Basic Theorems of Elliptic Curves

Jonathan Cass

10/14/09

Abstract

An elliptic curve E is the locus of solutions to a degree 3 equation. Many interesting results are obtained by looking at the structure of the solutions whose coordinates are in various fields. We denote the solutions to E with coordinates in a ring F by $E(F)$. We can introduce a group structure on $E(F)$, and then examine what kinds of groups we get when F is one of a number of different fields. In this talk we will discuss the structure of $E(\mathbb{C}), E(\mathbb{R}), E(\mathbb{Q})$, and $E(\mathbb{Z}_p)$. The main theorems to be proven are that $E(\mathbb{C})$ is isomorphic to a torus and that $E(\mathbb{Q})$ is a finitely generated group.

Content

The Basic Definition

An elliptic curve is the locus of solutions to a polynomial equation in two variables of the form

$$y^2 = x^3 + ax^2 + bx + c$$

or, as is birationally equivalent, of the form

$$y^2 = x^3 + ax + b$$

we will use either form as is convenient.

An elliptic curve is said to be non-singular if it has 3 distinct roots. This is equivalent to the condition that its discriminant $\Delta = -16(4a^3 + 27b^2)$ is non-zero (here the E is written in the second form, so a and b are the coefficient on x and the constant term, respectively).

The Group Law

One of the reasons that elliptic curves are interesting is that they have a natural group structure. To start with, we introduce an identity element $\mathcal{O}$ at infinity. What this means is that we are actually looking at elliptic curves in projective space, but we will work with homogeneous coordinates and simply keep in mind

that there is a "point at infinity" that is always present. Now, given two points P and Q on E, we can draw the line through them and this line will, in general, intersect E at a third point - say $P \oplus Q$. Then we define $P + Q = (P \oplus Q) \oplus \mathcal{O}$. It is not hard to prove that this operation preserves key properties of its points - specifically, if P and Q are rational, then so is $P + Q$ and the same goes for real points. Thus we get a chain of subgroups

$$\mathcal{O} \leq E(\mathbb{Q}) \leq E(\mathbb{R}) \leq E(\mathbb{C})$$

It will be useful and illustrative to give an exact formula for $P + Q$. First of all, note that $x(P + Q) = x(P \oplus Q)$ and $y(P + Q) = -y(P \oplus Q)$ since the line through the point at infinity is vertical. Now, let $P = (x_1, y_1)$ and $Q = (x_2, y_2)$ be distinct points on E. Say $P \oplus Q = (x_3, y_3)$. Next, define $\lambda = \frac{x_2 - x_1}{y_2 - y_1}$ and assume that the line through P and Q has equation $y = \lambda x + \nu$. Then the three points at the intersection of the line and E satisfy

$$y^2 = (\lambda x + \nu)^2 = x^3 + ax^2 + bx + c$$

so expanding and rearranging we get

$$0 = x^3 + (a - \lambda^2)x^2 + (b - 2\lambda\nu)x + c - \nu^2$$

The three roots of this cubic will be the x-coordinates of the points P, Q, and $P \oplus Q$ - namely, x_1, x_2, x_3. Thus

$$x^3 + (a - \lambda^2)x^2 + (b - 2\lambda\nu)x + c - \nu^2 = (x - x_1)(x - x_2)(x - x_3)$$

Comparing coefficients of the constant term, we obtain $a - \lambda^2 = -x_1 - x_2 - x_3$ so that

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

Then plugging this back into the equation for the line, we see that

$$y_3 = \lambda x_3 + \nu = \lambda^3 - (a + x_1 + x_2)\lambda + \nu$$

Therefore

$$P + Q = (\lambda^2 - a - x_1 - x_2, -\lambda^3 + (a + x_1 + x_2)\lambda - \nu$$

Of course, this only works if P and Q are distinct points, and if we are to turn E into a group, we must be able to compute $2P$. To compute $P + P$, we simply take the tangent line at P. The tangent line intersects (in most cases) with multiplicity 2, so it will intersect the curve in one other point, which we take to be $P \oplus P$. Then joining with $\mathcal{O}$ (or, equivalently, reflecting over the horizontal axis) gives $P + P$. This will not be of much concern in this talk, but I thought it might be useful to mention.

$E(\mathbb{C})$

Over $\mathbb{C}$, elliptic curves have a very simple and elegant form. They are precisely the surfaces of genus 1, also known as tori. To prove this we need a bit of complex analysis, but nothing too bad.

In the complex plane we have the notion of a "doubly periodic function". A function f is doubly periodic if there exist ω_1, ω_2 linearly independent such that $f(z+\omega_1) = f(z) = f(z+\omega_2)$. One of the simplest examples of these is called the Weierstrass $\wp$-function. Given a lattice $\mathbb{L} \in \mathbb{C}$ generated by ω_1, ω_2 we define

$$\wp(z) = \frac{1}{z^2} + \sum_{\substack{\omega \in \mathbb{L} \\ \omega \neq 0}} \frac{1}{(z-\omega)^2} - \frac{1}{\omega^2}$$

This is a meromorphic function in the complex plane, with poles of order 2 at each lattice point. It turns out that $\wp$ satisfies some very interesting differential equations - one of which is closely related to elliptic curves:

$$\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3$$

for all $z \notin \mathbb{L}$, where

$$\begin{aligned} g_2 &= 60 \sum_{\substack{\omega \in \mathbb{L} \\ \omega \neq 0}} \omega^{-4} \\ g_3 &= 140 \sum_{\substack{\omega \in \mathbb{L} \\ \omega \neq 0}} \omega^{-6} \end{aligned}$$

These properties are not too difficult to verify, but we will not go into the proofs here.

Let a non-singular elliptic curve E be given by

$$y^2 = x^3 + ax + b$$

Then we see that

$$(2y)^2 = 4x^3 + 4ax + 4b$$

Since E is non-singular, we know that $(-4a)^3 - 27(-4b)^2 \neq 0$ and therefore we can use the uniformization theorem to find a lattice $\mathbb{L}$ such that $g_2 = -4a$ and $g_3 = -4b$. Thus we have

$$(2y)^2 = 4x^3 - g_2x - g_3$$

a quick change of variables (and slight abuse of notation) gives

$$y^2 = 4x^3 - g_2x - g_3$$

Now we can look at the $\wp$-function associated with the lattice that we obtained by uniformization. This will satisfy the differential equation

$$\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3$$

with the same g_2 and g_3 that are now defining our elliptic curve. Therefore we have a map $\phi : \mathbb{C} \to E$ defined by $z \mapsto (\wp(z), \wp'(z))$.

We now make three claims:

1. ϕ is surjective.
2. ϕ is a homomorphism.
3. Ker $\phi = \mathbb{L}$.

To see that ϕ is surjective, take any point $(x_0, y_0) \in E$. Then $\frac{1}{\wp(z)-x_0}$ is an elliptic function and therefore has a pole (use Liouville's theorem) so that $\wp(z) - x_0$ has a zero. Therefore there is a $z_0 \in \mathbb{C}$ such that $\wp(z_0) = x_0$. So $\wp'(z_0)^2 = y_0^2$ and (taking $-z_0$ instead of z_0 if necessary) $\wp'(z_0) = y_0$. Thus we see that $\phi(z_0) = (x_0, y_0)$ as claimed.

The second claim is saying that

$$\phi(z_1 + z_2) = \phi(z_1) + \phi(z_2)$$

where the addition on the left side is the regular addition of complex numbers, while the addition on the right is the group law on the elliptic curve. This proposition turns out to be deeper than I originally thought and we will not be able to prove it here. It follows from the theory of divisors, however, and can be seen in Silverman's "The Arithmetic of Elliptic Curves".

Finally, it is clear that the kernel of ϕ is the lattice $\mathbb{L}$, since the identity on the curve is the point at infinity. Therefore Im $\phi = \mathbb{C}/\mathbb{L}$ and we are done.

$E(\mathbb{R})$

The real points on a curve, $E(\mathbb{R})$, are very easy to describe. Any elliptic curve has either 1 or 3 real roots. If E has precisely 1 root, then it will be isomorphic to the circle group - otherwise it will have 3 roots and will be isomorphic to a product of $\mathbb{Z}_2$ and the circle group.

$E(\mathbb{Q})$

The rational points are perhaps the most interesting. There are two main theorems that, together, describe the structure of $E(\mathbb{Q})$ very well - up to one mystery. We have the

Mordell-Weil Theorem.
Given a non-singular elliptic curve E, $E(\mathbb{Q})$ forms a finitely generated group.

So $E(\mathbb{Q}) \cong E_{tors} \oplus \mathbb{Z}^r$ for some r. The next theorem classifies $T(E)$:

Mazur's Theorem.

Given a non-singular elliptic curve E, the torsion subgroup of $E(\mathbb{Q})$ is one of the following fifteen groups:

$$\mathbb{Z}/N\mathbb{Z} \qquad N \in \{1, 2, ..., 10, 12\}$$

or

$$\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2N\mathbb{Z} \qquad 1 \leq N \leq 4$$

In fact, given an elliptic curve E it is not too hard to figure out which of these torsion subgroups E possesses. Another interesting theorem on the torsion subgroup is that it consists entirely of points with integer coordinates - this is the Nagell-Lutz theorem. The converse, however, is not true. It is important to note that the Nagell-Lutz theorem implies that if the rank of E is zero, then all rational points are integral.

Since we know that $E(\mathbb{Q})$ is finitely generated and what the torsion part looks like, the only question left about $E(\mathbb{Q})$ is how to compute the rank r. This turns out to be a very deep and difficult question and is the subject of the famous Birch Swinnerton-Dyer Conjecture. More on this later, if there is time.

Mazur's theorem is extremely difficult, and I haven't the slightest idea how it is proved - but I can find a reference for anyone who is interested. The Mordell-Weil theorem, on the other hand, is more tractable and we will discuss the proof here. We follow very closely the treatment given in Silverman and Tate's "Rational Points on Elliptic Curves". We won't have time to prove all the different parts of the theorem, but we will get the ideas of most of the proof.

Before we begin the proof, we need one more definition: given $r = \frac{a}{b} \in \mathbb{Q}$ written in lowest terms, we define the height $H(r) = \max\{a, b\}$. Then, since we want the height to act additively (rather than multiplicatively) we define $h(r) = \log H(r)$. We then define the height of a point $P = (x, y)$ as $h(P) = h(x)$.

To prove the Mordell-Weil theorem, we need the following four lemmas:

Lemma 1. *For every real number M, the set*

$$\{P \in E(\mathbb{Q}) : h(P) \leq M\}$$

is finite.

Lemma 2. *Let P_0 be a fixed rational point on E. There is a constant κ_0 depending on P_0 and E such that*

$$h(P + P_0) \leq 2h(P) + \kappa_0 \qquad \text{for all } P \in E(\mathbb{Q})$$

Lemma 3. *There is a constant κ, depending on E, such that*

$$h(2P) \geq 4h(P) - \kappa \qquad \text{for all } P \in E(\mathbb{Q})$$

Lemma 4. *The index* $(E(\mathbb{Q}) : 2E(\mathbb{Q}))$ *is finite.*

We will first indicate how the Mordell-Weil theorem is proved, using these lemmas - and then go back and prove them individually. So, assuming lemmas 1-4 we proceed as follows:

By lemma 4, we know that $(E(\mathbb{Q}) : 2E(\mathbb{Q}))$ is finite and so we can take a representative from each coset, say $Q_1, ..., Q_n$. Take $P \in E(\mathbb{Q})$, then there is an index i_1 such that

$$P - Q_{i_1} \in 2E(\mathbb{Q})$$

So for some $P_1 \in E(\mathbb{Q})$, we have

$$P - Q_{i_1} = 2P_1$$

We now iterate this process to get

$$\begin{aligned} P_1 - Q_{i_2} &= 2P_2 \\ P_2 - Q_{i_3} &= 2P_3 \\ &\vdots \\ P_{m-1} - Q_{i_m} &= 2P_m \end{aligned}$$

We can now represent P as

$$P = Q_{i_1} + 2P_1 = Q_{i_1} + 2Q_{i_2} + 4P_2 = Q_{i_1} + ... + 2^{m-1}Q_{i_m} + 2^m P_m$$

so that P is in the subgroup of $E(\mathbb{Q})$ generated by the Q_i and P_m. We will now use lemmas 2 and 3 to show that by choosing m large enough, $h(P_m)$ will be bounded, so that by lemma 1 there will only be finitely many needed. Therefore $E(\mathbb{Q})$ will be generated by the Q_i and these finitely many P_m.

Using lemma 2, we can find κ_i for $1 \leq i \leq n$ such that

$$h(P - Q_i) \leq 2h(P) + \kappa_i \qquad \text{for all } P \in E(\mathbb{Q})$$

Now take $\kappa' = \max\{\kappa_i | 1 \leq i \leq n\}$ so that

$$h(P - Q_i) \leq 2h(P) + \kappa' \qquad \text{for all } P \in E(\mathbb{Q}) \text{ and } 1 \leq i \leq n$$

We next use lemma 3 to find κ such that

$$h(2P) \geq 4h(P) - \kappa \qquad \text{for all } P \in E(\mathbb{Q})$$

then fix j and calculate

$$4h(P_j) \leq h(2P_j) + \kappa = h(P_{j-1} - Q_{i_j}) + \kappa \leq 2h(P_{j-1}) + \kappa' + \kappa$$

divide through by 4 and rearrange terms to get

$$\begin{aligned} h(P_j) &\leq \frac{1}{2}h(P_{j-1}) + \frac{\kappa' + \kappa}{4} \\ &= \frac{3}{4}h(P_{j-1}) - \frac{1}{4}(h(P_{j-1}) - (\kappa' + \kappa)) \end{aligned}$$

so, any time we have $h(P_{j-1}) \geq \kappa' + \kappa$, then

$$h(P_j) \leq \frac{3}{4}h(P_{j-1})$$

Therefore, for any starting point P, we can find an index m such that $h(P_m) \leq \kappa' + \kappa$ and

$$P = \sum_{i=1}^{n} a_i Q_i + 2^m P_m$$

so the set

$$\{Q_i\}_1^n \cup \{R : h(R) \leq \kappa' + \kappa\}$$

generates $E(\mathbb{Q})$ and since lemma 4 implies the first set is finite, and lemma 1 implies the second is finite, we know that $E(\mathbb{Q})$ is finitely generated.

□

So we now need to verify the four lemmas from above. Let E be a non-singular elliptic curve, given by $y^2 = x^3 + ax^2 + bx + c$. To prove lemma 1, we observe that it is equivalent to the fact that the set $\{P \in E(\mathbb{Q}) : H(P) \leq M\}$ is finite. But, given M, there are at most M^2 rational numbers with height up to M and therefore $2M^2$ rational points with these coordinates. So lemma 1 is clear.

Lemma 2 will require somewhat more work but is still not too bad. Essentially we use the addition formula derived when the group law was stated, and then the triangle inequality. Given distinct points $P = (x, y)$ and $P_0 = (x_0, y_0)$ we have

$$\begin{aligned} H(P + P_0) &= H(x(P + P_0)) \\ &= H\left(\frac{(y - y_0)^2}{(x - x_0)^2} - a - x - x_0\right) \\ &= H\left(\frac{(y - y_0)^2 - (x - x_0)^2(a + x + x_0)}{(x - x_0)^2}\right) \end{aligned}$$

In the numerator we will have $y^2 - x^3 = ax^2 + bx + c$ so that for some constants $A, ..., G$ we see that

$$H(P + P_0) = H\left(\frac{Ay + Bx^2 + Cx + D}{Ex^2 + Fx + G}\right)$$

If we write x and y in lowest terms then we will have $x = \frac{m}{e^2}$ and $y = \frac{n}{e^3}$, we make this substitution and clear denominators to get

$$H(P + P_0) = H\left(\frac{Ane + Bm^2 + Cme^2 + De^4}{Em^2 + Fme^2 + Ge^4}\right)$$

We do not know if this fraction is in lowest terms, but cancellation would only lessen the height. Therefore

$$H(P+P_0) \leq \max\{|Ane + Bm^2 + Cme^2 + De^4|, |Em^2 + Fme^2 + Ge^4|\}$$

It is evident that $e \leq H(P)^{1/2}$ and $m \leq H(P)$. Less apparent, but still true, is that $n \leq KH(P)^{3/2}$ for some K. This can be seen by substituting $(\frac{m}{e^2}, \frac{n}{e^3})$ into the equation for E

$$n^2 = m^3 + am^2e^2 + bme^4 + ce^6$$

so

$$|n^2| \leq |m^3| + |am^2e^2| + |bme^4| + |ce^6| \leq H(P)^3 + |a|H(P)^3 + |b|H(P)^3 + |c|H(P)^3$$

take $K = \sqrt{1 + |a| + |b| + |c|}$ and our claim is evident. So, we can then see that

$$\begin{aligned} |Ane + Bm^2 + Cme^2 + De^4| &\leq |Ane| + |Bm^2| + |Cme^2| + |De^4| \\ &\leq (|AK| + |B| + |C| + |D|)H(P)^2 \end{aligned}$$

and also

$$\begin{aligned} |Em^2 + Fme^2 + Ge^4| &\leq |Em^2| + |Fme^2| + |Ge^4| \\ &\leq (|E| + |F| + |G|)H(P)^2 \end{aligned}$$

therefore

$$H(P+P_0) \leq \max\{|AK| + |B| + |C| + |D|, |E| + |F| + |G|\}H(P)^2$$

and so, taking logarithms,

$$h(P+P_0) \leq 2h(P) + \kappa_0$$

Since this κ_0 does not depend on P, we are done with lemma 2.

Lemma 3 requires significant computation, but can be found in Silverman and Tate's book. In essence, it is more difficult because we need to ensure that there is not too much cancellation, which would lower the height considerably.

Finally, lemma 4 is sometimes called the Weak Mordell-Weil Theorem. It can be proven at many different levels of depth but is a fairly intense theorem in and of itself. We will not go into it here, but it can be looked up in many books.

$E(\mathbb{Z}_p)$

Attempts to understand $E(\mathbb{Z}_p)$ are still underway. The mysterious rank that was referenced earlier seems to depend heavily on these groups. The idea is that if an elliptic curve has a lot of primes p such that $\#E(\mathbb{Q})$ is large, then it is likely to have high rank. This is formalized in the Birch Swinnerton-Dyer conjecture, which actually says

$$\lim_{s\to 1}(s-1)^{-r}L(E,s) = \Omega \#Ш(E) 2^r R(E) (\#E_{tors}(\mathbb{Q}))^{-2} \prod_p c_p$$

Most of these terms are well beyond the scope of this talk (and my current understanding), but, fortunately, there are many other ways of talking about the conjecture. Specifically, there is a weaker form that we can discuss here.

First of all, we must define the L-series of an elliptic curve. We recall that we have continually required our elliptic curve to be non-singular. This could potentially cause problems when we reduce the equation for E, modulo p. The new discriminant Δ_p in $\mathbb{Z}_p$ will satisfy $\Delta \equiv \Delta_p (\text{mod } p)$ and therefore will be zero if and only if $p|\Delta$. We therefore define good and bad reduction:

We say that an elliptic curve E has bad reduction at a prime p if $p|\Delta$. Otherwise, E is said to have good reduction.

Now we let $\mathbb{P}_E$ be the set of primes at which E has good reduction. Then the L-function for E is defined as

$$L(E,s) = \prod_{p \in \mathbb{P}_E} \frac{1}{1 - \frac{p+1-N(p)}{p^s} + \frac{p}{p^{2s}}}$$

Since $L(E,s)$ is an infinite product, we immediately run into issues of convergence. It is not hard to show that $L(E,s)$ is convergent when $\text{Re}s > \frac{3}{2}$ - however, except for in special cases, not much more than that is known. The Birch Swinnerton-Dyer conjecture therefore splits into two parts:

1. $L(E,s)$ has an analytic continuation defined at $s = 1$.
2. The order of vanishing of $L(E,s)$ at 1 is the rank of E

So implicit in the conjecture is the fact that if $E(\mathbb{Q})$ is infinite (i.e. has non-zero rank) then $L(E,1) = 0$. Coates and Wiles have proven this in the case where E has complex multiplication. Complex multiplication is a topic of a lot of current research but has a fairly easy definition.

Given any elliptic curve E, for each integer m, there is a multiplication-by-m map from E to itself. For most elliptic curves, there are no other endomorphisms. Any curve E whose endomorphism ring is strictly greater than $\mathbb{Z}$ is said to have complex multiplication. For some reason this endows it with interesting properties - I am currently trying to learn more about this and maybe will give another talk on it when I understand some of it.