7 Point counting

7.1 Separable and inseparable endomorphisms

Recall that the Frobenius endomorphism π_E is inseparable. In order to prove Hasse's theorem we will need to use the fact that $\pi_E - 1$ is separable. This follows from a much more general result: adding a separable isogeny to an inseparable isogeny always yields a separable isogeny. Note that the sum of two separable isogenies need not be separable: in characteristic $p > 0$, if we have $a + b = p$ and both a and b prime to p, then $[a]$ and $[b]$ are both separable but $[a] + [b] = [a + b] = [p]$ is inseparable.

Lemma 7.1. *Let α and β be isogenies from E_1 to E_2, with α inseparable. Then $\alpha + \beta$ is inseparable if and only if β is inseparable.*

Proof. If β is inseparable then by Corollary 5.4 we can write $\alpha = \alpha_{\text{sep}} \circ \pi^m$ and $\beta = \beta_{\text{sep}} \circ \pi^n$, where π is the p-power Frobenius map and $m, n > 0$. We then have

$$\alpha + \beta = \alpha_{\text{sep}} \circ \pi^m + \beta_{\text{sep}} \circ \pi^n = (\alpha_{\text{sep}} \circ \pi^{m-1} + \beta_{\text{sep}} \circ \pi^{n-1}) \circ \pi,$$

which is inseparable (any composition involving an inseparable isogeny is inseparable because inseparable degrees multiply). If $\alpha+\beta$ is inseparable, then so is $-(\alpha+\beta)$, and $\alpha-(\alpha+\beta) = \beta$ is a sum of inseparable isogenies, which we have just shown is inseparable. □

Remark 7.2. Since the composition of an inseparable isogeny with any isogeny is always inseparable, Lemma 7.1 implies that the inseparable endomorphisms in $\text{End}(E)$ form an ideal (provided we view 0 as inseparable, which we do).

7.2 Hasse's Theorem

We are now ready to prove Hasse's theorem.

Theorem 7.3 (Hasse)**.** *Let $E/\mathbb{F}_q$ be an elliptic curve over a finite field. Then*

$$\#E(\mathbb{F}_q) = q + 1 - t,$$

where $t := \operatorname{tr} \pi_E$ is the trace of the Frobenius endomorphism π_E and $|t| \le 2\sqrt{q}$.

Proof. Recall that we defined $\mathbb{F}_q$ as the splitting field of $x^q - x$ over $\mathbb{F}_p$, where $p = \text{char}(\mathbb{F}_q)$, thus $\mathbb{F}_q = \{\alpha \in \overline{\mathbb{F}}_p : \alpha^q - \alpha = 0\} = \{\alpha \in \overline{\mathbb{F}}_q : \alpha^q - \alpha = 0\}$ is precisely the subfield of $\overline{\mathbb{F}}_q$ fixed by the q-power Frobenius automorphism $x \mapsto x^q$. The Frobenius endomorphism $\pi_E \colon E \to E$ is defined by $\pi_E(x : y : z) = (x^q : y^q : z^q)$, therefore

$$E(\mathbb{F}_q) = \{P \in E(\overline{\mathbb{F}}_q) : \pi_E(P) = P\} = \{P \in E(\overline{\mathbb{F}}_q) : \pi_E(P) - P = 0\} = \ker(\pi_E - 1),$$

where 1 denotes the multiplication-by-1 map $[1] \in \text{End}(E)$. The Frobenius endomorphism π_E is inseparable and -1 is separable, so by Lemma 7.1 the endomorphism $\pi_E - 1$ is separable, thus the cardinality of its kernel is equal to its degree (by Theorem 5.8). Therefore

$$\#E(\mathbb{F}_q) = \#\ker(\pi_E - 1) = \deg(\pi_E - 1) = \widehat{(\pi_E - 1)}(\pi_E - 1) = \hat{\pi}_E\pi_E + 1 - (\hat{\pi}_E + \pi_E) = q + 1 - t.$$

It remains only to show that $|t| \le 2\sqrt{q}$.

Consider the endomorphism $r\pi_E - s$ for $r, s \in \mathbb{Z}$ with $s \neq 0$. We have

$$\begin{aligned}\deg(r\pi_E - s) &= \widehat{(r\pi_E - s)}(r\pi_E - s) = (\hat{\pi}_E\hat{r} - \hat{s})(r\pi_E - s) = (\hat{\pi}_E r - s)(r\pi_E - s)\\ &= \hat{\pi}_E r^2\pi_E - \hat{\pi}_E rs - sr\pi_E + s^2 = r^2\hat{\pi}_E\pi_E - rs(\hat{\pi}_E + \pi_E) + s^2\\ &= r^2 \deg \pi_E - rs \operatorname{tr} \pi_E + s^2\\ &= r^2 q - rst + s^2,\end{aligned}$$

where we have used Lemmas 6.11 and 6.12, and the fact that $\mathbb{Z}$ is in the center of $\mathrm{End}(E)$. Dividing by s^2 and noting that $\deg(r\pi_E - s) \geq 0$ yields the inequality

$$q\,(r/s)^2 - t\,(r/s) + 1 \geq 0,$$

valid for all rational numbers r/s. Now $\mathbb{Q}$ is dense in $\mathbb{R}$, so we must have $qx^2 - tx + 1 \geq 0$ for all real numbers x. It follows that the discriminant $t^2 - 4q$ cannot be positive, which yields the desired bound $|t| \leq 2\sqrt{q}$. □

Recall that for an odd prime p the Legendre symbol $\left(\frac{a}{p}\right)$ is defined by

$$\left(\frac{a}{p}\right) = \left\{\begin{array}{ll} 1 & \text{if } y^2 = a \text{ has two solutions mod } p\\ 0 & \text{if } y^2 = a \text{ has one solution mod } p\\ -1 & \text{if } y^2 = a \text{ has no solutions mod } p\end{array}\right\} = \#\{\alpha \in \mathbb{F}_p : \alpha^2 = a\} - 1.$$

We extend the Legendre symbol to all finite fields $\mathbb{F}_q$ of odd characteristic by defining

$$\left(\frac{a}{\mathbb{F}_q}\right) = \#\{\alpha \in \mathbb{F}_q : \alpha^2 = a\} - 1 \in \{-1, 0, 1\}.$$

Thus $1 + \left(\frac{a}{\mathbb{F}_q}\right)$ counts the solutions to $y^2 = a$ in $\mathbb{F}_q$. It follows that if $E/\mathbb{F}_q$ is given by the Weierstrass equation $y^2 = x^3 + Ax + B$, then

$$\begin{aligned}\#E(\mathbb{F}_q) &= 1 + \sum_{x_0 \in \mathbb{F}_q}\left(1 + \left(\frac{x_0^3 + Ax_0 + B}{\mathbb{F}_q}\right)\right)\\ &= q + 1 + \sum_{x_0 \in \mathbb{F}_q}\left(\frac{x_0^3 + Ax_0 + B}{\mathbb{F}_q}\right). \qquad (1)\end{aligned}$$

Hasse's Theorem is equivalent to the statement that the sum in (1) has absolute value at most $2\sqrt{q}$. This is remarkable for a sum with q terms, almost all of which are ± 1. From a probabilistic point of view, one might expect that *on average* an $O(\sqrt{q})$ bound should hold, but Hasse's theorem guarantees that it *always* holds.

The bound in Hasse's theorem is the best possible. Later in the course we will see how to explicitly construct elliptic curves $E/\mathbb{F}_q$ with cardinalities matching every integer value in the *Hasse interval*

$$\mathcal{H}(q) := [q + 1 - 2\sqrt{q},\ q + 1 + 2\sqrt{q}] = [(\sqrt{q} - 1)^2, (\sqrt{q} + 1)^2]$$

when q is prime, and all but at most two integers when q is not prime.

7.3 Point counting

We now consider the problem of computing the cardinality of $E(\mathbb{F}_q)$, which is crucial to cryptographic applications; as we shall see, it is quite important to know the cardinality of the group one is working in. The most naïve approach one might take would be to evaluate the curve equation $y^2 = x^3 + Ax + B$ for E at every pair $(x_0, y_0) \in \mathbb{F}_q^2$, count the number of solutions, and add 1 for the point at infinity. This takes $O(q^2\,\mathsf{M}(\log q))$ time. Note that the input to this problem is the pair of coefficients $A, B \in \mathbb{F}_q$, which each have $O(n)$ bits, where $n = \log q$. Thus in terms of the size of its input, this algorithm takes

$$O\big(\exp(2n)\,\mathsf{M}(n)\big)$$

time, which is obviously exponential in n.

A slightly less naïve approach is to precompute a table of quadratic residues in $\mathbb{F}_q$ so that we can very quickly compute the extended Legendre symbol $\left(\frac{\cdot}{\mathbb{F}_q}\right)$. We can construct such a table in $O(q\,\mathsf{M}(\log q))$ time, and then compute

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

in $O(q\,\mathsf{M}(\log q))$ time, yielding a total running time of

$$O(\exp(n)\,\mathsf{M}(n)).$$

So far we have not taken advantage of Hasse's theorem which gives us an interval $\mathcal{H}(q)$ of width $4\sqrt{q}$ which we know must contain the integer $\#E(\mathbb{F}_q)$ we wish to determine.

7.4 Computing the order of a point

Before giving an algorithm to compute $\#E(\mathbb{F}_q)$ using Hasse's theorem, let us first consider an easier problem: computing the order $|P|$ of a single point $P \in E(\mathbb{F}_q)$. Since the order of the group $E(\mathbb{F}_q)$ lies in $\mathcal{H}(q)$, we know that $\mathcal{H}(q)$ contains at least one integer M for which $MP = 0$, namely $M = \#E(\mathbb{F}_q)$, and any such M is a multiple of $|P|$. To find such an M, we set $M_0 = \lceil(\sqrt{q} - 1)^2\rceil$, compute M_0P using double-and-add scalar multiplication, and then generate the sequence of points

$$M_0P,\ (M_0 + 1)P,\ (M_0 + 2)P,\ \ldots,\ MP = 0,$$

by adding P repeatedly. Note that M is bounded by $M_0 + 4\sqrt{q}$, so $4\sqrt{q}$ additions suffice.

We then compute the prime factorization $M = p_1^{e_1} \cdots p_w^{e_w}$ (easy, compared to the time to find M, we could even use trial division). To compute the exact order of the point P we use the following generic algorithm.

Algorithm 7.4. Given an element P of an additive group and the prime factorization $M = p_1^{e_1} \cdots p_r^{e_r}$ of an integer M for which $MP = 0$, compute the order of P as follows:

1. Let $m = M = p_1^{e_1} \cdots p_r^{e_r}$.
2. For each prime p_i, while $p_i | m$ and $(m/p_i)P = 0$, replace m by m/p_i.
3. Output m.

When this procedure is complete we know that $mP = 0$ and $(m/p)P \neq 0$ for every prime p dividing m; this implies that $m = |P|$. You will analyze the efficiency of this algorithm and develop several improvements to it in Problem Set 4, but the number of group operations is clearly polynomial in $\log M$, which is all we need for the moment.

The time to compute $|P|$ is thus dominated by the time to find a multiple of $|P|$ in $\mathcal{H}(q)$. This involves $O(\sqrt{q})$ operations in $E(\mathbb{F}_q)$, yielding a bit complexity of $O(\sqrt{q}\ \mathsf{M}(\log q))$ or

$$O\big(\exp(n/2)\,\mathsf{M}(n)\big),$$

assuming that we use projective coordinates to avoid field inversions when adding points.

We will shortly see how this can be further improved, but first let us consider how to use our algorithm for computing $|P|$ to compute $\#E(\mathbb{F}_q)$. If we are lucky (and if q is large we almost always will be), the first multiple M of $|P|$ that we find in $\mathcal{H}(q)$ will actually be the only multiple of $|P|$ in $\mathcal{H}(q)$. If this happens, then we must have $M = \#E(\mathbb{F}_q)$. Otherwise, we might try our luck with a different point P. If we can find a combination of points for which the least common multiple of their orders has a unique multiple in $\mathcal{H}(q)$, then we can determine the group order. Unfortunately this will not always be possible, but before addressing that issue, let us consider the question of how long it might take to compute the least common multiple of the orders of *all* the points in $E(\mathbb{F}_q)$, which is a lot less than one might expect.

7.5 The group exponent

Definition 7.5. For a finite group G, the *exponent* of G, denoted $\lambda(G)$, is defined by

$$\lambda(G) = \operatorname{lcm}\{|\alpha| : \alpha \in G\}.$$

Note that $\lambda(G)$ is a divisor of $\#G$ and is divisible by the order of every element of G. Thus $\lambda(G)$ is the maximal possible order of an element of G, and when G is abelian this maximum is achieved: there exists an element $\alpha \in G$ with order $|\alpha| = \lambda(G)$. To see this, note that the structure theorem for finite abelian groups allows us to decompose G as

$$G \simeq \mathbb{Z}/n_1\mathbb{Z} \oplus \mathbb{Z}/n_2\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/n_r\mathbb{Z},$$

with $n_i | n_{i+1}$ for $1 \leq i < r$. Thus $\lambda(G) = n_r$, and any $\alpha = (\alpha_1, \ldots, \alpha_r) \in G$ for which α_r is a generator for $\mathbb{Z}/n_r\mathbb{Z}$ will necessarily satisfy $|\alpha| = \lambda(G)$.

Rather than searching for a single α with maximal order, it is enough to find any set of elements $S \subseteq G$ for which $\operatorname{lcm}\{|\alpha| : \alpha \in S\} = \lambda(G)$. If we choose S randomly, how large does it need to be to have a good chance of determining $\lambda(G)$? The answer is surprisingly small: for $|S| = 2$ we already have a better than $50/50$ chance.

Theorem 7.6. *Let G be a finite abelian group with exponent $\lambda(G)$. Let α and β be uniformly distributed random elements of G. Then*

$$\Pr[\operatorname{lcm}(|\alpha|, |\beta|) = \lambda(G)] > \frac{6}{\pi^2}.$$

Proof. We first reduce to the case that G is cyclic. As noted above, $G \simeq \mathbb{Z}/n_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/n_r\mathbb{Z}$ with $n_i | n_{i+1}$ and $\lambda(G) = n_r$. Let α_r and β_r be the projections of α and β to $\mathbb{Z}/n_r\mathbb{Z}$. Then $\operatorname{lcm}(|\alpha_r|, |\beta_r|) = \lambda(G)$ certainly implies $\operatorname{lcm}(|\alpha|, |\beta|) = \lambda(G)$, thus

$$\Pr[\operatorname{lcm}(|\alpha|, |\beta|) = \lambda(G)] \geq \Pr[\operatorname{lcm}(|\alpha_r|, |\beta_r|) = \lambda(G)],$$

and in the worst case G is cyclic and this inequality is an equality, which we now assume.

So let $G = \langle\gamma\rangle$ and let $p_1^{e_1} \cdots p_k^{e_k}$ be the prime factorization of $|\gamma| = \lambda(G) = \#G$. Then $\alpha = a\gamma$, with $0 \le a < |\gamma|$, and unless a is divisible by p_i, which occurs with probability $1/p_i$, the order of α will be divisible by $p_i^{e_i}$ (and similarly for β). The two probabilities for α and β are independent, thus with probability $1 - 1/p_i^2$ at least one of α and β has order divisible by $p_i^{e_i}$. Call this event E_i. The events $E_1, \ldots, E_k$ are independent, since we may write G as a direct sum of cyclic groups of prime-power orders $p_1^{e_1}, \ldots, p_k^{e_k}$, and the projections of α and β to each of these cyclic groups are uniformly and independently distributed. Thus

$$\begin{aligned}\Pr[\operatorname{lcm}(|\alpha|, |\beta|) = \lambda(G)] &= \Pr[E_1 \cap \cdots \cap E_k] \\ &= \prod_{p|\lambda(G)} (1 - p^{-2}) > \prod_p (1 - p^{-2}) = \left(\sum_{n=1}^{\infty} \frac{1}{n^2}\right)^{-1} = \frac{1}{\zeta(2)} = \frac{6}{\pi^2},\end{aligned}$$

where $\zeta(s) = \sum n^{-s}$ is the Riemann zeta function. □

Theorem 7.6 implies that if we generate random points $P \in E(\mathbb{F}_q)$ and accumulate the least common multiple N of their orders, we should expect to obtain $\lambda(E(\mathbb{F}_q))$ within $O(1)$ iterations. Regardless of when we obtain $\lambda(E(\mathbb{F}_q))$, at every stage we know that N divides $\#E(\mathbb{F}_q)$, and if we ever find that N has a unique multiple M in the Hasse interval $\mathcal{H}(q)$, then we know that $\#E(\mathbb{F}_q) = M$.

Unfortunately this might not ever happen; it can happen that $\lambda(E(\mathbb{F}_q)) \le 4\sqrt{q}$, in which case it is possible for $\lambda(E(\mathbb{F}_q))$ to have more than one multiple in $\mathcal{H}(q)$. To deal with this problem we need to consider the *quadratic twist* of E, which you saw on Problem Set 1.

7.6 The quadratic twist of an elliptic curve

Suppose s is an element of $\mathbb{F}_q$ that is *not* a square, meaning that $\left(\frac{s}{\mathbb{F}_q}\right) = -1$. If we consider the elliptic curve $\widetilde{E}$ defined by $sy^2 = x^3 + Ax + B$, then the affine point (x, y) will lie on the curve if and only if $x^3 + Ax + B$ is *not* a square. Thus

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

and it follows that if $\#E(\mathbb{F}_q) = q+1-t$, then $\#\widetilde{E}(\mathbb{F}_q) = q+1+t$. The curve $\widetilde{E}$ is called the *quadratic twist* of E (by s). We can put the curve equation for $\widetilde{E}$ in standard Weierstrass form by substituting x/s for x and y/s^2 for y and then clearing denominators, yielding

$$y^2 = x^3 + s^2Ax + s^3B.$$

Notice that it does not matter which non-residue s we choose. As you showed in Problem Set 1, if s and s' are any two non-squares in $\mathbb{F}_q$, then the corresponding curves $\widetilde{E}$ and $\widetilde{E'}$ are isomorphic over $\mathbb{F}_q$; thus we refer to $\widetilde{E}$ as "the" quadratic twist of E.[1]

Our interest in the quadratic twist of E lies in the fact that

$$\#E(\mathbb{F}_q) + \#\widetilde{E}(\mathbb{F}_q) = 2q + 2.$$

Thus if we can compute either $\#E(\mathbb{F}_q)$ or $\#\widetilde{E}(\mathbb{F}_q)$, we can easily determine both values.

[1]This situation is specific to finite fields. Over $\mathbb{Q}$, for example, every elliptic curve has infinitely many quadratic twists that are not isomorphic over $\mathbb{Q}$ (of course they are all isomorphic over $\overline{\mathbb{Q}}$).

7.7 Mestre's Theorem

As noted above, it is not necessarily the case that the exponent of $E(\mathbb{F}_p)$ has a unique multiple in the Hasse interval. But if we also consider the quadratic twist $\widetilde{E}(\mathbb{F}_p)$, then a theorem of Mestre (published by Schoof in [4]) ensures that for all primes $p > 229$, either $\lambda(E(\mathbb{F}_p))$ or $\lambda(\widetilde{E}(\mathbb{F}_p))$ has a unique multiple in the Hasse interval $\mathcal{H}(p)$. A generalization of this theorem that works for arbitrary prime powers q can be found in [2], but we will restrict ourselves to the case of primes $p > 229$ for the sake of simplicity.

Theorem 7.7 (Mestre)**.** *Let $p > 229$ be prime, and let $E/\mathbb{F}_p$ be an elliptic curve with quadratic twist $\widetilde{E}/\mathbb{F}_p$. At least one of the integers $\lambda(E(\mathbb{F}_p))$ and $\lambda(\widetilde{E}(\mathbb{F}_p))$ has a unique multiple in the Hasse interval $\mathcal{H}(p) = [(\sqrt{p}-1)^2, (\sqrt{p}+1)^2]$.*

Proof. Let $E(\mathbb{F}_p) \simeq \mathbb{Z}/n\mathbb{Z} \oplus \mathbb{Z}/N\mathbb{Z}$ and $\widetilde{E}(\mathbb{F}_p) \simeq \mathbb{Z}/m\mathbb{Z} \oplus \mathbb{Z}/M\mathbb{Z}$, where $n|N$ and $m|M$. Let t be the trace of the Frobenius endomorphism π of E. We have $E[n] = E(\mathbb{F}_p)[n]$, so π fixes $E[n]$ and the matrix π_n corresponding to the restriction of π to $E[n]$ is the identity matrix. The matrix π_{n^2} then has the form

$$\pi_{n^2} = \begin{bmatrix} 1+an & bn \\ cn & 1+dn \end{bmatrix},$$

for some $a, b, c, d \in \mathbb{Z}/n\mathbb{Z}$, and we have

$$p \equiv \det \pi_{n^2} \equiv 1 + (a+d)n \bmod n^2,$$
$$t \equiv \operatorname{tr} \pi_{n^2} \equiv 2 + (a+d)n \bmod n^2.$$

It follows that $4p - t^2 \equiv 0 \bmod n^2$. The trace of Frobenius for $\widetilde{E}$ is $-t$, and we similarly obtain $4p - t^2 \equiv 0 \bmod m^2$. Thus $\operatorname{lcm}(m^2, n^2)$ divides $4p - t^2$. We also have $t = un + 2$ and $t = vm - 2$, for some integers u and v, and subtracting these equations yields $vm - un = 4$. This implies $\gcd(m, n) \le 4$, and therefore $\gcd(m^2, n^2) \le 16$. Thus

$$\frac{m^2 n^2}{16} \le \operatorname{lcm}(m^2, n^2) \le 4p - t^2 \le 4p. \tag{2}$$

Suppose for the sake of contradiction that $N = \lambda(E(\mathbb{F}_p))$ and $M = \lambda(\widetilde{E}(\mathbb{F}_p))$ both have more than one multiple in $\mathcal{H}(p)$. Then M and N are both less than $4\sqrt{p}$ and $MN < 16p$. Since mM and nN lie in $\mathcal{H}(p)$, both are greater than $(\sqrt{p}-1)^2$, and $mnMN > (\sqrt{p}-1)^4$. It follows that $mn > (\sqrt{p}-1)^4/(16p)$. Dividing by 4 and squaring both sides yields

$$\frac{m^2 n^2}{16} > \frac{(\sqrt{p}-1)^8}{4096p^2}. \tag{3}$$

Combining (2) and (3), we have

$$16384p^3 > (\sqrt{p}-1)^8. \tag{4}$$

This implies that if neither M nor N has a unique multiple in $\mathcal{H}(p)$, then $p < 17413$. An exhaustive computer search for $p < 17413$ finds that in fact we must have $p \le 229$. □

7.8 Computing the group order with Mestre's Theorem

We now give a complete algorithm to compute $\#E(\mathbb{F}_p)$ using Mestre's theorem, assuming that p is a prime greater than 229 (if p is smaller than this we can easily count points using one of our naïve algorithms); see [2] for an analogous algorithm that works for all prime powers $q > 49$. As usual, $\mathcal{H}(p) := [(\sqrt{p}-1)^2, (\sqrt{p}+1)^2]$ denotes the Hasse interval.

Algorithm 7.8. Given $E/\mathbb{F}_p$ with $p > 229$ prime, compute $\#E(\mathbb{F}_p)$ as follows:

1. Compute a quadratic twist $\widetilde{E}$ of E using a randomly chosen non-square $s \in \mathbb{F}_p$.
2. Let $E_0 = E$ and $E_1 = \widetilde{E}$, let $N_0 = N_1 = 1$, and let $i = 0$.
3. While neither N_0 nor N_1 has a unique multiple in $\mathcal{H}(p)$:
   a. Generate a random point $P \in E_i(\mathbb{F}_p)$.
   b. Find an integer $M \in \mathcal{H}(p)$ such that $MP = 0$.
   c. Factor M and compute $|P|$ via Algorithm 7.4.
   d. Replace N_i by $\operatorname{lcm}(N_i, |P|)$ and replace i by $1 - i$.
4. If N_0 has a unique multiple M_0 in $\mathcal{H}(p)$ return M_0, otherwise return $2p + 2 - M_1$, where M_1 is the unique multiple of N_1 in $\mathcal{H}(p)$ guaranteed by Mestre's theorem.

It is clear that the output of the algorithm is correct, and it follows from Theorems 7.6 and 7.7 that the expected number of iterations of step 3 is $O(1)$. We thus have a *Las Vegas* algorithm to compute $\#E(\mathbb{F}_p)$. Its running time is dominated by the time to find M in step 3b, and we obtain a total expected running time of $O(\sqrt{p}\ \mathsf{M}(\log p))$, or

$$O\big(\exp(n/2)\,\mathsf{M}(n)\big).$$

We now show how this complexity can be improved using the *baby-steps giant-steps* method to find a suitable M in step 3b.

7.9 Baby-steps giant-steps

The baby-steps giant-steps method is a generic group algorithm that was first introduced by Daniel Shanks in [5] and subsequently generalized by many authors. In the context of searching $\mathcal{H}(q)$ for an integer M such that $MP = 0$, it works as follows.

Algorithm 7.9. Given $P \in E(\mathbb{F}_q)$ compute $M \in \mathcal{H}(q)$ such that $MP = 0$:

1. Pick integers r and s such that $rs \geq 4\sqrt{q}$ and let $a := \lceil (\sqrt{q}-1)^2 \rceil = \min(\mathcal{H}(q) \cap \mathbb{Z})$.
2. Compute the set $S_{\text{baby}} := \{0, P, 2P, \ldots, (r-1)P\}$ of *baby steps*.
3. Compute the set $S_{\text{giant}} := \{aP, (a+r)P, (a+2r)P, \ldots, (a+(s-1)r)P\}$ of *giant steps*.
4. For each giant step $P_{\text{giant}} = (a+ir)P \in S_{\text{giant}}$, check whether $P_{\text{giant}} + P_{\text{baby}} = 0$ for some baby step $P_{\text{baby}} = jP \in S_{\text{baby}}$. If so, output $M = a + ri + j$.

Note that *every* integer in $\mathcal{H}(q)$ can be written in the form $a + ir + j$ with $0 \leq i < s$ and $0 \leq j < r$, and for at least one such M we must have

$$MP = (a+ri)P + jP = P_{\text{giant}} + P_{\text{baby}} = 0$$

for some $P_{\text{giant}} \in S_{\text{giant}}$ and $P_{\text{baby}} \in S_{\text{baby}}$; this shows that the algorithm is correct.

To implement this algorithm efficiently, we typically store the baby steps S_{baby} in a lookup table (such as a hash table or binary tree) and as each giant step P_{giant} is computed, we look up $-P_{\text{giant}}$ in this table. Alternatively, one may compute the sets S_{baby} and S_{giant} in their entirety, sort both sets, and then efficiently search for a match. In both cases, we need the points in S_{baby} and S_{giant} to be uniquely represented.

If we are using projective coordinates this means we must convert each point to affine form: the point $(x : y : z)$ is put in the form $(x/z : y/z : 1)$ by computing the inverse of z in $\mathbb{F}_q$. Done naïvely, this requires $r + s$ field inversions, which costs $O((r + s)\,\mathsf{M}(n) \log n)$, but by using the method described in the next section, it is possible to perform $r + s$ field inversions in $O((r + s)\,\mathsf{M}(n))$ time. Assuming this is done, if we choose $r \approx s \approx 2q^{1/4}$, then the running time of the algorithm above is $O(q^{1/4}\,\mathsf{M}(\log q))$.

Using the baby-steps giant-steps method to implement step 3b of Algorithm 7.8 thus allows us to compute $\#E(\mathbb{F}_q)$ in expected time

$$O\big(\exp(n/4)\,\mathsf{M}(n)\big).$$

7.10 Batching field inversions

Suppose we are given a list of elements $\alpha_1, \ldots, \alpha_m \in \mathbb{F}_q$ whose inverses we wish to compute. The following algorithm accomplishes this using just one field inversion.

Algorithm 7.10. Given $\alpha_1, \ldots, \alpha_m \in \mathbb{F}_q$ compute $\alpha_1^{-1}, \ldots, \alpha_m^{-1}$ as follows:

1. Set $\beta_0 = 1$ and $\beta_i = \beta_{i-1}\alpha_i$ for i from 1 to m. $[\beta_i = (\alpha_1 \cdots \alpha_i)]$
2. Compute $\gamma_m = \beta_m^{-1}$. $[\gamma_m = (\alpha_1 \cdots \alpha_m)^{-1}]$
3. For i from m down to 1:
   a. Compute $\alpha_i^{-1} = \beta_{i-1}\gamma_i$. $[\alpha_i^{-1} = (\alpha_1 \cdots \alpha_{i-1})(\alpha_1 \cdots \alpha_i)^{-1}]$
   b. Compute $\gamma_{i-1} = \gamma_i\alpha_i$. $[\gamma_{i-1} = (\alpha_1 \cdots \alpha_{i-1})^{-1}]$

The algorithm performs less than $3m$ multiplications in $\mathbb{F}_q$ and just one inversion in $\mathbb{F}_q$. Provided that $m = \Omega(\log n)$, its running time is $O(m\,\mathsf{M}(n))$.

In the context of Algorithm 7.9, if we are using a table of baby steps, we can compute all of the baby steps using projective coordinates, convert them to affine form using just one field inversion, and then construct the lookup table. For the giant steps we work in batches of size $m > \log n$, converting an entire batch to affine form using one field inversion and then performing table lookups.

7.11 Optimizations

There are a wide range of optimizations to the baby-steps giant-steps method that have been developed over the years. Here we mention just a few.

1. **Optimize expected time**: If we suppose that M is uniformly distributed over an interval of width N, then we should use $r \approx \sqrt{N/2}$ baby steps so that the average number of giant steps is $s/2 \approx \sqrt{2N}/2 = \sqrt{N/2}$.

2. **Optimize for known distribution**: In the case of elliptic curves we know that M is *not* uniformly distributed – it has a semi-circular distribution.[2] This means we should search from the middle outwards by taking our first giant step in the middle of the interval (at $q+1$), and then alternating steps on either side. We should also choose the number of baby steps to optimize the expected time, using the fact that the expected distance between M and the middle of the interval is $\frac{8}{3\pi}\sqrt{q}$.
3. **Fast inverses**: In groups such as $E(\mathbb{F}_q)$ where we can compute inverses very quickly (the inverse of the point (x,y) is just $(x,-y)$), it makes sense to compute $-P_{\text{giant}}$ at the same time we compute P_{giant} and see whether either matches a baby step; equivalently, whether $P_{\text{giant}} \pm P_{\text{baby}} = 0$ holds. This allows us to double the width of the giant steps and use half as many, or (better), reduce both the number of baby steps and giant steps by a factor of $\sqrt{2}$.
4. **Parity**: We can easily determine the parity of $\#E(\mathbb{F}_q)$ by checking whether it has a point of order 2. If the curve equation is $y^2 = f(x) = x^3 + Ax + B$, then $\#E(\mathbb{F}_q)$ has even parity if and only if $f(x)$ has a root in $\mathbb{F}_q$ (recall that points of order 2 have y-coordinate 0), which we can determine using a root-finding algorithm.[3] Once we know the parity of M we can modify Algorithm 8.1 to only use baby steps that correspond to multiples of P with the same parity (so if M is odd we compute baby steps $P, 3P, 5P, \ldots$, adding $2P$ to each previous step), and use giant steps with even parity. We should then reduce the number of baby steps by a factor of $\sqrt{2}$.

References

[1] William D. Banks and Igor E. Shparlinski, *Sato-Tate, cyclicity, and divisibility statistics on average for elliptic curves of small height*, Israel J. Math. **173** (2009), 253–277.

[2] John E. Cremona and Andrew V. Sutherland, *On a theorem of Mestre and Schoof*, Journal de Théorie des Nombres de Bordeaux **22** (2010), 353–358.

[3] Joachim von zur Gathen and Jürgen Gerhard, *Modern computer algebra*, third edition, Cambridge University Press, 2013.

[4] René Schoof, *Counting points on elliptic curves over finite fields*, Journal de Théorie des Nombres de Bordeaux **7** (1995), 219–254.

[5] Daniel Shanks, *Class number, a theory of factorization and genera*, in 1969 Number Theory Institute (Proc. Symp. Pure Math., Vol. XX), Amer. Math. Soc., 1971, 415–440.

[6] Joseph H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics **106**, second edition, Springer 2009.

[7] Lawrence C. Washington, *Elliptic Curves: Number theory and cryptography*, second edition, Chapman and Hall/CRC, 2008.

[2] This follows from results showing that the Sato-Tate conjecture holds "on average"; see [1].
[3] In fact we only need to check whether $\deg \gcd(x^q - x, f(x)) > 0$, so we can do this deterministically.