteristic equation

$$\pi_E^2 - t\pi_E + q = 0,$$

with $t = \text{tr } \pi$ and $q = \deg \pi$. Restricting to the ℓ -torsion subgroup $E[\ell]$ yields

$$\pi_\ell^2 - t_\ell \pi_\ell + q_\ell = 0, \tag{1}$$

which we view as an identity in $\text{End}(E[\ell])$. Here $t_\ell \equiv t \bmod \ell$ and $q_\ell \equiv q \bmod \ell$ can be viewed either as restrictions of the scalar multiplication maps $[t]$ and $[q]$, or simply as scalars in $\mathbb{Z}/\ell\mathbb{Z}$ multiplied by $[1]_\ell$, the restriction of $[1] \in \text{End}(E)$ to $E[\ell]$ (equivalently the multiplicative identity in the ring $\text{End}(E[\ell])$). We shall take the latter view, regarding

$$q_\ell = q_\ell \cdot [1]_\ell = [1]_\ell + \cdots + [1]_\ell$$

as the sum of q_ℓ copies of $[1]_\ell$, and similarly for t_ℓ . We can efficiently compute q_ℓ using our usual double-and-add method to perform scalar multiplication by q_ℓ , provided that we know how to explicitly represent and perform ring operations on elements of $\text{End}(E[\ell])$; this is the topic of the next section.

Our strategy for determining t_ℓ is simple: for $c = 0, 1, \dots, \ell - 1$ compute $\pi_\ell^2 - c\pi_\ell + q_\ell$ and check whether it is equal to 0. The following lemma shows that whenever this occurs (which it must, since (1) guarantees this for $c = t_\ell$) we must have $c = t_\ell \in \mathbb{Z}/\ell\mathbb{Z}$. In fact we will prove something stronger.

Lemma 8.2. *Let E/\mathbb{F}_q be an elliptic curve with Frobenius endomorphism π , let ℓ be a prime not dividing q , and let $P \in E[\ell]$ be nonzero. Suppose that for some integer c the equation*

$$\pi_\ell^2(P) - c\pi_\ell(P) + q_\ell(P) = 0$$

holds. Then $c \equiv t_\ell = \text{tr } \pi \bmod \ell$.

Proof. From equation (1) we have

$$\pi_\ell^2(P) - t_\ell \pi_\ell(P) + q_\ell P = 0,$$

and we are assuming that

$$\pi_\ell^2(P) - c\pi_\ell(P) + q_\ell P = 0.$$

Subtracting these equations yields $(c - t_\ell)\pi_\ell(P) = 0$. Since $\pi_\ell P$ is a nonzero element of $E[\ell]$ and ℓ is prime, the point $\pi_\ell(P)$ has order ℓ , which must divide $c - t_\ell$. So $c \equiv t_\ell \bmod \ell$. \square

8.3 Arithmetic in $\text{End}(E[\ell])$

Let $h = \psi_\ell(x, y)$ be the ℓ th division polynomial of E . We have assumed that ℓ is odd, so by Lemma 5.20, we in fact have $h \in \mathbb{F}_q[x]$ (no dependence on y). As we proved in Lecture 5, a nonzero point $P = (x_0, y_0) \in E(\overline{\mathbb{F}}_q)$ lies in $E[\ell]$ if and only if $h(x_0) = 0$; this follows from Corollary 4.28 and Theorem 5.21. To represent elements of $\text{End}(E[\ell])$ as rational maps, we can thus treat the polynomials appearing in these maps as elements of the ring

$$\mathbb{F}_q[x, y]/(h(x), y^2 - f(x)),$$

where $y^2 = f(x) = x^3 + Ax + B$ is the Weierstrass equation for E .

In the case of the Frobenius endomorphism, we have

$$\begin{aligned} \pi_\ell &= (x^q \bmod h(x), y^q \bmod (h(x), y^2 - f(x))) \\ &= (x^q \bmod h(x), (f(x)^{(q-1)/2} \bmod h(x)) y), \end{aligned} \quad (2)$$

and we also note that

$$[1]_\ell = (x \bmod h(x), (1 \bmod h(x)) y).$$

We can thus represent all of the nonzero endomorphisms that appear in equation (1) in the form $(a(x), b(x)y)$, where a and b are elements of the polynomial ring $R = \mathbb{F}_q[x]/(h(x))$ that we may uniquely represent as polynomials in $\mathbb{F}_q[x]$ of degree less than $\deg h = (\ell^2 - 1)/2$ by taking their remainders modulo h .

8.3.1 Multiplication in $\text{End}(E[\ell])$.

If $\alpha_1 = (a_1(x), b_1(x)y)$ and $\alpha_2 = (a_2(x), b_2(x)y)$ are two elements of $\text{End}(E[\ell])$, the product $\alpha_1\alpha_2$ in $\text{End}(E[\ell])$ is defined by the composition

$$\alpha_1 \circ \alpha_2 = (a_1(a_2(x)), b_1(a_2(x))b_2(x)y),$$

where we may reduce $a_3(x) = a_1(a_2(x))$ and $b_3(x) = b_1(a_2(x))b_2(x)$ modulo $h(x)$.

8.3.2 Addition in $\text{End}(E[\ell])$.

Addition of endomorphisms is defined pointwise in terms of addition on the elliptic curve. Given $\alpha_1 = (a_1(x), b_1(x)y)$ and $\alpha_2 = (a_2(x), b_2(x)y)$, to compute $\alpha_3 = \alpha_1 + \alpha_2$, we simply apply the formulas for point addition to the coordinate functions of α_1 and α_2 . Recall that the general formula for addition of non-opposite affine points $(x_3, y_3) = (x_1, y_1) + (x_2, y_2)$ on the elliptic curve $E: y^2 = x^3 + Ax + B$ is given by the formulas

$$x_3 = m^2 - x_1 - x_2, \quad y_3 = m(x_1 - x_3) - y_1,$$

where

$$m = \frac{y_1 - y_2}{x_1 - x_2} \quad (\text{if } x_1 \neq x_2), \quad m = \frac{3x_1^2 + A}{2y_1} \quad (\text{if } x_1 = x_2).$$

Using the coordinate functions $x_1 = a_1(x)$, $x_2 = a_2(x)$, $y_1 = b_1(x)y$, $y_2 = b_2(x)y$, in the case $x_1 \neq x_2$ we have

$$m(x, y) = \frac{b_1(x) - b_2(x)}{a_1(x) - a_2(x)} y = r(x)y,$$

where $r = (b_1 - b_2)/(a_1 - a_2)$, and when $x_1 = x_2$ we have

$$m(x, y) = \frac{3a_1(x)^2 + A}{2b_1(x)y} = \frac{3a_1(x)^2 + A}{2b_1(x)f(x)}y = r(x)y,$$

where now $r = (3a_1^2 + A)/(2b_1f)$. Noting that $m(x, y)^2 = (r(x)y)^2 = r(x)^2f(x)$, the sum $\alpha_1 + \alpha_2 = \alpha_3 = (a_3(x), b_3(x)y)$ is defined by

$$\begin{aligned} a_3 &= r^2f - a_1 - a_2, \\ b_3 &= r(a_1 - a_3) - b_1. \end{aligned}$$

In both cases, provided that the polynomial v in the denominator of the rational function $r = u/v$ is invertible in the ring $\mathbb{F}_q[x]/(h(x))$, we can express r as a polynomial $uv^{-1} \bmod h$ and write $\alpha_3 = (a_3(x), b_3(x)y)$ in our desired form, with $a_3, b_3 \in \mathbb{F}_q[x]/(h(x))$ uniquely represented by polynomials in $\mathbb{F}_q[x]$ of degree less than the degree of h .

But this may not always be possible, because the ℓ -division polynomial $h(x)$ need not be irreducible. Indeed, if ℓ divides $\#E(\mathbb{F}_q)$ it certainly will not be irreducible, since $h(x)$ will then have rational roots corresponding to the x -coordinates of rational points of order ℓ , and even when $\ell \nmid \#E(\mathbb{F}_q)$, if E admits a rational isogeny α of degree ℓ then $h(x)$ will be divisible by the polynomial of degree $(\ell - 1)/2$ whose roots are the x -coordinates of the nonzero points in the kernel of α . Thus the ring $\mathbb{F}_q[x]/(h(x))$ is not necessarily a field; it may contain zero divisors, and these elements are not invertible.

At first glance this might appear to be a problem, but in fact it can only help us. If we encounter a rational function $r = u/v$ whose denominator v is not invertible in $\mathbb{F}_q[x]/(h(x))$ then we can obtain a non-trivial factor of h by computing $\gcd(v, h)$: if $v = a_1 - a_2$ then v is nonzero and has degree less than h , since in this case $a_1 \neq a_2$ and $\deg(a_1 - a_2) < \deg(h)$, and if $v = 2b_1f$ then $\gcd(v, h)$ must divide b_1 , because h and f cannot share a common factor (the roots of $f(x)$ in $\overline{\mathbb{F}}_q$ are x -coordinates of 2-torsion points, the roots of $h(x)$ in $\overline{\mathbb{F}}_q$ are x -coordinates of ℓ -torsion points, and $\ell \neq 2$), and $b_1 \neq 0$ has degree less than h .

Our strategy in this situation is to simply replace h by $g = \gcd(v, h)$ and compute t_ℓ by working in the smaller quotient ring $\mathbb{F}_q[x]/(g(x))$, which will be faster because $\deg g < \deg h$; in fact in this situation we will always have $\deg g \leq (\ell - 1)/2$, which is much smaller than $\deg h = (\ell^2 - 1)/2$. Lemma 8.2 implies that we can restrict our attention to the action of π_ℓ on points $P \in E[\ell]$ whose x -coordinates are roots of $g(x)$, even if $\deg g = 1$.

8.4 Algorithm to compute the trace of Frobenius modulo ℓ

We now give an algorithm to compute t_ℓ , the trace of Frobenius modulo ℓ .

Algorithm 8.3. Given $E : y^2 = f(x)$ over \mathbb{F}_q and an odd prime ℓ , compute t_ℓ as follows:

1. Compute the ℓ th division polynomial $h = \psi_\ell \in \mathbb{F}_q[x]$ for E .
2. Compute $\pi_\ell = (x^q \bmod h, (f^{(q-1)/2} \bmod h)y)$ and $\pi_\ell^2 = \pi_\ell \circ \pi_\ell$.
3. Use scalar multiplication to compute $q_\ell = q_\ell[1]_\ell$, and then compute $\pi_\ell^2 + q_\ell$.
(If a non-invertible denominator arises, update h and return to step 2).
4. Compute $0, \pi_\ell, 2\pi_\ell, 3\pi_\ell, \dots, c\pi_\ell$, until $c\pi_\ell = \pi_\ell^2 + q_\ell$.
(If a non-invertible denominator arises, update h and return to step 2).
5. Output $t_\ell = c$.

Throughout the algorithm, elements of $\text{End}(E[\ell])$ are represented in the form $(a(x), b(x)y)$, with $a, b \in R = \mathbb{F}_q[x]/(h(x))$, and all polynomial operations take place in the ring R . If a non-invertible denominator v is found in either steps 3 or 4 we replace h with whichever of $\gcd(h, v)$ and $h/\gcd(h, v)$ has lower degree; this guarantees that the degree of h is reduced by at least a factor of 2 (but see the next section for a further discussion).

The correctness of the algorithm follows from equation (1) and Lemma 8.2. The algorithm is guaranteed to find some $c\pi_\ell = \pi_\ell^2 + q_\ell$ in step 4 with $c < \ell$, since we know that $c = t_\ell$ works. Although we may be working modulo a proper factor g of h , every root x_0 of g is a root of h and therefore corresponds to a pair of nonzero points $P = (x_0, \pm y_0) \in E[\ell]$ for which $\pi_\ell^2(P) - c\pi_\ell(P) + q_\ell P = 0$ holds (there is at least one such root, since $\deg g > 0$), and Lemma 8.2 implies that we must have $c = t_\ell$.

The computation of the division polynomial in step 1 of the algorithm can be efficiently accomplished using the double-and-add approach described in Problem Set 3. You will have the opportunity to do a careful complexity analysis Algorithm 8.3 in the next problem set, but it is easy to see that its running time is polynomial in $n = \log q$ and ℓ : every operation involves polynomials over \mathbb{F}_q of degree less than ℓ^2 , in step 4 we can have at most ℓ iterations, and we can return to step 2 at most $2 \log \ell$ times (in fact this can happen only once). A simple implementation of the algorithm can be found in this [Sage notebook](#).

8.5 Factors of the division polynomial

As we saw when running our implementation of Schoof's algorithm in Sage, we do occasionally encounter non-invertible denominators and thereby obtain a proper factor g of the ℓ -division polynomial $h = \psi_\ell$. This is not too surprising, since there is no reason why h should necessarily be irreducible, but it is worth noting that whenever this occurs the degree of g is always exactly $(\ell - 1)/2$. Why is this the case?

Any point $P = (x_0, y_0) \in E(\overline{\mathbb{F}}_q)$ for which $g(x_0) = 0$ lies both in $E[\ell]$ and in the kernel of an endomorphism α (since x_0 is a root of the denominator of a rational function defining α). The point P is nonzero, so it generates a cyclic group C of order ℓ which must be a subgroup of $\ker \alpha$. It follows that over $\overline{\mathbb{F}}_q$ the polynomial g has at least $(\ell - 1)/2$ roots, one for each pair of nonzero points $(x_i, \pm y_i)$ in C (note that ℓ is odd). If g has any other roots, then there is point Q that lies in the intersection of $E[\ell] \cap \ker \alpha$ but not in C , in which case we must have $\ker \alpha = E[\ell]$, since $E[\ell]$ has ℓ -rank 2; but this is impossible because g is a proper factor of the ℓ -division polynomial h (whose roots are distinct because $\ell \nmid q$). So g must have exactly $(\ell - 1)/2$ roots in $\overline{\mathbb{F}}_q$. Reducing the polynomials that define our endomorphism modulo g corresponds to working in the subring $\text{End}(C)$ of $\text{End}(E[\ell])$.

If we are lucky enough to find such a proper factor g of h , our algorithm then speeds up by at least a factor of ℓ , since we are working modulo a polynomial of degree $(\ell - 1)/2$ rather than $(\ell^2 - 1)/2$. While we are fairly unlikely to stumble across such a g by chance, it turns out that in fact such a g exists for half of the primes ℓ (asymptotically speaking). Not long after Schoof published his result, Noam Elkies found a way to directly compute these polynomials, whose roots are the x -coordinates of points $P = (x_0, y_0)$ that lie in the kernel of a rational isogeny of degree ℓ . We will learn about Elkies' technique later in the course when we discuss modular polynomials. There is another optimization due to A.O.L. Atkin that applies to primes ℓ for which Elkies' optimization does not; together these yield what is known as the Schoof-Elkies-Atkin (SEA) algorithm.

8.6 Some historical remarks

When Schoof originally developed this algorithm, it was not clear to him that it had any practical use. This is in part because he (and others) were unduly pessimistic about its practical efficiency, in part because robust implementations of fast integer and polynomial arithmetic were not as widely available then as they are now. Even the simple Sage implementation given in the worksheet is already noticeably faster than the baby-steps giant-steps algorithm for $q \approx 2^{80}$ and can readily handle computations over fields of cryptographic size (it might take a day or two for $q \approx 2^{256}$, but this could be improved by at least an order of magnitude using a lower-level implementation in C or C++).

To better motivate his algorithm, Schoof gave an application that is of purely theoretical interest: he showed that it could be used to deterministically compute the square root of an integer a modulo a prime p in time that grows polynomially in $\log p$ when a is held fixed; we will see exactly how this works when we cover the theory of complex multiplication. Previously, no deterministic polynomial-time algorithm was known for this problem, unless one assumes the extended Riemann hypothesis. But Schoof's square-root application is really of no practical use; as we have seen, there are fast probabilistic algorithms to compute square roots modulo a prime, and unless the extended Riemann hypothesis is false, there are even deterministic algorithms that are much faster than Schoof's approach.

By contrast, in showing how to compute $\#E(\mathbb{F}_q)$ in polynomial-time, Schoof solved a practically important problem for which the best previously known algorithms were fully exponential (including randomized algorithms), despite the efforts of many experts working in the field. While perhaps not fully appreciated at the time, this has to be regarded as a major breakthrough, both from a theoretical and practical perspective. Improved versions of Schoof's algorithm (the SEA algorithm) are now the method of choice for computing $\#E(\mathbb{F}_q)$ in fields of large characteristic. In particular, the PARI library that is used by Sage includes an implementation of the SEA algorithm, and over 256-bit fields it takes only a few seconds to compute $\#E(\mathbb{F}_q)$. Today it is feasible to compute $\#E(\mathbb{F}_q)$ even when q is a prime with 5,000 decimal digits (over 16,000 bits), which represents the current record [5].

References

- [1] Joachim von zur Gathen and Jürgen Garhard, *Modern computer algebra*, third edition, Cambridge University Press, 2013.
- [2] Takakazu Satoh, *On p -adic point counting algorithms for elliptic curves over finite fields*, ANTS V, LNCS **2369** (2002), 43–66.
- [3] René Schoof, *Elliptic curves over finite fields and the computation of square roots mod p* . Mathematics of Computation **44** (1985), 483–495.
- [4] René Schoof, *Counting points on elliptic curves over finite fields*, Journal de Théorie des Nombres de Bordeaux **7** (1995), 219–254.
- [5] Andrew V. Sutherland, *On the evaluation of modular polynomials*, in Proceedings of the Tenth Algorithmic Number Theory Symposium (ANTS X), Open Book Series **1**, Mathematical Science Publishers, 2013, 531–555.