Schoof's Algorithm for Counting Points on $E(\mathbb{F}_q)$

Gregg Musiker

December 7, 2005

1 Introduction

In this write-up we discuss the problem of counting points on an elliptic curve over a finite field. Here, an elliptic curve E is the zero locus of an algebraic equation of a special form. Over a field of characteristic $\neq 2, 3$, this equation can be written as

$$y^2 = x^3 + Ax + B. \tag{1}$$

This is commonly known as **Weierstraß form**. To be rigorous, we would have to introduce projective coordinates to define the entire zero locus. However, in the case of an elliptic curve, the zero locus includes all the points (x, y) satisfying equation (1), these are known as affine points, plus *exactly one* additional point P_∞ "at infinity". While there is also an extensive theory of elliptic curves over $\mathbb{Q}$ and $\mathbb{C}$, in this presentation, we will focus on elliptic curves $E(\mathbb{F}_q)$ over a finite field $\mathbb{F}_q$, where $q = p^k$ and p is a prime greater than 3.

One of the important properties of elliptic curves is the existence of a group law

$$\oplus : E \times E \to E.$$

Given E as above, the addition on the curve is given as follows: For $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$, $P_1 \oplus P_2 = P_3 = (x_3, y_3)$ where

1) If $x_1 \neq x_2$ then

$$x_3 = m^2 - x_1 - x_2 \quad \text{and} \quad y_3 = m(x_1 - x_3) - y_1 \quad \text{with} \quad m = \frac{y_2 - y_1}{x_2 - x_1}.$$

2) If $x_1 = x_2$ but ($y_1 \neq y_2$, or $y_1 = 0 = y_2$) then $P_3 = P_\infty$.

3) If $P_1 = P_2$ and $y_1 \neq 0$, then

$$x_3 = m^2 - 2x_1 \quad \text{and} \quad y_3 = m(x_1 - x_3) - y_1 \quad \text{with} \quad m = \frac{3x_1^2 + A}{2y_1}.$$

4) P_∞ acts as the identity element in this addition.

Among other things, this allows elliptic curves to be used in cryptography, usually taking advantage of the difficulty of the discrete logarithm problem

for group $E(\mathbb{F}_q)$. Because of this application, an important computation for cryptographic purposes is the cardinality of $E(\mathbb{F}_q)$, which also happens to be the order of the group used for encryption. If the order of the group has small primes dividing it, then it leads to a weaker encryption scheme. Thus, if one is going to use elliptic curves for cryptography, it is important to have an efficient method to test the strength of the encryption scheme for a given curve E over $\mathbb{F}_q$.

While a few methods exist for computing the size of $|E(\mathbb{F}_q)|$ (as described in [5]), e.g. directly counting after running through possible values of x and y in $\mathbb{F}_q$, applying the formula

$$|E(\mathbb{F}_q)| = q + 1 + \sum_{\alpha \in \mathbb{F}_q} \left(\frac{x^3 + Ax + B}{q} \right),$$

Shank's Baby Step-Giant Step Method; the fastest one for extremely large primes is a variant of Schoof's algorithm.

Here I will first describe Schoof's original algorithm, as presented in [8] and summarized in [11]. After this warm-up, we will describe an improvement of Elkies which leads to a more efficient algorithm. There is in fact an additional improvement from Atkin to make Schoof-Elkies-Atkin (SEA) one of the fastest algorithm for counting the number of points on E over a large prime field, but its description will be outside the scope of this paper. (This full algorithm has been implemented in the computer algebra systems Magma, PARI, and SAGE. Please see [7] for details on the PARI implementation. In SAGE, you can use the command $E.sea(p)$.)

To begin, we note that the size of $E(\mathbb{F}_q)$ is very close to $q + 1$. This is not completely surprising since half of the elements of $\mathbb{F}_q^\times$ are squares, and the equation $y^2 = \alpha$ has two solutions if α is a square, one solution if $\alpha = 0$, and zero solutions if α is not a square in $\mathbb{F}_q^\times$. Thus $y^2 = x$ would have exactly q solutions of form (x_0, y_0) plus the point at infinity. Unfortunately, it is not always true that quantity $x^3 + Ax + B$, $x \in \mathbb{F}_q$, is a square half the time, but we at least expect it to be a square *close* to half the time. By analogous reasoning, we get that $y^2 = x^3 + Ax + B$ will have approximately $q + 1$ solutions, including P_∞.

We now proceed to make this statement more precise.

Theorem 1 (Hasse 1934). *Letting N_1 denote $|E(\mathbb{F}_q)|$, we obtain that*

$$|N_1 - q - 1| \leq 2\sqrt{q}.$$

This is a deep theorem, and can actually be shown to be equivalent to the Riemann Hypothesis for genus one function fields. We will give a partial proof of this theorem in Section 2. More details are in [11]. We have now simplified our problem by narrowing down $|E(\mathbb{F}_q)|$ to a finite (albeit large) set of possibilities.

Rene Schoof's insight was exploiting the fact we know that there is a finite range of possible values for the cardinality of $E(\mathbb{F}_q)$. Hence, if we had some way of computing the size of $E(\mathbb{F}_q)$ modulo N where $N > 4\sqrt{q}$, that would

be sufficient for determining $|E(\mathbb{F}_q)|$. There is not an efficient way to compute $|E(\mathbb{F}_q)|$ mod N directly for general N, however there is an efficient way to compute $|E(\mathbb{F}_q)|$ mod l for l prime.

Hence, we will compute $|E(\mathbb{F}_q)|$ mod l for l in a set S of primes such that $\prod_{l \in S} l = N > 4\sqrt{q}$. The Chinese Remainder Theorem, found described in numerous sources including [3], allows us to compute $|E(\mathbb{F}_q)|$ mod N given $|E(\mathbb{F}_q)|$ mod l for all $l \in S$.

We now proceed to describe how to efficiently compute $|E(\mathbb{F}_q)|$ mod l for a prime $l \neq$p. Note that requiring $p \notin S$ is no loss since we can pick a bigger prime to take its place to ensure the product is big enough.

To efficiently compute $|E(\mathbb{F}_q)|$ mod l, we will acutally find it easier to compute the values

$$t_l \equiv q + 1 - |E(\mathbb{F}_q)| \pmod{l}.$$

This computation will require a digression into the theory of the Frobenius map and Division Polynomials so we provide those now.

2 The Characteristic Equation for the Frobenius Endomorphism

Given an elliptic curve $E(\mathbb{F}_q)$ we consider $\overline{E}$ to be the curve with same defining equation as E, but now allowing points with coordinates in $\overline{\mathbb{F}_q}$, the algebraic closure of $\mathbb{F}_q$. Given curve $\overline{E}$, there exists a map, in fact its an endomorphism back onto $\overline{E}$, defined as

$$\pi : (x, y) \mapsto (x^q, y^q).$$

This map has several useful properties:

- For $P \in \overline{E}$, $\pi(P) \in \overline{E}$, thus it is in fact an endomorphism of $\overline{E}$ as claimed.
- For $P \in \overline{E}$, $\pi(P) = P$ if and only if $P \in E(\mathbb{F}_q)$. In particular $\pi(P_\infty) = P_\infty$.
- For $P, Q \in \overline{E}$, $\pi(P \oplus Q) = \pi(P) \oplus \pi(Q)$, thus π is compatible with the addition on the curve and in fact we can think of acting by π as multiplication in our endomorphism ring since the distributive law holds.

On $E(\mathbb{F}_q)$, the Frobenius map satisfies the characteristic equation $\pi - 1 = 0$. This begs the question of what kind of characteristic equation π satisfies on all of $\overline{E}$. In fact in satisfies a quadratic equation.

Theorem 2. *Frobenius map π satisfies the characteristic equation*

$$\pi^2 - t\pi + q = 0 \tag{2}$$

where $t = q + 1 - |E(\mathbb{F}_q)|$.

The proof of Theorem 2 is outside the scope of this presentation, although please see [11] for details. Nonetheless, we can use Theorem 2 to give a sketch of the proof of Theorem 1.

Proof. (Sketch) It can be shown that not only is $x^2 - tx + q$ the characteristic equation for the Frobenius map π, but furthermore, for all rational numbers $r/s \in \mathbb{Q}$,

$$(r/s)^2 - t(r/s) + q \geq 0. \tag{3}$$

This of course implies that $x^2 - tx + q \geq 0$ for all real numbers $x \in \mathbb{R}$. One way to show inequality (3) is by using π's characteristic equation to prove that for all $r, s \in \mathbb{Z}$, the quantity $r^2 - trs + qs^2$ equals deg $(r - s\pi)$, which is in $\mathbb{Z}_{\geq 0}$ by definition of degree. Once we have $x^2 - tx + q \geq 0$ for all real numbers $x \in \mathbb{R}$, Theorem 1 quickly follows by noting that the discriminant cannot be positive, hence

$$t^2 - 4q \leq 0 \Rightarrow |t| \leq 2\sqrt{q}.$$

□

Theorem 2 also immediately implies the following identity on the curve $\overline{E}$:

$$\text{For all } P = (x, y) \in \overline{E}, \quad (x^{q^2}, y^{q^2}) \oplus q(x, y) = t(x^q, y^q)$$

where scalar multiplication by t (or q) signifies adding a point to itself t (or q) times. We now spend the rest of this section, as well as the next two, on the problem of determining t_l, defined as t mod l, for a given prime $l \neq 2,\ p$. We therefore, refer to l as if it's been fixed, unless otherwise specified.

If point (x, y) is in the torsion subgroup $\overline{E}[l]$, meaning that $l(x, y) = P_\infty$, then the point $qP = \overline{q}P$ where $\overline{q}$ signifies q mod l, choosing $\overline{q}$ so that $|\overline{q}| < l/2$. Since $\pi(P_\infty) = P_\infty$, and $r \in \mathbb{Z}$ implies $r \cdot \pi(P) = \pi(rP)$ by repeated use of the distributive law, $\pi(P)$ will have the same order as P. Thus for $(x, y) \in \overline{E}[l]$, we also have $t(x^q, y^q) = \overline{t}(x^q, y^q)$ where $\overline{t}$ is t mod l. Hence we have reduced our problem to solving the equation

$$(x^{q^2}, y^{q^2}) \oplus \overline{q}(x, y) \equiv \overline{t}(x^q, y^q) \pmod{l} \tag{4}$$

for $\overline{t}$ (mod l).

The idea now is to explicitly compute $(x^{q^2}, y^{q^2}) \oplus \overline{q}(x, y)$ as a pair of rational functions (x', y') in terms of x and y. One method for finding t_l would involve plugging in $\overline{t} = 0, 1, 2, \ldots, l-1$ and find $\overline{t}$ such that the pair of rational functions given by $\overline{t}(x^q, y^q)$ are the same as the pair of rational functions on the left-hand-side, thus determining t_l. However, to do this more efficiently, we use division polynomials to allow us to compute multiples of a point P, and work with polynomials with bounded degrees instead of rational functions.

3 Division Polynomials

We define a sequence of polynomials in $\mathbb{Z}[x, y, A, B]$ via the following initial conditions and recurrence equations:

$$\begin{aligned}
\psi_0 &= 0 \\
\psi_1 &= 1 \\
\psi_2 &= 2y \\
\psi_3 &= 3x^4 + 6Ax^2 + 12Bx - A^2 \\
\psi_4 &= 4y(x^6 + 5Ax^4 + 20Bx^3 - 5A^2x^2 - 4ABx - 8B^2 - A^3) \\
&\cdots \\
\psi_{2m+1} &= \psi_{m+2}\psi_m^3 - \psi_{m-1}\psi_{m+1}^3 \text{ for } m \geq 2 \\
\psi_{2m} &= (\frac{\psi_m}{2y}) \cdot (\psi_{m+2}\psi_{m-1}^2 - \psi_{m-2}\psi_{m+1}^2) \text{ for } m \geq 2
\end{aligned}$$

The polynomial ψ_n is known as the nth Division Polynomial [6, 11]. These polynomials turn out to have the remarkable property that all of the finite n-torsion points (x_0, y_0), i.e. elements of $\overline{E}[n] \setminus \{P_\infty\}$, satisfy $\psi_n^2(x_0, y_0) = 0$. (Note that it can be shown using a variant of the characteristic equation (2) that the number of finite torsion points is exactly $n^2 - 1$.)

Additionally, we can define the multiple of a point, $r \cdot (x, y)$, as a pair of rational functions in terms of x and y using the ψ_n's. In particular, we have the following:

Proposition 1. *Let $P = (x, y)$ be a point on the elliptic curve $y^2 = x^3 + Ax + B$ over some field of characteristic $\neq 2$. Then for any positive integer n, $nP = P \oplus P \oplus P \oplus \cdots \oplus P$ is given by*

$$nP = \left(\frac{\phi_n(x)}{\psi_n^2(x)},\ \frac{\omega_n(x, y)}{\psi_n^3(x, y)}\right) = \left(x - \frac{\psi_{n-1}\psi_{n+1}}{\psi_n^2(x)},\ \frac{\psi_{2n}(x, y)}{2\psi_n^4(x)}\right).$$

$$-nP = \left(\frac{\phi_n(x)}{\psi_n^2(x)},\ -\frac{\omega_n(x, y)}{\psi_n^3(x, y)}\right) = \left(x - \frac{\psi_{n-1}\psi_{n+1}}{\psi_n^2(x)},\ -\frac{\psi_{2n}(x, y)}{2\psi_n^4(x)}\right)$$

where the polynomials ϕ_n and ω_n are defined as

$$\begin{aligned}
\phi_m &= x\psi_m^2 - \psi_{m+1}\psi_{m-1} \\
\omega_m &= \frac{\psi_{m+2}\psi_{m-1}^2 - \psi_{m-2}\psi_{m+1}^2}{4y}.
\end{aligned}$$

Note that using the equivalence relation $y^2 \equiv x^3 + Ax + B$ and the recurrence relations for ψ_{2m} and ψ_{2m+1}, we can inductively prove that

$$\psi_n^2,\ \frac{\psi_{2n}}{y},\ \psi_{2n+1},\ \text{and } \phi_n \text{ are all functions in terms of } x.$$

As a corollary, the x-coordinate of nP is a rational function strictly in terms of x, and the y-coordinate has the form $y \cdot \Theta(x)$. Proposition 1 is commonly proved using the theory of the Weirestraß $\mathfrak{P}$-function, as given in [6] and [11, Chapter 9].

We can summarize these results as follows: ψ^2 is a function in x alone and has degree n^2-1, which equals the number of finite n-torsion points. The degree of ψ^2 is easily verified via the above recurrence relations. Furthermore, if n is odd and $(x_0, y_0) \in \overline{E} \setminus \{P_\infty\}$, then

$$\psi_n(x_0) = 0 \text{ if and only if } (x_0, y_0) \in \overline{E}[n]. \tag{5}$$

If n is even, E defined by equation $y^2 = (x-\alpha_1)(x-\alpha_2)(x-\alpha_3)$ over $\overline{\mathbb{F}_q}$, and $(x_0, y_0) \in \overline{E} \setminus \{P_\infty,\ (\alpha_1, 0),\ (\alpha_2, 0),\ (\alpha_3, 0)\}$, then

$$\frac{\psi_n}{y}(x_0) = 0 \text{ if and only if } (x_0, y_0) \in \overline{E}[n]. \tag{6}$$

4 The Remainder of the Original Algorithm

With preliminaries out of the way, we proceed to present Schoof's algorithm for computing t_l, which is defined as $q+1-|E(\mathbb{F}_q)| \pmod{l}$, for $l \neq 2,\ p$. We recall our definition of $\overline{q}$ as a specific integer satisfying $\overline{q} \equiv q \pmod{l}$ and $|\overline{q}| < l/2$. First we use division polynomials to rewrite $\overline{q}(x, y)$ as a pair of rational functions in x and y:

$$(x_{\overline{q}}, y_{\overline{q}}) = \overline{q}(x, y) = \left(x - \frac{\psi_{\overline{q}-1}\psi_{\overline{q}+1}}{\psi_{\overline{q}}^2(x)},\ \frac{\psi_{2\overline{q}}(x, y)}{\psi_{\overline{q}}^4(x)} \right).$$

We then encounter an obstacle because the formula for computing

$$(x^{q^2}, y^{q^2}) \oplus \overline{q}(x, y),$$

given rational function representations of each of (x^{q^2}, y^{q^2}) and $(x_{\overline{q}}, y_{\overline{q}}) = \overline{q}(x, y)$, will depend on which of three cases, (e.g. are the points distinct, do they share the same x-coordinate), we are in. There is not an efficient enough way to determine which case we are to justify taking the time to check whether or not we happen to fall into the two uncommon special cases. Thus we might as well assume we happen to be in the case $(x^{q^2}, y^{q^2}) \neq \pm\overline{q}(x, y)$. If our guess is correct, which it will be most of the time, the proceeding algorithm will find t_l for us. If our guess is incorrect, our algorithm will find that no such $\overline{t}$ that will satisfy (4), hence alerting us that we are in the case $(x^{q^2}, y^{q^2}) = \overline{q}(x, y)$ or $(x^{q^2}, y^{q^2}) = -\overline{q}(x, y)$. We will prove the above assertions, and give the algorithm for these two cases in Section 4.1.

Thus, as prescribed, we now assume that $(x^{q^2}, y^{q^2}) \neq \pm\overline{q}(x, y)$ for some $(x, y) \in \overline{E}[l] \setminus \{P_\infty\}$. Hence, if we wish to compute

$$(x', y') = (x^{q^2}, y^{q^2}) \oplus \overline{q}(x, y),$$

we can use the usual addition formula and obtain

$$x' = \left(\frac{y^{q^2} - y_{\overline{q}}}{x^{q^2} - x_{\overline{q}}}\right)^2 - x^{q^2} - x_{\overline{q}}.$$

We get an analogous formula for y', but it is more computationally efficient to first use the x-coordinate to narrow down the choice of $\bar{t}$ to two possibilities. We must then determine which square-root of

$$\pm\sqrt{x'^3 + Ax' + B}$$

to use.

We recall from Section 3 that the y-coordinate of $\overline{q}(x, y)$, which we denote as $y_{\overline{q}}$, equals $y\Theta(x)$ and factor

$$(y^{q^2} - y\Theta(x))^2 = y^2(y^{q^2-1} - \Theta(x))^2 = (x^3+Ax+B)\left((x^3+Ax+B)^{\frac{q^2-1}{2}} - \Theta(x)\right)^2$$

to discover that x' is a rational function in variable x alone.

We will soon use the following fact: If $x' = x_{\bar{t}}^q$ for one point P in $\overline{E}[l] \setminus \{P_\infty\}$, then $\bar{t}$ satisfies

$$\pi^2(P) \ominus \bar{t}\pi(P) \oplus qP = P_\infty.$$

But since $\bar{t}$ in the characteristic equation is fixed, we obtain that this choice of $\bar{t}$ must indeed be t_l and we have proven $x' = x_{\bar{t}}^q$ for all points P in $\overline{E}[l] \setminus \{P_\infty\}$.

We recall that our goal is to solve (4) for $\bar{t}$. Taking x-coordinates of both sides, we obtain that the left-hand-side of (4) is x' and the right-hand-side is $x_{\bar{t}}^q$. We use our fact that for odd prime l, point $(x_0, y_0) \in \overline{E}[l] \setminus \{P_\infty\}$ if and only if $\psi_l(x_0) = 0$. Furthermore, the roots of ψ_l are simple by an easy counting argument. In particular, the degree of ψ_l^2 is exactly the number of finite l-torsion points.

Thus the equality $x_0' = x_{0_{\bar{t}}}^q$ at all points $(x_0, y_0) \in \overline{E}[l] \setminus \{P_\infty\}$ is equivalent to the statement

$$\psi_l(x) \Big| (x' - x_{\bar{t}}^q)$$

as a polynomial in x, modulo l. In conclusion, to solve (4) for $\bar{t}$, we need only solve

$$x' - x_{\bar{t}}^q \equiv 0 \pmod{\psi_l} \tag{7}$$

for $\bar{t}$.

This computation can be done efficiently for given $\bar{t} \in \{1, 2, \ldots, \frac{l-1}{2}\}$ by computing the power $x_{\bar{t}}^q$ by **successive squaring**. (Note that we can stop checking $\bar{t}$ at $\frac{l-1}{2}$ since we are only currently determining t_l up to additive inverse modulo l.) We create a table of

$$x_{\bar{t}}^{2^i} \pmod{\psi_l} \quad \text{for } i = 0, 1, 2, \ldots, \log_2 q,$$

each of which will be a polynomial of degree less than $\frac{l^2-1}{2} = \deg\ \psi_l$. We then can compute $x_{\bar{t}}^q \equiv x_{\bar{t}}^{2^{i_1}} \cdots x_{\bar{t}}^{2^{i_k}} \pmod{\psi_l}$ where $q = 2^{i_1} + \cdots + 2^{i_k}$ is the binary expansion of q.

Once we find $\bar{t}$ such that (7) is satisfied, we have found $\bar{t}$ such that

$$(x^{q^2}, y^{q^2}) \oplus \overline{q}(x, y) = \pm\bar{t}(x^q, y^q) \pmod{l}.$$

In particular $-(x^q, y^q) = (x^q, -y^q)$ thus the x-coordinate can only narrow down the possibilities of t_l to two.

To find whether t_l is $+\bar{t}$ or $-\bar{t}$ we check whether or not

$$(y' - y_{\bar{t}}^q)/y \equiv 0 \pmod{\psi_l}. \tag{8}$$

If so, then we choose $+\bar{t}$, otherwise $t_l = -\bar{t}$. This is sufficient for exactly the same reason as the sufficiency of (7).

4.1 The Cases $(x^{q^2}, y^{q^2}) = \pm\overline{q}(x, y)$

Our first important assertion is that if we falsely assumed points (x^{q^2}, y^{q^2}) and $\overline{q}(x, y) = (x_{\overline{q}}, y_{\overline{q}})$ had different x-coordinates, then the above procedure would not have found $\bar{t}$ such that (4) was satisfied.

Suppose on the contrary, that $\bar{t} \in \{1, 2, \ldots, \frac{l-1}{2}\}$ has satisfied (7) but that we have $(x^{q^2}, y^{q^2}) = \pm\overline{q}(x, y)$. Then

$$\bar{t}(x^q, y^q) \ominus (x^{q^2}, y^{q^2})$$

would be the same point as $\overline{q}(x, y) = \pm(x^{q^2}, y^{q^2})$, modulo l. But, this will lead to an immediate contradiction in the addition law.

So once we have failed to find a $\bar{t}$ we know that we have $(x^{q^2}, y^{q^2}) = \overline{q}(x, y)$ or $(x^{q^2}, y^{q^2}) = -\overline{q}(x, y)$. So now assume we are in this case. Let us assume further that we have

$$(x^{q^2}, y^{q^2}) = +\overline{q}(x, y). \tag{9}$$

Like before, our guess might be incorrect, but we will soon find out so. For $P = (x, y) \in \overline{E}[l] \setminus \{P_\infty\}$ satisfying (9), we obtain via the characteristic equation (2) modulo l that

$$\bar{t}\pi(P) = 2\overline{q}P.$$

By (9), we also have

$$\overline{t^2}\overline{q}P = \overline{t^2}\pi^2(P) = \bar{t}\pi(\bar{t}\pi(P))$$

which equals $(2\overline{q})^2 P$. Consequently,

$$\bar{t}^2\overline{q}P \equiv (2q)^2 P \pmod{l}$$

and thus $\overline{q}$ is a square modulo l unless $\bar{t} \equiv 0 \pmod{l}$. However, that would imply $2\overline{q}P = P_\infty$, which is a contridiction since $P \in \overline{E}[l]$ and $\gcd(q, l)$ assumed to be 1.

Once we know q is a square, the rest follows easily. We find q's square-roots over $\mathbb{F}_l$ efficiently by using $\gcd(x^2 - q, x^l - x)$ or otherwise. If $q \equiv w^2 \pmod{l}$, we find t_l will be $\pm 2w \bmod l$ depending on the y-coordinate, i.e. the rational function y_w.

$$(X \pm w)(X \pm w) = X^2 \pm 2wX + q. \tag{10}$$

If y_w matches y^q, then $t_l = 2w$, otherwise $t_l = -2w$.

If q happened not to be a square, then our second assumption was also false, meaning we were in the third case

$$(x^{q^2}, y^{q^2}) = -\overline{q}(x, y).$$

Once we know we are in this csae, it is also easy since we immediately find

$$(x^{q^2}, y^{q^2}) \oplus \overline{q}(x, y) = P_\infty$$

since they were additive inverses. Hence, $t_l \equiv 0$ in this case.

However, we still have to worry about the case of a false positive, i.e. it is possible for q to be a square mod l but $(x^{q^2}, y^{q^2}) = -\overline{q}(x, y)$ nonetheless. It is sufficient to check whether or not either square root $\pm w$ satisfies $\pi(P) = \pm wP$, for $P \in \overline{E}[l] \setminus \{P_\infty\}$, hence leading to the factorization of the characteristic equation for π as in (10).

The simple calculation of

$$\gcd(\text{numerator}(x^q - x_w), \phi_l)$$

and the test of whether or not it is 1 suffices. We have $\pi(P) = \pm wP$ for some $P \in \overline{E}[l] \setminus \{P_\infty\}$ if and only if the gcd $\neq 1$.

4.2 Case of l=2

Note that the above description actually works only for odd $l \neq p$. However, we can also run an analogous procedure for $l = 2$. This allows us to choose set S to contain slightly smaller primes. Since we assume q odd,

$$q + 1 - t \equiv t \pmod{2}$$

and in particular, $t_2 \equiv 0 \pmod{2}$ if and only if $E(\mathbb{F}_q)$ has an element of order 2. By definition of $\oplus$, any element of order 2 must be of the form $(x_0, 0)$. Thus $t_2 \equiv 0 \pmod{2}$ if and only if $x^3 + Ax + B$ has a root in $\mathbb{F}_q$. An efficient way to check this is by taking the $\gcd(x^q - x, x^3 + Ax + B)$. Here we compute x^q efficiently by taking successive squares modulo $x^3 + Ax + B$. In summary $t_2 \equiv 0$ if and only if $\gcd(x^q - x, x^3 + Ax + B) \neq 1$.

4.3 Summary

We can summarize the procedure as follows:

1) Choose a set of primes S, with $p \notin S$ such that $\prod_{l \in S} l > 4\sqrt{q}$.

2) For $l = 2$, we find $t_l \equiv 0$ if and only if $\gcd(x^q - x, x^3 + Ax + B) \neq 1$.

3) For $l \in S \setminus \{2\}$, do the following:

a) Let $\overline{q}$ be the unique integer satisfying $\overline{q} \equiv q \pmod{l}$ and $|\overline{q}| < \frac{l}{2}$.

b) Compute univariate rational function x' as defined above. We can even work modulo ψ_l so we can use polynomials of bounded degree instead.

c) For $\overline{t} \in \{1, 2, \ldots, \frac{l-1}{2}\}$ do:

i) Check if $x' - x_{\overline{t}}^q \equiv 0$ modulo ψ_l. If so, go to step (ii). Otherwise, try the next value of $\overline{t}$. If you have unsuccessfully tried $\frac{l-1}{2}$, go to step (d) instead.

ii) Check whether or not $(y' - y_{\overline{t}}^q)/y \equiv 0$ modulo ψ_l. If so, then $t_l \equiv \overline{t}$. Otherwise, $t_l \equiv -\overline{t}$.

d) Find q's square roots modulo l, if they exist. If they don't exist then $t_l \equiv 0$. Otherwise write $q \equiv w^2 \pmod{l}$.

e) Check whether or not $\gcd(\text{numerator}(x^q - x_w), \psi_l) = 1$. If so, $t_l \equiv 0$. Otherwise go to step (f).

f) Check whether or not $\gcd(\text{numerator}((y^q - y_w)/y), \psi_l) = 1$. If this gcd is 1, then $t_l \equiv -2w$. Otherwise, $t_l \equiv 2w$.

4) We now have computed t_l for all $l \in S$. Thus we know $|E(\mathbb{F}_q)| \equiv 1 + q - t_l \pmod{l}$ for every l in S. Using the Chinese Remainder Theorem, we obtain $|E(\mathbb{F}_q)|$ modulo N, where $N = \prod_{l \in S} l$. Since set S was chosen so that $N > 4\sqrt{q}$, by Hasse's theorem, we in fact know $|E(\mathbb{F}_q)|$ precisely.

As described in the introduction, Schoof's algorithm was a huge improvement over previous methods, having an asymptotic running times of $(\log q)^8$. In the 1990's Noam Elkies discovered an improvement to this algorithm that yields a faster running time. The trick is to restrict set S to primes of a certain type, now known as Elkies primes, and to use modular polynomials instead of division polynomials. Before describing this improvement, we will first introduce modular polynomials.

5 Modular Polynomials

Modular Polynomials come from the theory of modular forms and an interpretation of elliptic curves over $\mathbb{C}$ as lattices. Being a lattice, the matrix group $SL_2(\mathbb{Z})$ naturally acts on $\mathbb{C}/(\mathbb{Z} + \mathbb{Z}\tau)$. Two lattices are considered to be equivalent if there is a natural way to re-scale and rotate one to get the other. It turns out that two elliptic curves are isomorphic as groups if their associated $\mathbb{C}$-lattices are equivalent. Thus a certain invariant, known as the **j-invariant**, of lattice theory is in fact an invariant of isomorphic elliptic curves as well. For

E given in Weierstraß form (1), we have

$$j(E) = 1728\frac{4A^3}{4A^3 + 27B^2}.$$

5.1 Thinking in terms of cosets

We will now go back and forth between $j(E)$ and $j(z)$, where we use $j(z)$ to refer to the j-invariant of a lattice with basis $\{\omega_1, \omega_2\}$ such that $z = \frac{\omega_1}{\omega_2} \in \mathbb{C}$. The matrix groups $GL_2(\mathbb{Z})$ and $SL_2(\mathbb{Z})$ naturally act on $\mathbb{C}$ so it makes sense to use notation such as $j(M \cdot z)$ for $M \in GL_2(\mathbb{Z})$ or $SL_2(\mathbb{Z})$.

Proposition 2. *[10] If M is a 2×2 integer matrix with $\det M = m \in \mathbb{Z}_{>0}$ then $j(M \cdot z)$ and $j(z)$ are algebraically related, meaning there exists a bivariate polynomial $\Phi_m(x, y) \in \mathbb{Z}[x, y]$ such that $\Phi_m\Big(j(M \cdot z), j(z)\Big) = 0$.*

Proof. (Sketch)

If we consider the group M_m of all matrices of determinant m, we obtain that M_m has a left-coset decomposition as

$$M_m = \bigcup_{a,b,d} SL_2(\mathbb{Z}) \cdot \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$$

where we add the restrictions $a > 0,\ d > 0,\ ad = m$, and $0 \leq b < d$.

Since the set of j-invariants

$$\Big\{ j(B \cdot z) \ : \ B \in M_m \Big\} = \Big\{ j(BA \cdot z) \ : \ B \in M_m \Big\} \text{ for } A \in SL_2(\mathbb{Z}),$$

we find that

$$\Phi_m(x, j(z)) = \prod_{M \in M_m} \Big(x - j(M \cdot z) \Big) = \Big(x - j\Big(\frac{az + c}{d} \Big) \Big)$$

satisfies the desired properties. Note that it can be shown that this polynomial is integral. □

6 Sketch of Elkies' Improvement

Given the characteristic equation for the Frobenius map,

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

we define a prime l to be an **Elkies prime** if this equation splits over $\mathbb{F}_l$. This is equivalent to whether or not the discriminant $t^2 - 4q$ is a square modulo l. We define l to be an **Atkin prime** if it is not an Elkies prime.

Remark 1. *Note that the condition of being an Elkies or Atkin prime is contingent on the choice of curve E and field $\mathbb{F}_q$ which determine a specific characteristic equation for π. Since we have assumed a fixed choice of E and $\mathbb{F}_q$ all along, we will from here on out simply refer to primes l as Elkies or Atkin primes, without further specification, for the sake of expository convenience.*

Since our goal is to compute t, which is required to test what type a given prime is, this definition has yet to improve the algorithm; however this is where the modular polynomial Φ_l for prime l plays a role. The following deep result plays a key role. We omit its proof but please see [1] or [2] for details.

Proposition 3. *A prime l is an Elkies prime for curve E over $\mathbb{F}_q$ if and only if $\Phi_l(x, j(E))$ has a root for $x \in \mathbb{F}_q$, where $j(E)$ is the j-invariant for curve E.*

It is efficient to test whether or not prime l is an Elkies prime or not: we take the gcd of $\Phi_l(x, j(E))$ with $x^q - x$ as in Section 4.2. If l is not an Elkies prime, we do not include it in our set S. Furthermore, roughly half the primes are Elkies primes (a non-trivial but known fact), so this will not increase the size of the primes in S too significantly. (In all this, we always assume that $l \neq 2, p$ where p is the characteristic of $\mathbb{F}_q$.) The added efficiency for computing $t \pmod{l}$ for only Elkies primes l will dominate the reduced efficiency for using larger primes l.

In particular, when l is an Elkies prime, we will be able to solve the equation $\pi(x, y) = (x^q, y^q) = \lambda(x, y)$ for λ using a degree $\frac{l-1}{2}$ factor of the degree $\frac{l^2-1}{2}$ division polynomial ψ_l. By virtue of being an Elkies prime, the characteristic equation for π splits modulo l, and it follows that

$$\pi^2 - t\pi + q = (\pi - \lambda t)(\pi - \overline{\lambda} t)$$

such that $\lambda\overline{\lambda} = q$. Consequently,

$$t_l \equiv \lambda + \frac{q}{\lambda} \pmod{l}.$$

Thus to compute t_l, it is sufficient to compute an eigenvalue $\lambda \in \{1, 2, \ldots, l-1\}$ of the Frobenius map, and hence that is now our current goal. For this purpose, we consider a cyclic subgroup of $E(\mathbb{F}_q)$, which we denote by C which is fixed by the action of the Frobenius map. Thus the polynomial

$$F_l(x) = \prod_{\pm P_i \in C \setminus \{P_\infty\}} (x - (P_i)_X)$$

has integral coefficients, hence is defined over $\mathbb{F}_q$. Here this product is meant to only include one out of each of the pairs $\pm P_i$, since both have the same x-coordinate. The notation $(P_i)_X$ signifies the x-coordinate of point P_i. Note that there are efficient methods for explicitly computing $F_l(x)$, we refer the reader to Section VIII.4 of [1].

We now can work modulo $F_l(x)$ instead of modulo ψ_l which has smaller degree, $\frac{l-1}{2}$, hence is more efficient; i.e. instead of needing to solve equation (7), it suffices to solve

$$x' - x_\lambda^q \equiv 0 \pmod{F_l} \tag{11}$$

for λ.

We end with the note that Atkin's improvement involves the efficient extraction of information about t_l even for l a non-Elkies prime, i.e. an Atkin prime. However the description of this algorithm is a topic for another day. Please see [1] or [9] for details.

References

[1] I. Blake, G. Seroussi, and N. Smart. *Elliptic Curves in Crytography*, volume 265 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2000. Reprint of the 1999 original.

[2] L. Dewaghe. Remarks on the Schoof-Elkies-Atkin Algorithm. *Math. Comp.,* **67**(223):1247-1252, 1998.

[3] D. Dummit and R. Foote. *Abstract Algebra, 2nd Edition.* Prentice Hall, Upper Saddle River, 1999.

[4] A. Enge. *Elliptic Curves and their Applications to Cryptography: An Introduction.* Kluwer Academic Publishers, Dordrecht, 1999.

[5] N. Harris. *Math 168 Project.* University of California, San Diego, Fall 2005.

[6] S. Lang. *Elliptic Curves: Diophantine Analysis.* Springer-Verlag, Berlin, 1978.

[7] http://pari.math.u-bordeaux.fr/archives/pari-announce-05/msg00002.html

[8] R. Schoof. Elliptic Curves over Finite Fields and the Computation of Square Roots mod *p*. *Math. Comp.,* **44**(170):483-494, 1985.

[9] R. Schoof. Counting Points on Elliptic Curves over Finite Fields. *J. Thèor. Nombres Bordeaux* **7**:219-254, 1995.

[10] H. Stark. *Lecture notes for Math 204.* University of California, San Diego, Fall 2005.

[11] L. C. Washington, *Elliptic Curves: Number Theory and Cryptography.* Chapman & Hall/CRC, New York, 2003.