DRAFT

Enumeration of Rational Points on Elliptic Curves over Finite Fields

Leonard S. Charlap
Raymond Coley
David P. Robbins

Abstract

We describe a variant of Elkies' improvement of the Schoof algorithm for finding the number of rational points on an elliptic curve over a finite field. Elkies' method reduces the cost by doing some preliminary one-time work the results of which can subsequently be used for all curves. We give a modification which has the advantage that it depends on much more elementary considerations. However our modification may require more one-time work than does his method. These algorithms reduce the work for finding the number of points for curves defined over GF(p) to the order of $\log^6 p$ from order $\log^8 p$.

1. Introduction.

Let F be a field of characteristic $\neq 2, 3$ and let $\bar{F}$ its algebraic closure. For the purposes of this paper we may regard an elliptic curve E defined over F as the set of solutions in $\bar{F}$ of an equation of the form

$$y^2 = x^3 + Ax + B, \tag{1.1}$$

where A and B are in F and satisfy $4A^3 + 27B^2 \neq 0$. By convention the curve contains in addition to all the solutions to (1), its unique point at infinity denoted O. A point on the curve is said to be rational if its coordinates lie in F. By convention O is also considered a rational point. The problem that this paper addresses is that of determining the number of rational points on the curve when F is a finite field, For simplicity we treat mainly the cases in which the finite field is a prime field of characteristic $p > 3$. However, many of our results do extend easily to other finite fields. The first algorithm requiring work polynomial in $\log p$ was given by Schoof [Sc]. It requires an amount of work proportional to $\log^8 p$ using classical algorithms for polynomial algebra. Recently Elkies [E] gave a significant improvement. His method requires some preliminary "almost one-time work". The sense in which this work is one-time will be described later. Its cost has not been analyzed completely. Once this work is done, however, his method requires an amount of work proportional to $\log^6 p$ (using classical algorithms) with the constant of proportionality of the same magnitude as in Schoof's algorithm.

The one-time part of Elkies' method requires considerable facility and ingenuity with modular functions. In this paper we describe a more elementary variant of Elkies' method which costs about twice as much as his after the one-time work. The preliminary work for this variant may well be more expensive than in Elkies' method, but our method may be useful since no ingenuity is required to use it.

The one-time work for our modification involves the calculation of some rather large integers. It turns out that it is possible to calculate these integers modulo a set of primes and, by estimating their size, recover them with the help of the Chinese remainder theorem. However another possibility presents itself. When we compute the number of rational points on a curve over GF(p), these universal integers are only used as elements of GF(p). Thus if we only want to compute numbers of rational points on curves defined over one or

just a few prime fields, then we have the option of doing a less expensive kind of one-time work which is one-time for each p. For some p it turns out that this one-time work for each p is not very much more than the rest of the work that we do on a single curve over GF(p), so that we can nearly afford to do the one-time work every time. However, if we wish to run our algorithm with many curves over the same field, we benefit by not repeating the one-time part of the work for the first curve. Also if we wish to run our algorithm with many fields we can combine the one-time portions with the Chinese Remainder Theorem to produce the universal integers which will work for all fields thus making negligible any future work on these parts of the algorithm.

A combination of these two types of one-time work may be best. See section 8 for a more detailed discussion.

2. Background on elliptic curves and Schoof's algorithm

In this section we give background material on elliptic curves and Schoof's algorithm that will be used below. We will only state the relevant facts. Except where noted a relatively elementary exposition, including proofs, of all these facts is given in [CR].

The best methods for enumerating the rational points on an elliptic curve E over a finite field depend on the abelian group structure of E which is conventionally written additively. This is defined as follows. First the identity element of the group is O. For ordinary points (x, y) one defines $-(x, y) = (x, -y)$. For two points (x_1, y_1) and (x_2, y_2) which are not negatives we define $(x_1, y_1) + (x_2, y_2) = (x_3, y_3)$ where

$$x_3 = -x_1 - x_2 + \lambda^2$$

$$y_3 = -y_1 + \lambda(x_1 - x_3)$$

and λ is given by

$$\lambda = \frac{y_1 - y_2}{x_1 - x_2} \quad \text{if } x_1 \neq x_2,$$

and

$$\lambda = \frac{3x_1^2 + A}{2y_1} \quad \text{if } x_1 = x_2 \text{ and } y_1 \neq 0.$$

In geometric terms the definition of addition says that a straight line always meets the curve in points which add up to O if we take into account multiplicities at the points of intersection of the line and the curve.

Given an integer $m > 0$ the set $E[m]$ of points Q of E with $m \cdot Q = O$ are known as m-torsion points. These points comprise a subgroup of E. It is known that this subgroup is isomorphic to the direct sum of two cyclic groups of order m when the characteristic of F is either 0 or prime to m.

A rational function on E is a function which can be expressed as a rational function of the coordinates x and y. More formally the rational functions can be regarded as the quotient field of the polynomial ring $F[X, Y]$ modulo the ideal generated by $Y^2 - X^3 - AX - B$. There is a sensible definition of the order of zeroes and poles of rational functions.

The "division polynomials" for E are the rational functions ψ_k defined by

$$\begin{aligned}
\psi_1 &= 1 \\
\psi_2 &= 2y, \\
\psi_3 &= 3x^4 + 6Ax^2 + 12Bx - A^2, \\
\psi_4 &= 4y(x^6 + 5Ax^4 + 20Bx^3 - 5A^2x^2 - 4ABx - 8B^2 - A^3), \\
\psi_{2n} &= \psi_n(\psi_{n+2}\psi_{n-1}^2 - \psi_{n-2}\psi_{n+1}^2)/2y \quad \text{if } n \geq 3, \\
\psi_{2n+1} &= \psi_{n+2}\psi_n^3 - \psi_{n-1}\psi_{n+1}^3 \quad \text{if } n \geq 2.
\end{aligned} \tag{2.1}$$

For m prime to the characteristic of F, the function ψ_m has simple zeroes exactly at the set of m-torsion points $\neq O$. One may easily verify that if m is odd, then ψ_m is a polynomial in x, A and B, that its degree in x is $(m^2-1)/2$, and that its leading coefficient is m. Thus it has a simple zero at the x-coordinates of all the m-torsion points. If $m = 2n$ is even, then ψ_m is y times a polynomial in x, A and B.

Suppose that (x, y) is a point of E and that n is a positive integer. Then $n \cdot (x, y) = (g_n(x, y), h_n(x, y))$ for suitable rational functions g_n and h_n. These are given explicitly by

$$\begin{aligned}
g_n &= x - \frac{\psi_{n-1}\psi_{n+1}}{\psi_n^2}, \\
h_n &= \frac{\psi_{n+2}\psi_{n-1}^2 - \psi_{n-2}\psi_{n+1}^2}{4y\psi_n^3}.
\end{aligned} \tag{2.2}$$

We shall make use later of the functions g_n. It is easily verified that g_n is a rational function $g_n(x)$ of x alone. It follows that $g_m(g_n(x)) = g_{mn}(x)$ for all $m, n \geq 1$.

If F has q elements, then the mapping ϕ sending (x, y) to (x^q, y^q) is a mapping from E to E which fixes precisely the set of the rational points. The mapping ϕ is an automorphism of the group structure of E.

The main theoretical tool which allows the computation of the number of rational points on a curve is the following theorem.

Theorem 1: Let E be an elliptic curve defined over GF(q). Denote by $|E|$ the number of rational points on E. Then there exists an integer t such that

1. $\phi^2 - t\phi + q = 0$ (in the endomorphism ring of E);
2. $|E| = 1 - t + q$;
3. $|t| \leq 2\sqrt{q}$.

Also ϕ has trace t and determinant q as a linear operator on the 2-dimensional GF(l)-space $E[l]$.

Schoof's algorithm applies Theorem 1 to curves over GF(p) as follows. For each of a suitable set of small primes l we choose a point P of E with order l. Then we can find, by a simple search, an integer t_1 in the interval from 0 to $l-1$ which satisfies

$$\phi^2(P) - t_1 \cdot \phi(P) + p \cdot P = O.$$

We may then conclude, from part 1 of Theorem 1, that $t \equiv t_1 \bmod l$. We do this for enough small primes l to recover $|E|$ with the help of parts 2 and 3 of Theorem 1 and the Chinese remainder theorem.

We carry this out as follows. We find a polynomial $f(x)$ one of whose roots is the x-coordinate of an l-torsion point. Then there is an l-torsion point $P = (x, y)$ which satisfies $f(x) = 0$ and $y^2 = x^3 + Ax + B$. Using the two preceding algebraic relations we can calculate $Q = \phi^2(P) + p \cdot P$ and then search for a t_1 in the range $0, \ldots, l-1$ such that $t_1 \cdot \phi(P) = Q$.

It will suffice here to note that the work to carry out this computation for a prime l is dominated by the work required to compute x^{p^2} and y^{p^2} mod $f(x)$. This work is somewhat more than $4 \log_2 p$ polynomial multiplications modulo $f(x)$ which, with classical algorithms, requires $8d^2 \log_2 p$ multiplications in F if the degree of $f(x)$ is d. (The obvious way of multiplying polynomials modulo a degree d polynomial requires d^2 field operations for the multiplication and d^2 for the division. For large primes p the cost of multiplication mod p is much higher than the cost of addition, but the number of additions is generally about the same as the number of multiplications. We can therefore usually ignore the costs of additions.) Note that the lower the degree of the polynomial $f(x)$, the faster the algorithm will be. In his original algorithm Schoof suggested the use of the division polynomials described above which have degree $(l^2-1)/2$ which made the work to process a prime l proportional to $l^4 \log_2 p$ field operations in GF(p).

It is not hard to see that for those values l for which $t^2 - 4p$ is a quadratic residue mod l, the Frobenius automorphism ϕ would have an invariant 1-dimensional subspace in $E[l]$. The set of the non-O x-coordinates of the points in such a subspace would be invariant under ϕ, and therefore the degree $(l-1)/2$ polynomial $f(x)$, with these x-coordinates as roots, would have coefficients in F. However no method was known for finding $f(x)$ with little enough work to make this idea useful. The essence of Elkies' improvement is that he shows how to obtain these polynomials, for these l's, for the additional cost of some one-time work, depending only on l, but independent of the choice of elliptic curve, and the cost, for each curve, of finding a root in F of a polynomial of degree $l+1$.

We can outline the cost of these algorithms. In general for Schoof's algorithm we need to process enough primes l to determine the order of the group with the Chinese Remainder Theorem. To have enough primes their product has to be comparable to $\sqrt{p}$ or, equivalently, the sum of their natural logarithms has to be comparable to $\frac{1}{2} \log_e p$. We may regard the prime number theorem as stating roughly that a random integer l is prime with probability $1/\log_e l$. On the other hand, when l is prime, it contributes $\log_e l$ to the sum of logarithms. Thus a random integer l contributes about 1 to the sum of logarithms. It follows that the largest prime that we need to consider with Schoof's algorithm should be about $\frac{1}{2} \log_e p$. With the new algorithms we need to go about twice as high since only about half the primes satisfy the quadratic residue condition. To process a prime l with Elkies' method (or our variant) we need first to form a certain degree $l+1$ polynomial and then to find its roots if any in GF(p). With the help of the one-time work, the formation of the degree $l+1$ polynomial requires negligible time.

We can now explain the precise nature of the one-time work. Strictly speaking this is one-time for each l. For example, in our variant, for each odd prime l, there is a

single polynomial $U_l(T, A, B)$ in $\mathbf{Z}[T, A, B]$ with the property that, when the coefficients are reduced mod p, and the values of the coefficients A and B of the given curve are substituted for the indeterminates A and B, then the resulting polynomial is the one that we need to construct. However the larger the value of p, the more l's we will have to process so the work is not strictly one-time since it does increase to ∞ with p. Nevertheless the U_l for the primes l up to say l_* will suffice to enumerate points on most elliptic curves with up to about e^{l_*} rational points. This amounts to a negligible amount of work per curve. We shall also see that we have the option of computing U_l mod p instead. With this method the work is one-time for each combination of l and p.

The cost of finding the roots in GF(p) of the polynomial U_l is about $2\log_2 p$ polynomial multiplications modulo this U_l or about $2l^2 \log_2 p$ field operations. If U_l has no roots in GF(p), we may proceed to the next prime. If U_l has roots in GF(p), then we use these roots to construct a degree $(l-1)/2$ factor of the l-th division polynomial. Again the one-time work makes negligible the work required to construct this factor. Once we find the factor, the work required for computing t mod l is about $2\log_2 p$ polynomial multiplications modulo this factor. Thus the work for using the factor of the division polynomial is about half the work required to construct it or about $l^2 \log_2 p$ field operations in F.

Our main point here is that the primes that we process will generally be somewhat smaller than $\log_2 p$ and that the work required for processing each l is about $l^2 \log_2 p$ field operations.

The figure of $\log^6 p$ from the introduction is derived from the preceding by counting $\log^2 p$ binary operations per field operation and by using the very conservative approximation that we need to process $\log p$ primes l all about $\log p$ in size.

3. A method which works in principle.

In this section we give an argument based on the Galois theory of the division polynomials to show why a method such as Elkies gives must exist.

Let $F = \mathbf{Q}(A, B)$ where $\mathbf{Q}$ is the field of rational numbers and A and B are indeterminates, and let E be the elliptic curve defined by

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

This is the "generic" elliptic curve. We shall see that certain formulas which we can prove for this curve will hold for any elliptic curve over any field.

Let l be an odd prime. The division polynomial ψ_l, regarded as a polynomial in x, A and B is irreducible over $\mathbf{Q}$. We will include a proof in the Appendix. Let Q be an l-torsion point of E and, for i prime to l, let x_i be the x-coordinate of the point $i \cdot Q$. We have $x_i = x_{-i}$ and $x_i = x_{l+i}$. Since ψ_l is irreducible, the field $K = \mathbf{Q}(A, B, x_1)$, generated over $\mathbf{Q}$ by A, B and x_1, is an algebraic extension of F of degree $(l^2-1)/2$. For each i prime to l, we have $x_i = g_i(x_1)$. It follows that the x_i are elements of K and roots of $\psi_l(x) = 0$. Thus, for each i prime to l, there is an automorphism ρ_i of K which sends x_1 to x_i. We then have

$$(\rho_i \rho_j)(x_1) = g_i(g_j(x_1)) = g_{ij}(x_1) = \rho_{ij}(x_1).$$

Thus $\rho_i \rho_j = \rho_{ij}$ so that $i \to \rho_i$ is a homomorphism from the multiplicative group of GF(l) to the automorphism group of K. The kernel of this mapping is the set ± 1. The group

of automorphisms generated by the ρ's therefore has $(l-1)/2$ elements. Let L denote the fixed field of this group of automorphisms. This will be a subfield of K of index $(l-1)/2$ and hence an extension of F of degree $l+1$.

It is easy to construct elements of the field L. We introduce the abbreviation $d = (l-1)/2$. Any polynomial symmetric in $x_1, \ldots, x_d$ with coefficients in F will be in L, since the x's are permuted by the ρ's. We shall be particularly concerned with the power sums p_k of the x's defined by $p_0 = d$ and, for $k \geq 1$ by

$$p_k = x_1^k + \cdots + x_d^k.$$

We shall show in the Appendix that p_1 generates L over F. It follows that all other symmetric functions of the $x_1, \ldots, x_d$ can be expressed as rational functions of p_1 and A and B.

Now here is a plan for producing factors of the division polynomials. We first find the monic minimum polynomial U_l satisfied by p_1 over F. Next for each $k = 2, \ldots, d$, we express p_k as a rational function of A, B, p_1 and possibly also $p_2, \ldots, p_{k-1}$. (We use power sums here rather than elementary symmetric functions because the computations are somewhat simpler. We can convert between power sums and elementary symmetric functions with the help of Newton's formulas (3.1) below.)

Each of these formulas which we derive over $\mathbf{Q}$ can be written as rational functions with integer coefficients. Thus they make sense over any field. We will be able to show later that the ones in which we are particularly interested are valid over any field.

In principle the finding of the polynomial U_l and these rational expressions could be done as one-time work. In fact, we will not carry this plan out in every detail. In particular it turns out that our best methods for finding some of the rational relations are so inefficient that we are better off with an alternate plan which we describe in section 7. Others are so efficient that we can afford to do them every time as described in section 5.

Now suppose that F is GF(p) and that we are given a specific curve and therefore a specific A and B in this field. When $t^2 - 4p$ is a quadratic residue mod l, we know from Theorem 1 that the Frobenius operator will leave invariant at least one 1-dimensional GF(l)-subspace of $E[l]$. Let $x_1, \ldots, x_d$ be the x-coordinates of the non-O points in this subgroup. Then the symmetric functions of these x's will be in F. In particular their sum will be a root in F of $U_l = 0$ regarded as a polynomial over F. We can use standard methods to determine whether such a root exists and what the roots are when they do exist. When there are no roots in F, we proceed to the next l. If we find a (non-repeated) root in F, this can only arise from an eigenspace V of $E[l]$ for the operator ϕ. Indeed, suppose that a one-dimensional subspace V of $E[l]$, not invariant under the action of ϕ, happened to have a sum of x-coordinates that was in F. Then the space $\phi(V)$ would have the same sum of x-coordinates so this root in F would be repeated. Returning to the non-repeated root arising from the eigenspace V, we now form the other power sums of x-coordinates from this subspace using our known rational expressions. We then convert the power sums to elementary symmetric functions $s_0, \ldots, s_d$ of $x_1, \ldots, x_d$ using $s_0 = 1$ and Newton's formula

$$s_k = \frac{1}{k} \sum_{j=1}^{k} (-1)^{j-1} p_j s_{k-j} \qquad k \geq 1. \tag{3.1}$$

These elementary symmetric functions are the coefficients of a factor of degree d with coefficients in F of the division polynomial $\psi_l(x)$. This can be used in place of $\psi_l(x)$ in Schoof's algorithm with a considerable reduction in running time because of its lower degree.

This explains how the method works in principle. What problems might arise? One conceivable problem is that, once we having solved for p_1, the rational expressions for the other symmetric functions might evaluate to 0/0. In this case our method would fail. At first glance this would appear to be an unlikely occurrence. The main problem is that the universal polynomials U_l and the universal rational relations might be prohibitively hard to calculate or to store if they can be calculated.

Elkies has given two given two essentially separate contributions to this problem. The easier to understand is his method for expressing $p_4, \ldots, p_d$ in terms of p_1, p_2, p_3. Here we will simply give an exposition of his method. The problem of finding p_1, p_2 and p_3 is more complicated. Here we describe our own method. The idea behind both methods is to find the desired relations relations that hold for elliptic curves over the complex numbers. When these relations are uniquely determined, they will be the same as the ones described above over $\mathbf{Q}(A, B)$.

4. Elliptic Curves over the complex numbers.

The facts about complex elliptic curves which we use in this section are standard, appearing in a great many books. Thus we give no proofs. All the proofs are, for example, given in the first few pages of [A].

For elliptic curves over the complex numbers the exposition is easier if we begin with a description of the group structure and find the equation of the curve later. To define an elliptic curve over the complex numbers we start with two complex numbers ω_1 and ω_2 with ω_2/ω_1 having positive imaginary part. Let ω stand for a variable element of the lattice L of points $m\omega_1 + n\omega_2$, where m and n vary over the integers. Then the quotient group $E = \mathbf{C}/L$ is an elliptic curve over $\mathbf{C}$. This defines the group structure. To identify this group with the groups we have defined above we need to define the coordinate functions x and y on the curve which satisfy the equation (1).

The standard way to do this is with the Weierstrass function $\wp(z)$ defined by

$$\wp(z) = \frac{1}{z^2} + \sum_{\omega \neq 0} \left[\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} \right], \tag{4.1}$$

which is invariant under translations by elements of L and so defines a function on E. Define $x(z) = \wp(z)$ and $y(z) = \frac{1}{2}\wp'(z)$. Using (4.1) we can expand $x(z)$ in powers of z

$$x(z) = \frac{1}{z^2} + \sum_{k=1}^{\infty} c_k z^{2k}$$

with

$$c_i = (2i+1) \sum_{\omega \neq 0} \frac{1}{\omega^{2i+2}} \tag{4.2}$$

One may then prove that

$$y^2 = x^3 + Ax + B \tag{4.3}$$

where

$$A = -5c_1 \quad \text{and} \quad B = -7c_2. \tag{4.4}$$

(One checks that with these A and B the difference of the two sides of (4.3) is doubly periodic with no poles and maps 0 to 0. This implies that the difference function is itself zero.) Thus these curves that arise from lattices are examples of the type we defined at the outset. It can also be shown, conversely, that, provided $4A^3 + 27B^2 \neq 0$, there exists a lattice which gives rise to a curve with these values of A and B.

Differentiating (4.3) with respect to z, we find that

$$\frac{d^2x}{dz^2} = 6x^2 + 2A. \tag{4.5}$$

By comparing power series expansions of the two sides of (4.5) we can calculate $c_3, c_4, \ldots$ in terms of c_1 and c_2. In fact we find that, for $k \geq 3$, we have

$$c_k = \frac{3}{(k-2)(2k+3)} \sum_{h=1}^{k-2} c_h c_{k-1-h}. \tag{4.6}$$

We can also continue differentiating (4.3) with respect to z, finding

$$\frac{d^3x}{dz^3} = 24xy$$

and

$$\frac{d^4x}{dz^4} = 120x^3 + 72Ax + 48B, \quad \text{etc.}$$

We find that all even derivatives of x with respect of z are polynomials in x

$$\frac{d^{2k}x}{dz^{2k}} = \mu_k(k+1)x^{k+1} + \cdots + \mu_k(0). \tag{4.7}$$

with $\mu_k(k+1) = (2k+1)!$. If A and B are known, it is a simple matter to compute the $\mu_k(j)$ for small k and j.

5. Finding $p_4, p_5, \ldots$, in terms of p_1, p_2, p_3.

We begin with a method which works for elliptic curves over the complex numbers. Let E be the curve associated to the lattice spanned by ω_1 and ω_2. Let l be an odd prime and $d = (l-1)/2$. Let P be the l-torsion point ω_1/l, and let $x_j = x(j \cdot P)$ for j prime to l, where x is the Weierstrass $\wp$-function described in the preceding section. Again let

$$p_k = x_1^k + \cdots + x_d^k$$

be the k-th power sum of $x_1, \ldots, x_d$. We assume that we know p_1, p_2 and p_3 and show how to compute the other p_k's.

Let E' be the elliptic curve whose lattice is spanned by ω_1/l and ω_2 and let x' and y' be the associated coordinate functions. This is the "isogenous" curve referred to in the title of [E]. Then we have

$$x'(z) = x(z) + \sum_{j=1}^{l-1} [x(z + j \cdot P) - x(j \cdot P)]$$

since the function on the right side has periods ω_1/l and ω_2, has double poles exactly at the points of the lattice L', and has constant term 0 in its power series in powers of z. The corresponding y-coordinate is given by

$$y' = \frac{1}{2}\frac{dx'}{dz} = \sum_{j=0}^{l-1} y(z + j \cdot P).$$

Suppose that the equation of this curve is

$$(y')^2 = (x')^3 + A'x' + B'.$$

We may expand x' in powers of z:

$$x' = \frac{1}{z^2} + c_1' z^2 + c_2' z^4 + \cdots.$$

The function

$$f(z) = x'(z) - x(z) = \sum_{j=1}^{l-1} [x(z + j \cdot P) - x(j \cdot P)]$$

is analytic at $z = 0$. We also have from (4.7)

$$\left.\frac{d^{2k} f}{dz^{2k}}\right|_{z=0} = \sum_{h=1}^{l-1} \sum_{j=0}^{k+1} \mu_k(j) x^j (h \cdot P) = 2 \sum_{j=0}^{k+1} \mu_k(j) p_j.$$

On the other hand the derivative is $(2k)!$ times the coefficient of z^{2k} in the power series expansions of $f = x' - x$. This leads to the identity

$$(2k)!(c_k' - c_k) = 2[\mu_k(0)p_0 + \cdots + \mu_k(k+1)p_{k+1}] \quad \text{for } k \geq 1. \tag{5.1}$$

Since we know A and B, we can compute all the c's and μ's in the preceding equation. The instances $k = 1$ and $k = 2$ of (5.1) combined with (4.3) yield

$$A' - A = 5(p_2 + Ap_0) \tag{5.2}$$

and

$$B' - B = 7(5p_3 + 3Ap_1 + 2Bp_0) \tag{5.3}$$

from which we can determine A' and B' since we know A, B, p_0, p_1, p_2 and p_3. Next we use the formulas (4.6) again with the curve E' to compute all the c'_k. Finally if, inductively, we already know $p_1, \ldots, p_k$ we can then use (5.1) to find p_{k+1}. Thus all subsequent p_{k+1}'s can be found.

This entire method for calculating our factor of the division polynomial from p_1 p_2 and p_3 makes sense over any field, and is, in fact, valid over any field, as the discussion below will show. It is also quite inexpensive since all the steps require on the order of l^2 field operations. Thus this part of the work could and should be done every time rather than as one-time work.

Now we discuss the validity of this method for elliptic curves over other fields. We can eliminate A' and B' from the relations that we just derived for curves over the complex numbers and replace them by equivalent polynomial relations in $p_1, \ldots, p_k$ for $k > 3$ with coefficients in $\mathbf{Z}[A, B]$ in which the only appearance of p_k is the term $(2k+1)!p_k$. We wish to show that these relations make sense and are valid for any field.

Thus we let E be an elliptic curve over any field and let x_1 be the x-coordinate of an l-torsion point P. Also for j prime to l let x_j be the x-coordinate of $j \cdot P$. Recall that from (2.2)

$$x_j = x_1 - \frac{\psi_{j-1}(x_1)\psi_{j+1}(x_1)}{\psi_j(x_1)^2}.$$

Then the power sums p_k of $x_1, \ldots, x_d$ are given by

$$p_k = \sum_{j=1}^{d} \left[x_1 - \frac{\psi_{j-1}(x_1)\psi_{j+1}(x_1)}{\psi_j(x_1)^2} \right]^k. \tag{5.4}$$

If we could show that our relations derived above for p_k's of curves over C were valid for these p_k's, then these same relations can be used to calculate p_k for all k for which $(2k+1)! \neq 0$. Since we actually need only $p_1, \ldots, p_d$, we will be able to calculate all the p_k's that we need provided that $p > l$. (This condition will certainly hold for all cases of interest if we are finding the number of rational points on elliptic curves over GF(p). However it appears to fail when we attempt to compute the number of rational points on an elliptic curve defined over a large finite field with small characteristic.)

Using (5.4) the relations derived can be written as rational relations among the ψ's. If we then multiply through by the product of the denominators, we obtain an equivalent relation which states that a certain polynomial combination $H(x, \psi_1(x), \ldots, \psi_{d+1}(x))$ with coefficients in $\mathbf{Z}[A, B]$ is zero for $x = x_1$, the x-coordinate of an l-torsion point. (The condition is equivalent since the l-torsion points will not be zeroes of the denominators.) Since we know that our relation is valid for the complex numbers, we may conclude that $H(x, \psi_1(x), \ldots, \psi_{d+1}(x))$ is divisible by $\psi_l(x)$ for all complex A and B satisfying $4A^3 + 27B^2 \neq 0$. Since $\psi_l(x)$ has constant leading coefficient l, the quotient, $G = H/\psi_l$, regarded as a polynomial in x, and obtained with the standard division algorithm, will have coefficients which are polynomials in $\mathbf{Z}[A, B, 1/l]$. This implies that the identity $G\psi_l = H$ holds for all complex A and B with $4A^3 + 27B^2 \neq 0$ and consequently holds as a polynomial identity in $\mathbf{Z}[A, B, x, 1/l]$. Thus our identity will hold in any field of characteristic 0. We may now also reduce mod p to obtain an identity with coefficients in GF(p) which will

therefore hold in any field of characteristic p. Since the division polynomials also reduce to division polynomials mod p, we can conclude that $G\psi_l = H$ is valid over any field of characteristic $p \neq l$. Reversing our steps, we see that H, which makes sense over any field of characteristic $\neq l$, will be zero at the x-coordinate of any l-torsion point. Thus our relations will be valid over any field of characteristic $\neq l$.

6. Finding the polynomial satisfied by p_1.

In this section we discuss the calculation of p_1. In section 3 we showed that, over the field $\mathbf{Q}(A, B)$, p_1 will satisfy a certain polynomial U_l of degree $l+1$. Here we show how to compute these polynomials. Then for a given curve we substitute our known values of A and B and solve for p_1. Throughout this section we assume that l is an odd prime and $d = (l-1)/2$.

We have three methods for finding the polynomials U_l. They are presented in order of increasing complexity. However the more complex methods are also computationally more efficient.

Method 1.

The simplest method is to express p_1 as a rational function in the x-coordinate x_1 of some l-torsion point using the case $k = 1$ of (5.4). We also know that x_1 satisfies the l-th division polynomial. Eliminating x_1 from these two expressions using the resultant yields a polynomial satisfied by p_1. However the polynomial obtained is actually $U_l^{(l-1)/2}$ rather than U_l itself. Thus this method seems impractical although it can be carried out by hand if $l = 3$ when there is essentially nothing to do ($U_3 = \psi_3/3$) or with MACSYMA, say, when $l = 5$.

Method 2.

For our second method we begin with curves defined over the complex numbers where the lattice is generated by 1 and τ, a complex number with positive imaginary part. Then the point $1/l$ is an l-torsion point. We take p_1 to be the sum of x-coordinates from the 1-dimensional subspace of $E[l]$ spanned by $1/l$. We can express A and B and p_1 as power series in powers of $q = e^{2\pi i\tau}$. We have

$$\begin{aligned} p_1 &= 4l\pi^2\left(\frac{l-1}{24} + \sum_{n=1}^{\infty}\sigma_1'(n)q^n\right), \\ A &= -\frac{1}{3}\pi^4\left(1 + 240\sum_{n=1}^{\infty}\sigma_3(n)q^n\right), \\ B &= -\frac{2}{27}\pi^6\left(1 - 504\sum_{n=1}^{\infty}\sigma_5(n)q^n\right), \end{aligned} \tag{6.1}$$

where $\sigma_k(n)$ denotes the sum of the kth powers of the divisors of n and $\sigma_1'(n)$ is the sum of the divisors of n that are prime to l. The last two formulas appear in many books. See for example [A]. The first formula is derived in [E] using (6.3) below. We also indicate another derivation at the end of section 6.

Before we proceed farther we need to derive some properties of the polynomials U_l that will help with their computation. First notice that in the defining equation (1.1) for

elliptic curves all the terms have the same degree if we assign to x, y, A, B the "degrees" 2, 3, 4 and 6. We will call this the *generalized degree* of a term. A polynomial in these variables will be called homogeneous if all the terms have the same generalized degree. With this terminology one may verify that the division polynomials ψ_n are homogeneous of generalized degree n^2-1. We may extend this grading by generalized degree to rational functions by assigning to a quotient of two polynomials a generalized degree equal to the difference of the degrees of the numerator and denominator. The rational functions g_n and h_n of (2.2) are homogeneous of generalized degrees 2 and 3 respectively. Because ψ_l is homogeneous, the field K of section 3 has a natural grading which assigns the degrees 2, 4 and 6 to x_1, A and B. Thus in K the $g_k(x_1)$ and therefore also

$$p_1 = g_1(x_1) + g_2(x_1) + \cdots + g_d(x_1)$$

have the degree 2. Also since $g_1(x_1), \ldots, g_d(x_1)$ are all zeroes of $\psi_l(x)$, which has leading coefficient l, lp_1 is integral over $\mathbb{Z}[A,B]$ (satisfies a monic polynomial with coefficients in $\mathbb{Z}[A,B]$.) Thus the monic minimum polynomial $U_l(T)$ satisfied by p_1 can be written

$$U_l(T) = T^{l+1} + \frac{1}{l}C_1(A,B)T^l + \cdots + \frac{1}{l^{l+1}}C_l(A,B)$$

where $C_1(A,B), \ldots, C_l(A,B)$ are polynomials in A and B with integer coefficients. A simple homogeneity argument then shows that $C_j(A,B)$ is a polynomial in A and B homogeneous of generalized degree $2j$. Thus U_l is an integral linear combination of the monomials $T^{i_1}A^{i_2}B^{i_3}$ with $i_1, i_2, i_3 \geq 0$ and $i_1 + 2i_2 + 3i_3 = l+1$. Let $N(l)$ be the number of these monomials; that is, the number of solutions in nonnegative integers i_1, i_2 and i_3 to the equation $i_1 + 2i_2 + 3i_3 = l+1$.

Now let us return to to our complex curve. To find U_l we need only look for a rational linear combination of the monomials $p_1^{i_1}A^{i_2}B^{i_3}$ which is zero. To find these coefficients we compute the first $N(l)-1$ terms of the q-series of each of the monomials. Now we let the coefficients in the linear combination that we are seeking be unknowns. For the correct linear combination the coefficients of its q-series must be 0. If we add the extra condition that the coefficient of the monomial p_1^{l+1} be 1 (which amounts to requiring that U_l be monic), then we obtain a system of $N(l)$ linear equations in $N(l)$ unknowns. Of course this system will always have a solution which arises from the coefficients of U_l. For all the l we have tried the system has been non-singular so that we can use this system to find U_l. If one of these systems were singular, we could probably still determine U_l by using more terms of the q-series which would give extra equations satisfied by the same unknowns.

We can simplify the calculation somewhat as follows. Since all the monomials have the same generalized degree, the polynomial we are seeking is unchanged if we multiply the expressions for p_1, A and B by λ^2, λ^4 and λ^6 for any non-zero constant λ. In practice we choose our "standard scaling" $\lambda^2 = 3/\pi^2$ so that our three basic series above are replaced

by

$$p_1 = \frac{l(l-1)}{2} + 12l \sum_{n=1}^{\infty} \sigma_1'(n) q^n$$

$$A = -3 - 720 \sum_{n=1}^{\infty} \sigma_3(n) q^n = -3E_2 \tag{6.2}$$

$$B = -2 + 1008 \sum_{n=1}^{\infty} \sigma_5(n) q^n = -2E_3$$

which have integer coefficients. In general the standard scaling of an expression of generalized degree $2k$ is obtained by multiplying it by $(3/\pi^2)^k$. We will scale other series for homogeneous functions of A and B with this same standard scaling. We shall see that this will lead to series with integer coefficients for all the quantities of interest. We shall make use later of the "Eisenstein series" E_2 and E_3 defined in (6.2), which have the desirable property that their series have constant term 1.

The coefficients in this system of equations equation are rather large integers so that multiple precision arithmetic is necessary to compute the solutions over the rationals. It is easier to calculate modulo several primes of convenient size and use the Chinese remainder theorem to compute the final answer. To use the Chinese Remainder theorem we first need to know that the coefficients that we are finding are integers. Our discussion of the integrality of p_1 above tells us that the $C_j(A,B)$ is in $\mathbb{Z}[A,B]$. In fact we shall see in our discussion of Method 3 below that if $l \geq 5$, then $C_j'(A,B) = C_j(A,B)/l^j$ is always an integer. We can save some precision by solving for the C''s instead of the C's.

Of course we need a bound on the size of the integers that we are determining in order to be sure that we have the correct answers. We postpone this discussion to section 8.

Next we give some work estimates for solving these systems of equations modulo a prime p. The number $N(l)$ can be interpreted as the number of points in the Cartesian plane with integral coordinates in the triangle with vertices $(0,0)$, $((l+1)/2, 0)$ and $(0, (l+1)/3)$. Thus its value is approximately $l^2/12$. This implies that the work required to solve the equations for the coefficients modulo p is approximately

$$\frac{1}{3}\left(\frac{l^2}{12}\right)^3 = \frac{l^6}{5184}$$

field operations in GF(p). The work required to generate the equations, that is, the work required to compute the truncated power series for the needed monomials in p_1, A and B is about $\frac{1}{2}N(l)^2$ field operations for each polynomial multiply mod $q^{N(l)-1}$. The q-series can be obtained with one polynomial multiply each by computing and saving the monomials involving only A and B first which can serve for all the l's we are interested in. Then the powers of the various p_1's can be multiplied in later. Thus the work for this part is $\frac{1}{2}N(l)^3 \approx l^6/3456$. Since the polynomials that we are multiplying are of high degree, we can use fast multiplication techniques to make this part of the work small compared to the work for solving the equations.

Method 3.

We used Method 2 to calculate the $U_l(T)$ for $l \leq 43$ and observed that there was a fairly clear pattern to the coefficients of the high powers of T, and that, more generally, there was a recognizable pattern for all the coefficients of T when $A = -3$, $B = -2$. These are the values that A and B take on when $q = 0$ in their q-series expansions. This led to the idea of computing the entire q-series expansion of each of the coefficients which in turn led to the method described next.

As above let l be an odd prime and $d = (l-1)/2$. Each of the $l+1$ sums of x-coordinates from each of the $l+1$ possible one-dimensional GF(l)-subspaces of $E[l]$ are roots of U_l. Thus one approach to finding U_l is to construct the polynomial with these sums as roots. In general this observation is of no help since we do not know the sums of x-coordinates. However, in the case of our curves over C, we can expand each of the $l+1$ sums of x-coordinates as q-series and then form the polynomial with these q-series as zeroes. (We shall see that, strictly speaking, the series for these sums are series in powers of $w = q^{1/l}$ but that symmetric functions of these series involve only powers of q.) This will give us the q-series for the coefficients of U_l and we can then solve a system of linear equations to express these q-series as polynomials in A and B.

We again prefer to work with power sums so that we actually express the power sums of the zeroes of U_l first as q-series and then as polynomials in A and B. When we wish to use the results of this one-time work on a particular curve, we use our formulas to obtain the power sums of the zeroes of U_l for the particular curve and then use Newton's formulas to convert these to elementary symmetric functions which give us the coefficients of U_l for that curve.

We now turn to the computation of the q-series for the power sums. Let E be the complex curve associated with the lattice generated by 1 and τ. Then the $l+1$ points $1/l$ and $(k+\tau)/l$, $k = 0, \ldots, l-1$, generate the $l+1$ one-dimensional GF(l)-subspaces of $E[l]$. We need to find the q-series for the sum of x-coordinates from each of these subspaces. In the case of the subspace spanned by $1/l$ this sum has already been given in (6.1) above. Let us calculate the other sums $S_0, \ldots, S_{l-1}$ associated with $(k+\tau)/l$, $k = 0, \ldots, l-1$. Here we make use of the q-series form (which is really the Fourier series expansion) of the Weierstrass function given, for example, in [L, p. 46]:

$$x(z) = 4\pi^2 \left[\frac{-1}{12} + 2 \sum_{n=1}^{\infty} \frac{q^n}{(1-q^n)^2} - \sum_{n=-\infty}^{\infty} \frac{q^n q_z}{(1-q^n q_z)^2} \right], \tag{6.3}$$

where $q = e^{2\pi i \tau}$ and $q_z = e^{2\pi i z}$. Note that the first two terms are independent of z. S_k is given by

$$\frac{1}{2} \sum_{j=1}^{l-1} x((jk + j\tau)/l).$$

When $z = (jk + j\tau)/l$, we have $q_z = \zeta^{kj} w^j$ where $w = e^{2\pi i \tau / l}$ is a "distinguished" l-th root of q and $\zeta = e^{2\pi i / l}$ is a "distinguished" l-th root of unity. The factor $1/2$ is present to adjust the sum to just the first $d = (l-1)/2$ terms as we did for p_1. From here on we

use the scaled S's obtained by multiplying the ones defined above by $3/\pi^2$. It follows that

$$S_k = \frac{(1-l)}{2} + 12(l-1)\sum_{n=1}^{\infty}\frac{q^n}{(1-q^n)^2} - 6\sum_{j=1}^{l-1}\sum_{n=-\infty}^{\infty}\frac{q^n\zeta^{kj}w^j}{(1-q^n\zeta^{kj}w^j)^2}.$$

If we replace q^n by w^{ln}, the last term becomes

$$6\sum_{\gcd(j,l)=1}\frac{\zeta^{kj}w^j}{(1-\zeta^{kj}w^j)^2}.$$

If we note that the expression $t(1-t)^{-2}$ remains unchanged when t is replaced by $1/t$, then we obtain

$$S_k = \frac{(1-l)}{2} + 12l\sum_{n=1}^{\infty}\frac{q^n}{(1-q^n)^2} - 12\sum_{n=1}^{\infty}\frac{\zeta^{kn}w^n}{(1-\zeta^{kn}w^n)^2}.$$

Finally we can expand the two sums in powers of q and w respectively to obtain

$$S_k = l\left[\frac{-1}{2} + 12\sum_{n=1}^{\infty}\sigma_1(n)q^n\right] - \left[\frac{-1}{2} + 12\sum_{n=1}^{\infty}\sigma_1(n)\zeta^{kn}w^n\right].$$

These are the expressions that we were seeking for the other sums. Note that they are really series in powers of w.

We still need to compute the q-series of the power sums of $p_1, S_0, \ldots, S_{l-1}$. We require only the first few terms of these series so we need not compute these series in closed form. Nevertheless we can rewrite the power sums in a more efficient form which makes clear that they are power series in q rather than w. We treat the powers of p_1 and the S's separately and add the results. The powers of p_1 which are needed are easily computed directly. The power sums of the remaining terms $S_0, \ldots, S_{l-1}$ are more difficult. Let

$$\gamma(u) = \frac{-1}{2} + 12\sum_{n=1}^{\infty}\sigma_1(n)u^n$$

and define $c(j,n)$ by

$$\gamma^j(u) = \sum_{n=0}^{\infty}c(j,n)u^n.$$

Then $S_k = l\gamma(q) - \gamma(\zeta^k w)$ so that the r-th power sum of the S's is given by

$$\begin{aligned}
\sum_{k=0}^{l-1} S_k^r &= \sum_{k=0}^{l-1}\sum_{j=0}^{r}(-1)^j\binom{r}{j}l^{r-j}\gamma^{r-j}(q)\gamma^j(\zeta^k w) \\
&= \sum_{j=0}^{r}\left[(-1)^j\binom{r}{j}l^{r-j}\gamma^{r-j}(q)\sum_{k=0}^{l-1}\gamma^j(\zeta^k w)\right] \\
&= \sum_{j=0}^{r}\left[(-1)^j\binom{r}{j}l^{r-j}\gamma^{r-j}(q)\sum_{n=0}^{\infty}c(j,n)\sum_{k=0}^{l-1}\zeta^{nk}w^n\right] \qquad (6.4) \\
&= l\sum_{j=0}^{r}\left[(-1)^j\binom{r}{j}l^{r-j}\gamma^{r-j}(q)\sum_{n=0}^{\infty}c(j,nl)q^n\right] \\
&= l\sum_{j=0}^{r}(-1)^j\binom{r}{j}l^{r-j}\gamma^{r-j}(q)\delta_{jl}(q),
\end{aligned}$$

where $\delta_{jl}(q) = \sum_{n=0}^{\infty} c(j,nl)q^n$. These equations give a relatively efficient way to compute the needed terms of the power sums of the S's.

Now we have expressed the power sums of the zeroes of U_l as q-series. However we need these power sums expressed in terms of A and B. We do this by solving a system of linear equations. The k-th power sum is a symmetric function of the zeroes of U_l of generalized degree $2k$. Thus we can express it as a linear combination of the monomials $A^{j_1}B^{j_2}$ with $j_1, j_2 \geq 0$ and $2j_1 + 3j_2 = k$. Let $M(k)$ be the number of these monomials. Given $M(k)$ terms of the q-series expansion of the k-th power sum and $M(k)$ terms of the q-series expansion of each of these monomials, we can form a system of $M(k)$ equations in $M(k)$ unknowns to find the coefficients in this linear combination. One sees easily that $M(k) \approx k/6$, so these systems are much smaller than those which arose in Method 2. However these systems are simpler in another way which we now describe.

Let

$$\Delta = 12^{-3}(E_2^3 - E_3^2) = -2^{-8}3^{-6}(4A^3 + 27B^2),$$

a quantity of generalized degree 12. The scaled (see the discussion of scaling in the previous section) q-series expansion is known [Se$_7$ p.95], to be

$$\Delta = q\prod_{n=1}^{\infty}(1-q^n)^{24},$$

which has integer coefficients and no constant term. Suppose that $k = 2m$, $m \geq 1$. Then $M(k) = [m/3] + 1$. The polynomials

$$E_2^{m-3j}\Delta^j, \quad j = 0, \ldots, [m/3] \qquad (6.5)$$

are all generalized degree $2k$ polynomials in A and B and they are linearly independent since the term indexed by j in this sequence has a q-series beginning with q^j. (Here we

are essentially following a remark in [Se, p.105].) Now, if we express the k-th power sum as a linear combination of terms of this basis, our system of equations will be triangular with 1's on the diagonal which is easier to solve than a random system. If $k = 2m + 3$ we can use the basis

$$E_3 E_2^{m-3j} \Delta^j, \quad j = 0, \ldots, [m/3] \tag{6.6}$$

instead. This basis is just as easy to use with a particular curve. We simply compute $E_2 = -A/3$, $E_3 = -B/2$ and Δ from the known values of A and B and then use our formulas expressing the power sums of the roots of U_l in terms of E_2, E_3 and Δ. Finally we can compute the coefficients of U_l with the help of Newton's formulas.

We discuss the cost of the one-time work in section 8.

Here are the first few U_l's.

$$\begin{aligned}
U_3 &= x^4 + 2Ax^2 + 4Bx - A^2/3 \\
U_5 &= x^6 + 20Ax^4 + 160Bx^3 - 80A^2x^2 - 128ABx - 80B^2 \\
U_7 &= x^8 + 84Ax^6 + 1512Bx^5 - 1890A^2x^4 \\
&\quad - 9072ABx^3 + 644A^3x^2 - 21168B^2x^2 + 5832A^2Bx - 567A^4
\end{aligned}$$

Remark:

The (scaled) power series for p_1 and the S's have integer coefficients. This implies that the q-series expansions for their elementary symmetric functions will also have integer coefficients. Thus ultimately the coefficients of U_l when expressed as polynomials in E_1, E_2 and Δ as described above have integer coefficients. Substituting the expressions for E_1, E_2 and Δ in terms of A and B, we find that U_l is in $\mathbf{Z}[A, B, 1/2, 1/3]$. But we also know from our integrality argument that U_l is in $\mathbf{Z}[A, B, 1/l]$. This shows that if $l > 3$, then U_l is really in $\mathbf{Z}[A, B]$ as claimed in the discussion of Method 2.

Remark: There are no monomials in A and B of generalized degree 2. Therefore the sum of the zeroes of $\psi_l(x)$, which is a polynomial in A and B of generalized degree 2, is zero. This implies that

$$p_1 + S_0 + S_1 + \cdots + S_{l-1} = 0.$$

Now the preceding expressions for S_k can be used to give another derivation of the formula for p_1.

Remark: It is possible using much the same style argument as we gave in section 5 to show that the minimum polynomials which we find for p_1 which hold over the complex numbers will hold over any field of characteristic > 3 and not equal to l. We omit the details.

7. The polynomials satisfied by A' and B'.

Our plan described in section 3 calls now for expressing p_2 and p_3 as rational functions of p_1, A and B. While this is certainly possible in principle, and would be desirable in practice, we know only one rather ad hoc method for finding these relations (described in section 10) and we are uncertain what the cost of this method is. Thus it is easier to find p_2 and p_3 using the same methods that we used for p_1. However, it turns out to be

somewhat more convenient to compute (in the notation of section 3) A' and B' instead of p_2 and p_3. Knowledge of one pair can be converted to the other with the formulas (5.2) and (5.3).

Note first that, from (5.2) and (5.3), A' and B' are integral over $\mathbf{Q}[A, B]$ since p_2 and p_3 are integral. We shall see that, using the methods of the previous section, we can compute the monic minimum polynomials V_l and W_l of A' and B'. In the cases that U_l has a zero in our base field, we find the zeroes in the base field of V_l and W_l.

We expect that the every-time work for our algorithm is concentrated primarily in solving these polynomial equations. We have to solve one equation of degree $l+1$ for half the l's and three equations for the other half. Elkies' method requires solving one equation of degree $l+1$ for every l. This is why we estimate that our method will require about twice as much work. Our method has some additional complications such as those described immediately below, but these should not cost us much extra work.

When we solve the equations for A' and B', in the case of favorable l, we expect these polynomials to have two roots each in the field F (one for each eigenspace with respect to ϕ). We are then faced with the problem of determining how to associate the zeroes of V_l and W_l with those of U_l. We can choose an arbitrary root of U_l for p_1, but then there remain four possibilities for matching this with values of A' and B'. We know of no procedure better than guessing which of the four is correct. Having made the guess we then construct the putative factor of ψ_l according to the method of section 5. We can then test whether the resulting polynomial is a factor of ψ_l by computing ψ_l and dividing. However, the work involved just in computing ψ_l is on the order of l^4. A simple way to avoid this is to use the recursion (2.1) from scratch but compute the ψ's modulo the putative factor. This way the work should be a fraction of l^3 so we can afford to do it every time. Note also that that this search through the four possible polynomials will only be necessary in the half of the cases that l is favorable. On these occasions we will need to try $2\frac{1}{2}$ possibilities on average.

We now proceed with a discussion of finding the minimum polynomials of A' and B'. We will treat the case of A' carefully and give only the results for B'. All the methods of the previous section are again available but we will discuss only the analog of Method 3 since it is much more efficient.

Let us denote by V_l the monic minimum polynomial for A'. It turns out that the computation of V_l is really simpler than that of U_l. We compute V_l for the complex curve E whose associated lattice is generated by 1 and τ by finding the power sums of its roots as q-series and then solving a triangular system of linear equations to express the power sums as polynomials in E_1, E_2 and Δ. The $l+1$ subgroups of E of order l are generated by the $l+1$ complex numbers $1/l$ and $(k+\tau)/l$, $k = 0, \ldots, l-1$. The roots of V_l are the q-series for the "A-coefficient" of the curves whose associated lattices are spanned by $1/m$ and τ, in the first case, and by 1 and $(k+\tau)/l$, $k = 0, \ldots, l-1$ in the other l cases. For the first of these curves the A-coefficient is, from (4.2), l^4 times the coefficient associated to the lattice spanned by 1 and $l\tau$. Therefore from (6.2) its series (conventionally scaled) is

$$-3l^4\left(1 + 240\sum_{n=1}^{\infty}\sigma_3(n)q^{ln}\right).$$

The series (in powers of $w = q^{1/l}$) for the coefficients for the other l curves can be obtained by substituting $(k+\tau)/l$ for τ in (6.2). They are

$$3\left(1+240\sum_{n=1}^{\infty}\sigma_3(n)\zeta^{kn}w^n\right), \quad k=0,\ldots,l-1$$

where as in the last section $w = e^{2\pi i\tau/l}$ and $\zeta = e^{2\pi i/l}$.

We again use Newton's formulas and compute the power sums of these roots rather than their elementary symmetric functions. Let

$$\alpha(u) = 1+240\sum_{n=1}^{\infty}\sigma_3(n)u^n$$

and suppose that $a(r,n)$ are defined by

$$\alpha^r(u) = \sum_{n=0}^{\infty}a(r,n)u^n.$$

The r-th power sum is

$$\begin{aligned} &3^r\left(l^{4r}\alpha^r(q^l)+\sum_{k=0}^{l-1}\alpha^r(\zeta^k w)\right) \\ =&3^r\left(l^{4r}\sum_{n=0}^{\infty}a(r,n)q^{ln}+l\sum_{n=0}^{\infty}a(r,ln)q^n\right) \end{aligned} \tag{7.1}$$

In this case we need to expand the q-series of the r-th power sum to $M(2r)$ terms since the A's have degree 4. If l_* is the largest prime that we wish to process, all the q-series that we need can be easily be computed from the first $l_*M(2l_*+2)$ coefficients of the α^j, $j = 1,\ldots,l_*+1$. The work for this part is $l_*^5/18$ field operations with classical algorithms. However, this is essentially all the one-time work there is, and it is all concentrated in these polynomial multiplications of high degree polynomials. Thus it is possible to save time with fast multiplication techniques.

Remark: The l-th *Hecke operator* applied to a modular form $H(\tau)$ of degree d is

$$l^{d-1}H(l\tau)+\frac{1}{l}\sum_{k=0}^{l-1}H\left(\frac{k+\tau}{l}\right).$$

It is interesting to note that the r-th power sum of the roots of V_l, which is given in (7.1) above is actually (l times) the l-th Hecke operator applied to the modular form A^r. We do not see how to make use of this observation.

Here are the corresponding formulas for B'. Let us denote by W_l the monic minimum polynomial satisfied by B' over $\mathbf{Q}[A, B]$. Let

$$\beta(u) = 1 - 504 \sum_{n=1}^{\infty} \sigma_5(n) u^n$$

and suppose that $b(r, n)$ are defined by

$$\beta^r(u) = \sum_{n=0}^{\infty} b(r, n) u^n.$$

Then the r-th power sum of the roots of W_l is given by

$$2^r \left(l^{6r} \sum_{n=0}^{\infty} b(r, n) q^{ln} + l \sum_{n=0}^{\infty} b(r, ln) q^n \right) \tag{7.2}$$

Here the r-th power sum has generalized degree $6r$ so that we need to compute $M(3r)$ terms of their q-series. In the end the work involved is about $l_*^5/8$ field operations.

8. The size of the integers needed in a one-time computation.

The one-time work of our algorithm is essentially the work involved in expressing the power sums of the roots of U_l, V_l, W_l as polynomials in A, B (or more efficiently as polynomials in E_1, E_2 and Δ.) The most efficient method is the last method of section 6.

At any stage in the one-time work we have the option of reducing all our computed results modulo p and doing the rest of our work over $\mathrm{GF}(p)$. As long as we are calculating over the integers we will calculate modulo a collection of single-precision primes and combine results in the end with the Chinese Remainder Theorem. So long as we proceed this way we need to know maximum size of the final integers that we are computing. It may be posssible to estimate this by doing the calculations in floating point arithmetic. Here we may have some annoying, but not insurmountable, problems with overflow. Another possibly much more serious difficulty is round-off error.

Probably the most expensive part of the one-time work is the calculation of the powers of the series

$$\beta(u) = 1 - 504 \sum_{n=1}^{\infty} \sigma_5(n) u^n.$$

A practical method is to find first the powers of

$$\lambda(u) = \sum_{n=1}^{\infty} \sigma_5(n) u^n$$

find the powers of $\beta(u)$ later with the help of the binomial theorem.

Let us estimate the size of the largest coefficient that we will obtain. We have

$$\sigma_5(n) = \sum_{d|n} \frac{n^5}{d^5} \leq n^5 \sum_{d=1}^{\infty} \frac{1}{d^5} = \zeta(5) n^5.$$

Therefore $\lambda(u)$ is dominated term by term by $\zeta(5)\sum_{n=1}^{\infty} n^5 u^n$ which is in turn dominated by

$$C\sum_{n=0}^{\infty}\binom{n+4}{5}u^n = Cu(1-u)^{-6},$$

where $C = 5!\zeta(5)$. It follows that $\lambda^m(u)$ is dominated by $C^m(1-u)^{-6m}$ The highest term we will need in λ^m is about $m^2/2$. We can obtain an approximation to the size of the coefficient of $u^{m^2/2}$ in $\lambda^m(u)$:

$$C^m\binom{m^2/2+6m-1}{6m-1} \approx C^m\frac{(m^2/2+6m)^{6m}}{(6m/e)^{6m}} = C^m\left(\frac{e}{12}\right)^{6m}(m+12)^{6m}.$$

The largest value m that we need here is about $\log_e p$. This shows that the number of bits needed for the largest coefficient is order of magnitude $6(\log_e p)\log_2(\log_e p)$. A similar analysis could be done of the sizes of the other large integers that are needed for one-time work.

9. Finding rational relations.

We now consider briefly our original plan of expressing p_2 and p_3 as rational functions of p_1, A and B. A method something like Method 2 of section 6 is possible. We simply look for linear relations among monomials in p_1, p_2 A and B, which are of degree at most one in p_2. We may restrict our search to homogeneous linear combinations of monomials. The problem is that we do not know what the lowest degree is where we can expect to find non-trivial linear relations. We have had to experiment to find out what this lowest degree is for each l. If we attempt to solve the resulting homogeneous systems of linear equations with rational arithmetic, we need high precision. Thus we would like to have recourse to solving modulo a set of primes. Now, however, we have the added complication that we do not know how to scale the solutions so that we have integer coefficients and the integers, when we do find them, may well be larger than the coefficients we needed to compute in the previous section.

We still have the option of solving modulo the characteristic of F. This is probably a practical alternative. Here the main objection is that the systems of equations are larger than those we needed to solve in Method 3 of section 6. Also it does not seem possible to produce the systems in triangular form as we did above.

An alternative to looking for rational relations is just to look for relations among all monomials of lowest generalized degree in, say, p_1, p_2, A and B. The resulting equations are not necessarily linear in p_2 but they are of lower degree so it is possible that there will be only one rational solution for p_2 given the values of p_1, A and B allowing us to avoid the awkwardness of having to guess the match-ups. Also less work is involved in solving the lower degree relations for p_2.

Appendix I.

Here we prove the facts claimed in section 3. These must certainly be present in the literature of elliptic curves. Our object here is to provide easily accessible proofs.

Theorem: Suppose that l is an odd prime. Then $\psi_l(x)$ regarded as a polynomial in $\mathbf{Q}[A, B, x]$, is irreducible.

Proof: Let l be an odd prime. Regarding $\psi_l(x)$ as a polynomial in x with coefficients in $\mathbf{Z}[A,B]$, we note that it has no factors which are polynomials in A and B since the leading coefficient is l. Suppose that it factored into two polynomials of positive degree in x. Then it would factor with integer coefficients. Thus it suffices to prove that there exists a prime p and elements A and B of GF(p) such that $\psi_l(x)$ is irreducible mod p. Choose integers a and b such that the polynomial $\lambda^2 + a\lambda + b$ is primitive mod l. Choose p with $p \equiv b$ mod l and $|a| < 2\sqrt{p}$. Then according to [U, Theorem 4] there exists a curve $|E|$ defined over GF(p) with $t = |E_p| - 1 - p = a$. For this curve then the Frobenius operator ϕ will act as a single cycle of length $l^2 - 1$ on the l-torsion points. It will also therefore act as a single cycle on the x-coordinates of the l-torsion points. It follows that $\psi_l(x)$ is a power of an irreducible polynomial modulo this prime. However we also know that $\psi_l(x)$ has distinct roots mod p for any $p \neq l$ and any A and B with $4A^3 + 27B^2 \neq 0$. This proves the theorem.

Theorem: Suppose that l is an odd prime. Then the polynomial $U_l(x)$ regarded as a polynomial in $\mathbf{Q}[A,B,x]$ is irreducible.

Proof: Let l be an odd prime. U_l is the monic polynomial whose zeroes are the sums of x-coordinates from 1-dimensional subspaces of $E[l]$. If $l = 3$, $U_l = \psi_l/3$ which is irreducible. If $l > 3$, from the discussion section 6, we know that the coefficients of U_l are polynomials in $\mathbf{Z}[A,B]$. Thus it suffices to prove that, for some A and B and $p > 3$, U_l is irreducible mod p. The argument of the previous theorem shows that there exist p, A and B for which the Frobenius operator is transitive on the roots of U_l so that, for this choice of p, A and B, U_l is a power of an irreducible polynomial.

Thus we need to show that, for at least one of these curves, the polynomial U_l has distinct roots. We know that the polynomial has distinct roots over the complex numbers since our q-series expansions exhibit distinct roots in the field of formal Laurent series in $w = q^{1/l}$ with rational coefficients. Thus the discriminant is non-zero over fields of characteristic zero. However the discriminant is also a homogeneous function of A and B of generalized degree $2l(l+1)$ since it is the product of the squares of the differences of the $l+1$ roots and each root has generalized degree 2. Thus there are at most $2l(l+1)$ distinct values of the the ratio of A^3 to B^2 over any field for which the U_l has repeated roots. We shall show in the Lemma below that there are at most six possible values for the cardinality of the set of rational points for a curve over GF(p) if we are given the ratio of A^3 to B^2. Thus there are at most $12l(l+1)$ possible cardinalities for the rational points of a curve over GF(p) for which U_l has repeated roots.

Assuming the Lemma we now choose $p \equiv b$ mod l such that there are more than $12l(l+1)$ integers t with $t \equiv a$ mod l and $|t| \leq 2\sqrt{p}$. Thus at least one of these values of t arises from a curve for which U_l has distinct roots. For any such curve U_l will be irreducible. This completes the proof of the theorem.

Lemma: Given p, and the ratio of A^3 to B^2, there are at most two possible values for the number of rational points on curve $y^2 = x^3 + Ax + B$ if $A, B \neq 0$, at most 4 values if $A \neq 0$, $B = 0$ and at most 6 values if $A = 0$, $B \neq 0$.

Proof of Lemma: We consider the case $A, B \neq 0$. It suffices to show that if we are given two curves, one determined by A, B and the other by A', B', and $A^3/B^2 = A'^3/B'^2$

and A/B has the same quadratic character as A'/B', then the two curves have the same number of elements. In this case we can choose $\lambda \neq 0$ such that $\lambda^2 A/B = A'/B'$. It is now routine to check that the A, B-curve can be converted to the A', B'-curve by replacing x by x/λ^2 and y by y/λ^3.

Similarly if $A \neq 0$, $B = 0$, the cardinality depends only on the coset of A modulo the fourth powers in the multiplicative group of F, while if $A = 0$, $B \neq 0$, the cardinality depends only on the coset of B modulo the sixth powers in the multiplicative group of F. This proves the Lemma.

These two theorems prove some claims made in section 3. In particular the last theorem implies that p_1 generates the field L over the field F in the notation of section 3 as claimed. Indeed we now know that p_1 is an element of L and the root of an irreducible polynomial of degree $l+1$ with coeffcients in F, which is the same as the degree of the extension L/F.

Appendix II.

Eric Liverance has suggested an improvement to our method of finding a factor of the division polynomial. This improvement is based on a paper of Stark ([St]). We think that this idea could be used to shorten and clarify the exposition. It may also shorten the running time of our version of Schoof's algorithm.

Suppose we are given an elliptic curve E with equation $y^2 = x^3 + Ax + B$. Given a small prime l, on page 9 we define a curve E' over the complex numbers whose coordinate functions are x' and y', and whose equation is

$$(y')^2 = (x')^3 + A'x' + B'.$$

We can consider x' as a function on E, and since it is even (i.e. $x'(P) = x'(-P)$), x' is a rational function of x alone. Write

$$x'(x) = f(x)/g(x)$$

where f and g are relatively prime polynomials. It is easy to see that this implies that $g = h^2$, and h is a polynomial of degree $\frac{1}{2}(l-1)$ which is a factor of the division polynomial ψ_l. Furthermore the degree of f is l.

To find f and g in the complex case we would write

$$x'(x) \cdot g(x) = f(x).$$

Then we could replace x and x' by the Laurent series for the Weierstrass $\wp$-functions for E and E' respectively. These series only depend on A, B, A', and B'. We are given A and B, and we can compute A' and B' using the methods of Section 7. Comparing terms in the above equation gives a system of linear equations which must determine f and g. The polynomial g is then the square of the desired factor of ψ_l.

For the case of positive characteristic p, the coefficients of these Laurent series make sense up to about degree p. Since f has degree l and g has degree $l-1$, if l is small compared to p, the technique described above should work in this case too.

If all of the details work out, we would not need the polynomial U_l, nor would we need Elkies' method of finding $p_4, p_5, \ldots$ from p_1, p_2, and p_3, i.e. we would not need Sections 5 and 6. Since the work in our version of Schoof's algorithm is dominated by finding the roots of V_l and W_l, there would be little improvement in running time because of this. However, the discussion on page 18 could then be revised. We would have to solve one equation of degree $l+1$ half of the time and two equations for the other half. Hence our method with Liverance's improvement should require about 50% more work than Elkies' method.

References.

[A] Tom Apostol, Modular Functions and Dirichlet Series in Number Theory, Springer-Verlag, New York, 1976.

[CR] L. S. Charlap and David P. Robbins, An Elementary Introduction to Elliptic Curves, CRD Expository Report No. 31, IDA–CRD Log No. 82299, December, 1988.

[E] Noam Elkies, Explicit Isogenies, preprint.

[L] Serge Lang, Elliptic Functions, Addison-Wesley, Reading, 1973.

[Sc] Rene Schoof, Elliptic Curves over Finite Fields and the Computation of Square Roots mod p, Math Comp. 44(1985), 483–494.

[Se] Jean-Pierre Serre, A Course in Arithmetic, Springer-Verlag, New York, 1973.

[St] H. M. Stark, Class-Numbers of Complex Quadractic Fields in Modular Functions of One Variable I, Lecture Notes in Math. 320, Springer-Verlag, New York, 1973.

[U] Emmanuela Ughi, On the Number of Points of Elliptic Curves over a Field and a Problem of B. Segre, European Journal of Combinatorics (1983) 4,263–270.