

18.783 Elliptic Curves

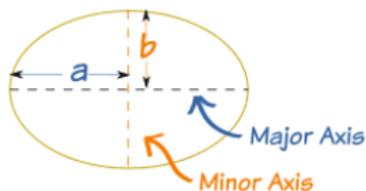
Lecture 1

Andrew Sutherland

February 3, 2015

What is an elliptic curve?

The equation $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$ defines an **ellipse**.



An ellipse, like all conic sections, is a curve of genus 0. It is **not an elliptic curve**. Elliptic curves have genus 1.

The area of this ellipse is πab . What is its circumference?

The circumference of an ellipse

Let $y = f(x) = b\sqrt{1 - x^2/a^2}$.

Then $f'(x) = -rx/\sqrt{a^2 - x^2}$, where $r = b/a < 1$.

Applying the arc length formula, the circumference is

$$4 \int_0^a \sqrt{1 + f'(x)^2} dx = 4 \int_0^a \sqrt{1 + r^2 x^2 / (a^2 - x^2)} dx$$

With the substitution $x = at$ this becomes

$$4a \int_0^1 \sqrt{\frac{1 - e^2 t^2}{1 - t^2}} dt,$$

where $e = \sqrt{1 - r^2}$ is the eccentricity of the ellipse.

This is an **elliptic integral**. The integrand $u(t)$ satisfies

$$u^2(1 - t^2) = 1 - e^2 t^2.$$

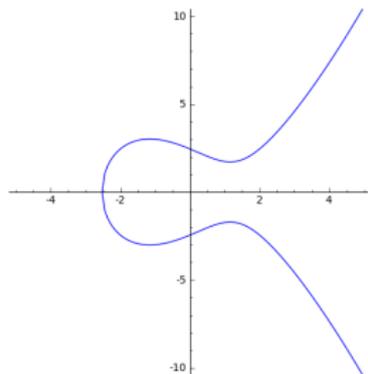
This equation defines an **elliptic curve**.

An elliptic curve over the real numbers

With a suitable change of variables, every elliptic curve with real coefficients can be put in the standard form

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

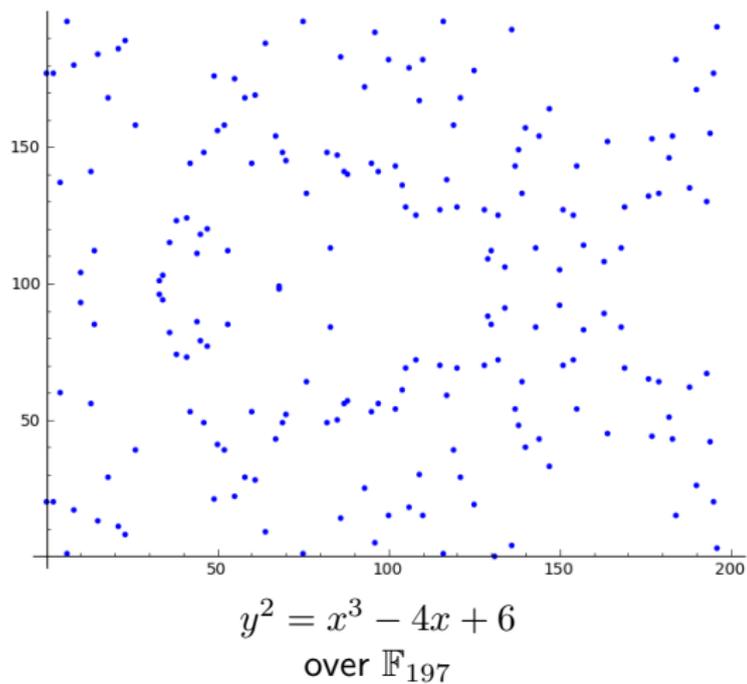
for some constants A and B . Below is an example of such a curve.



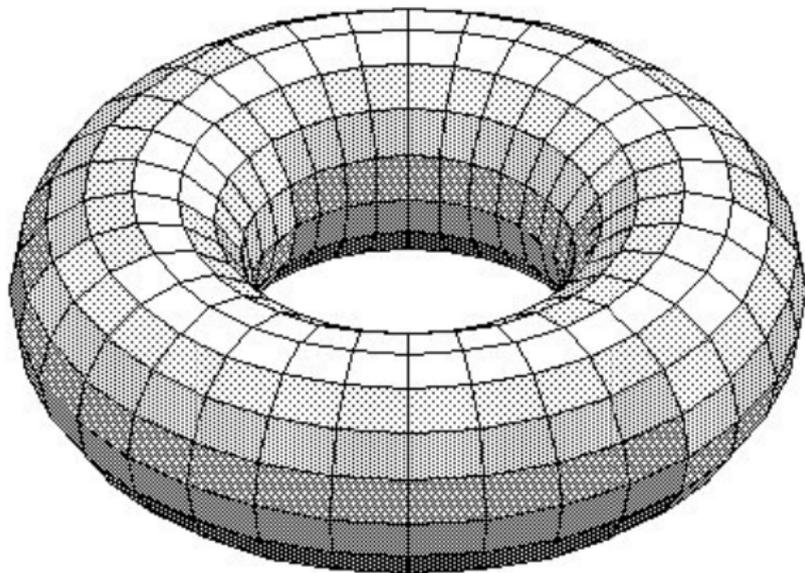
$$y^2 = x^3 - 4x + 6$$

over \mathbb{R}

An elliptic curve over a finite field



An elliptic curve over the complex numbers



An elliptic curve over \mathbb{C} is a compact manifold of the form \mathbb{C}/L , where $L = \mathbb{Z} + \omega\mathbb{Z}$ is a 2-d lattice in the complex plane.

Definitions

Definition

An **elliptic curve** is a smooth projective curve of genus 1 with a distinguished point.

Definition (more explicit)

An **elliptic curve** (over a field k) is a smooth projective curve of genus 1 (defined over k) with a distinguished (k -rational) point.

Not every smooth projective curve of genus 1 is an elliptic curve, it needs to have at least one rational point!

For example, the curve defined by $y^2 = -x^4 - 1$ is not an elliptic curve over \mathbb{Q} , even though it is a smooth projective curve of genus 1.

The projective plane

Definition

The **projective plane** is the set $\mathbb{P}^2(k)$ of all nonzero triples (x, y, z) in k^3 modulo the equivalence relation $(x, y, z) \sim (\lambda x, \lambda y, \lambda z)$.

The **projective point** $(x : y : z)$ is the equivalence class of (x, y, z) .

Points of the form $(x : y : 1)$ are called **affine points**.

They form an affine (Euclidean) plane $\mathbb{A}^2(k)$ embedded in $\mathbb{P}^2(k)$.

Points of the form $(x : y : 0)$ are called **points at infinity**.

These consist of the points $(x : 1 : 0)$ and the point $(1 : 0 : 0)$, which form the **line at infinity**, a projective line $\mathbb{P}^1(k)$ embedded in $\mathbb{P}^2(k)$.

Plane projective curves

Definition

A **plane projective curve** C_f/k is a homogeneous polynomial $f(x, y, z)$ with coefficients in k .¹

For any field K containing k , the **K -rational points** of C_f form the set

$$C_f(K) = \{(x : y : z) \in \mathbb{P}^2(K) \mid f(x, y, z) = 0\}.$$

A point $P \in C_f(K)$ is **singular** if $\frac{\partial f}{\partial x}$, $\frac{\partial f}{\partial y}$, $\frac{\partial f}{\partial z}$ all vanish at P .

C_f is **smooth** (or **nonsingular**) if there are no singular points in $C_f(\bar{k})$.

Every polynomial equation $g(x, y) = h(x, y)$ of degree d determines a projective curve C_f of degree d with $f(x, y, 1) = g(x, y) - h(x, y)$.

We often specify projective curves with affine equations, but we always mean to define a **projective curve**.

¹Fine print: up to scalar equivalence and with no repeated factors in $\bar{k}[x, y, z]$.

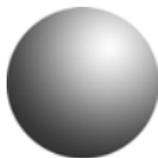
Examples of plane projective curves over \mathbb{Q}

affine equation	$f(x, y, z)$	points at ∞
$y = mx + b$	$y - mx - bz$	$(1 : m : 0)$
$x^2 + y^2 = 1$	$x^2 + y^2 - z^2$	none
$x^2 - y^2 = 1$	$x^2 - y^2 - z^2$	$(1 : 1 : 0), (1, -1, 0)$
$y^2 = x^3 + Ax + B$	$y^2z - x^3 - Axz^2 - Bz^3$	$(0 : 1 : 0)$
$x^2 + y^2 = 1 - x^2y^2$	$x^2z^2 + y^2z^2 - z^4 + x^2y^2$	$(1 : 0 : 0), (0 : 1 : 0)$

The first four curves are smooth (provided that $4A^3 + 27B^2 \neq 0$).
 The last curve is singular (both points at infinity are singular).

Genus

Over \mathbb{C} , an irreducible projective curve is a compact manifold that is topologically a sphere with handles. The number of handles is the genus.



genus 0



genus 1



genus 2



genus 3

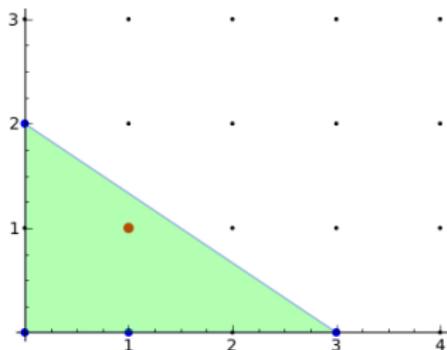
In fact, the genus can be defined algebraically over any field, not just \mathbb{C} .

Newton polytopes

Definition

The **Newton polytope** of a polynomial $f(x, y) = \sum a_{ij}x^i y^j$ is the convex hull of the set $\{(i, j) : a_{ij} \neq 0\}$ in \mathbb{R}^2 .

An easy way to compute the genus of a (sufficiently general) irreducible curve defined by an affine equation $f(x, y) = 0$ is to count the integer lattice points in the interior of its Newton polytope:



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

Weierstrass equations

Let $A, B \in k$ with $4A^3 + 27B^2 \neq 0$, and assume $\text{char}(k) \neq 2, 3$.

The (short/narrow) **Weierstrass equation** $y^2 = x^3 + Ax + B$ defines a smooth projective genus 1 curve over k with the rational point $(0 : 1 : 0)$.

In other words, an elliptic curve!

Up to isomorphism, **every** elliptic curve over k can be defined this way.

The general Weierstrass equation

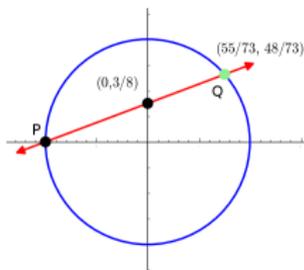
$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$$

works over any field, including those of characteristic 2 and 3.

Rational points in genus 0

Let C be a smooth projective curve over \mathbb{Q} of genus 0 (a unit circle, say), with a rational point P (let's use $(-1, 0, 1)$).

Any line ℓ with rational slope t that passes through P intersects C in exactly one "other" point $Q \in C(\mathbb{Q})$ (when ℓ is a tangent, $Q = P$). Conversely, for every $Q \in C(\mathbb{Q})$ the line \overline{PQ} is either vertical or has a rational slope t .



Treating the vertical line as the point at infinity on the projective line $\mathbb{P}^1(\mathbb{Q})$, there is a rational map from $C(\mathbb{Q})$ and $\mathbb{P}^1(\mathbb{Q})$, and vice versa.

In fact every genus 0 curve with a rational point is isomorphic to $\mathbb{P}^1(\mathbb{Q})$.

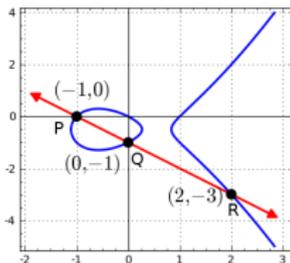
All genus 0 curves with a rational point are essentially the same!
(and this is true for any field, not just \mathbb{Q})

Rational points in genus 1

Now let E be an elliptic curve over \mathbb{Q} defined by a Weierstrass equation.

If P is a rational point and ℓ is a line through P with rational slope, it is not necessarily true that ℓ intersects E in another rational point.

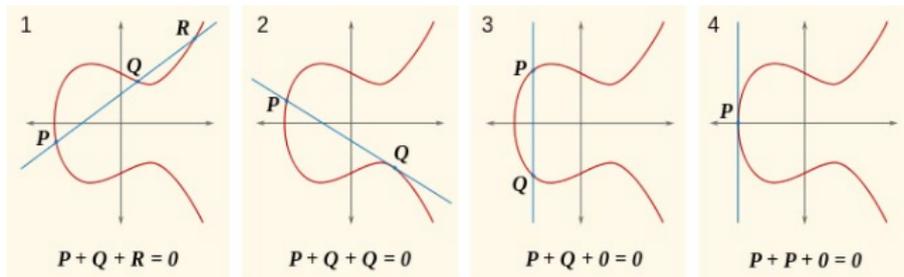
However, if P and Q are two rational points on E , then the line \overline{PQ} intersects E in a third rational point R (this follows from Bezout's theorem and a little algebra). This allows us to generate many new rational points from old ones (but not necessarily all of them!).



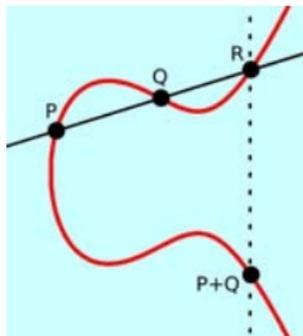
Even better, it allows us to define a group operation on $E(\mathbb{Q})$, or on $E(k)$ for any elliptic curve defined over any field k .

The elliptic curve group law

Three points on a line sum to zero.



Zero is the point at infinity.



The elliptic curve group law

With addition defined as above, the set $E(k)$ becomes an abelian group.

- ▶ The point $(0 : 1 : 0)$ at infinity is the identity element 0 .
- ▶ The inverse of $P = (x : y : z)$ is the point $-P = (x : -y : z)$.
- ▶ Commutativity is obvious: $P + Q = Q + P$.
- ▶ Associativity is not so obvious: $P + (Q + R) = (P + Q) + R$.

The computation of $P + Q = R$ is purely algebraic. The coordinates of R are rational functions of the coordinates of P and Q , and can be computed over any field.

By adding a point to itself repeatedly, we can compute $2P = P + P$, $3P = P + P + P$, and in general, $nP = P + \cdots + P$ for any positive n .

We also define $0P = 0$ and $(-n)P = -nP$.

Thus we can perform **scalar multiplication** by any integer n .

The group $E(k)$

When $k = \mathbb{C}$, the group operation on $E(\mathbb{C}) \simeq \mathbb{C}/L$ is just addition of complex numbers, modulo the lattice L .

When $k = \mathbb{Q}$ things get much more interesting. The group $E(\mathbb{Q})$ may be finite or infinite, but in every case it is **finitely generated**.

Theorem (Mordell 1922)

The group $E(\mathbb{Q})$ is a finitely generated abelian group. Thus

$$E(\mathbb{Q}) \simeq T \oplus \mathbb{Z}^r,$$

*where the **torsion subgroup** T is a finite abelian group corresponding to the elements of $E(\mathbb{Q})$ with finite order, and r is the **rank** of $E(\mathbb{Q})$.*

It may happen (and often does) that $r = 0$ and T is the trivial group. In this case the only element of $E(\mathbb{Q})$ is the point at infinity.

The group $E(\mathbb{Q})$

The torsion subgroup T of $E(\mathbb{Q})$ is well understood.

Theorem (Mazur 1977)

The torsion subgroup of $E(\mathbb{Q})$ is isomorphic to one of the following:

$$\mathbb{Z}/n\mathbb{Z} \quad \text{or} \quad \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2m\mathbb{Z},$$

where $n \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12\}$ and $m \in \{1, 2, 3, 4\}$.

The ranks of elliptic curves over \mathbb{Q}

The rank r of $E(\mathbb{Q})$ is **not** well understood.

Here are some of the things we do not know about r :

1. Is there an algorithm that is guaranteed to correctly compute r ?
2. Which values of r can occur?
3. How often does each possible value of r occur, on average?
4. Is there an upper limit, or can r be arbitrarily large?

We do know a few things about r . We can compute r in most cases where r is small. When r is large often the best we can do is a lower bound; the largest example is a curve with $r \geq 28$ due to Elkies (2006).

The ranks of elliptic curves over \mathbb{Q}

The most significant thing we know about r is a bound on its average value over all elliptic curves (suitably ordered).

The following result is very recent and is still being improved.

Theorem (Bhargava, Shankar 2010-2012)

The average rank of all elliptic curves over \mathbb{Q} is less than 1.

In fact we now know the average rank is greater than 0.2 and less than 0.9; it is believed to be exactly $1/2$ (half rank 0, half rank 1).

The group $E(\mathbb{F}_p)$

Over a finite field \mathbb{F}_p , the group $E(\mathbb{F}_p)$ is necessarily finite.

On average, the size of the group is $p + 1$, but it varies, depending on E .

The following theorem of Hasse was originally conjectured by Emil Artin.

Theorem (Hasse 1933)

The cardinality of $E(\mathbb{F}_p)$ satisfies $\#E(\mathbb{F}_p) = p + 1 - t$, with $|t| \leq 2\sqrt{p}$.

The fact that $E(\mathbb{F}_p)$ is a group whose size is not fixed by p is unique to genus 1 curves. This is the basis of many useful applications.

For curves C of genus $g = 0$, we always have $\#C(\mathbb{F}_p) = p + 1$.

For curves C of genus $g > 1$, the set $C(\mathbb{F}_p)$ does not form a group.

Reducing elliptic curves over \mathbb{Q} modulo p

Let E/\mathbb{Q} be an elliptic curve defined by $y^2 = x^3 + Ax + B$, and let p be a prime that does not divide the **discriminant** $\Delta(E) = -16(4A^3 + 27B^2)$.

The elliptic curve E is then said to have **good reduction** at p .

If we reduce A and B modulo p , we obtain an elliptic curve $\overline{E} = E \bmod p$ defined over the finite field $\mathbb{F}_p \simeq \mathbb{Z}/p\mathbb{Z}$.

Thus from a single curve E/\mathbb{Q} we get an infinite family of curves, one for each prime p where E has good reduction.

Now we may ask, how does $\#\overline{E}(\mathbb{F}_p)$ vary with p ?

The Sato-Tate conjecture

We know $\#\overline{E}(\mathbb{F}_p) = p + 1 - a_p$ for some integer a_p with $|a_p| \leq 2\sqrt{p}$.
So let $x_p := a_p/\sqrt{p}$. Then x_p is a real number in the interval $[-2, 2]$.

What is the distribution of x_p as p varies? Is it uniform?

<http://math.mit.edu/~drew>

The Sato-Tate conjecture, open for nearly 50 years, was recently proven. Taylor received the 2014 Breakthrough Prize in Mathematics for work that led to this the proof (and other results).

Theorem (Taylor et al., 2006 and 2008)

*Let E/\mathbb{Q} be an elliptic curve without complex multiplication.
Then the x_p have a semi-circular distribution.*

The Birch and Swinnerton-Dyer conjecture

There is believed to be a relationship between the infinite sequence of integers a_p associated to an elliptic curve E/\mathbb{Q} and the rank r .

The **L -function** $L_E(s)$ of an elliptic curve E/\mathbb{Q} is a function of a complex variable s that “encodes” the infinite sequence of integers a_p .

For the “bad” primes that divide $\Delta(E)$, one defines a_p to be 0, 1, or -1 , depending on the type of singularity E has when reduced mod p .

$$L_E(s) = \prod_{\text{bad } p} (1 - a_p p^{-s})^{-1} \prod_{\text{good } p} (1 - a_p p^{-s} + p^{1-2s})^{-1} = \sum_{n=0}^{\infty} a_n n^{-s}$$

The Birch and Swinnerton-Dyer conjecture

Based on extensive computer experiments (back in the 1960s!), Bryan Birch and Peter Swinnerton-Dyer made the following conjecture.

Conjecture (Birch and Swinnerton-Dyer)

Let E/\mathbb{Q} be an elliptic curve with rank r . Then

$$L_E(s) = (s - 1)^r g(s),$$

for some complex analytic function $g(s)$ with $g(1) \neq 0, \infty$. In other words, r is equal to the **order of vanishing** of $L_E(s)$ at 1.

They later made a more precise conjecture that also specifies the constant coefficient a_0 of $g(s) = \sum_{n \geq 0} a_n (s - 1)^n$.

Fermat's Last Theorem

Theorem (Wiles et al. 1995)

$x^n + y^n = z^n$ has no positive integer solutions for $n > 2$.

It suffices to consider n prime.

Suppose $a^n + b^n = c^n$ with $a, b, c > 0$ and $n > 3$ (the case $n = 3$ was proved by Euler). Consider the elliptic curve E/\mathbb{Q} defined by

$$y^2 = x(x - a^n)(x - b^n).$$

Serre and Ribet proved that E **is not modular**.

Wiles (with Taylor) proved that every semistable elliptic curve over \mathbb{Q} , including E , **is modular**. So no solution $a^n + b^n = c^n$ can possibly exist. We now know that all elliptic curves E/\mathbb{Q} are modular.

Applications of elliptic curves over finite fields

There are several factors that make elliptic curves over finite fields particularly well suited to practical applications:

- ▶ There are many groups available, even when the finite field is fixed.
- ▶ The underlying group operation can be made very efficient.
- ▶ There are techniques to construct a group of any desired size.
- ▶ The representation of group elements appears to be “opaque”.

There are three particular applications that we will explore in detail:

1. factoring integers
2. primality proving
3. cryptography

Factoring integers with elliptic curves

The elliptic curve factorization method (ECM), due to Lenstra, is a randomized algorithm that attempts to factor an integer n using random elliptic curves E/\mathbb{Q} with a known point $P \in E(\mathbb{Q})$ of infinite order.

For each curve E , the algorithm attempts to find a scalar multiple of P equivalent to zero in $\overline{E}(\mathbb{F}_p)$, for some unknown prime p dividing n .

The algorithm will succeed when $\#\overline{E}(\mathbb{F}_p)$ is sufficiently **smooth**, meaning that all its prime factors are small.

The expected running time is subexponential in $\log p$ and otherwise polynomial in $\log n$. No other algorithm with this property is known.

When p is large (say $\log p > \log^{2/3} n$), faster algorithms are known, but they typically use ECM as a subroutine.

Primality proving with elliptic curves

Elliptic curve primality proving (ECP) was introduced by Goldwasser and Kilian and later improved by Atkin and Morain (and Bach).

Let n be an integer that we believe to be prime and let $b = \sqrt{n}$.

Suppose one can find E/\mathbb{Q} with the following property: for every prime $p|n$, the group $\overline{E}(\mathbb{F}_p)$ contains a point of order $m > b + 1 + 2\sqrt{b}$.

Then the Hasse bound implies that $p > b = \sqrt{n}$ for all primes $p|n$.

Therefore n can have only one prime divisor, itself!

Heuristically, the expected running time of ECP is quasi-quartic in $\log n$, and, in practical terms, it is the fastest general purpose algorithm known for primality proving.

The deterministic AKS algorithm has been **proven** to run in polynomial time, and randomized versions of AKS have expected running times that are quasi-quartic in $\log n$. But they are still slower than ECP in practice.

The discrete log problem

Problem: Given a point $P \in E(\mathbb{F}_q)$ and $Q = nP$, determine n .

This is known as the **discrete log problem**, a term that originates from the analogous problem in the multiplicative group \mathbb{F}_q^* : given $a \in \mathbb{F}_q^*$ and $b = a^n$, determine $n = \log_a b$.

In the group \mathbb{F}_q^* , this problem can be solved in time that is subexponential in $\log q$, but no comparable result is known for the group $E(\mathbb{F}_q)$.

In fact, the best known algorithm for solving the discrete log problem in $E(\mathbb{F}_q)$ takes time $\Omega(\sqrt{q})$, which is fully exponential in $\log q$.

This allows cryptographic systems based on the elliptic curve discrete log problem to use key sizes that are much smaller than other systems.

Of course we do not have any proof that the elliptic curve discrete log problem is hard (just as we have no proof that factoring integers is hard).

Diffie-Hellman key exchange

Diffie and Hellman proposed a method for two parties to establish a secret key over a public network, based on the discrete log problem. Their method is generic, it works in a cyclic subgroup of any given group.

Let E/\mathbb{F}_p be an elliptic curve with a point $P \in E(\mathbb{F}_p)$.

Alice and Bob, who both know E and P , establish a secret S as follows:

1. Alice chooses a random integer a and sends $Q_a = aP$ to Bob.
2. Bob chooses a random integer b and sends $Q_b = bP$ to Alice.
3. Alice computes $aQ_b = abP = S$ and Bob computes $bQ_a = baP = S$.

The coordinates of S depend on the random integer ab and can be hashed to yield a shared secret consisting of $\log_2 ab$ random bits.²

An eavesdropper may know E , P , Q_a and Q_b , but not a , b , or S . It is believed that computing S from the known values is as hard as computing $a = \log_P Q_a$ and $b = \log_P Q_b$ (but this is not proven).

²As written, this protocol is vulnerable to a man-in-the-middle attack.

Ephemeral Diffie-Hellman

With *ephemeral* Diffie-Hellman (ECDHE) the elliptic curve E is fixed, but a new base point P is chosen for each key exchange.

This provides what is known as *perfect forward secrecy*, which compartmentalizes the security of each communication session (breaking one session should not make it easier to break others).

The ECDHE protocol was adopted by Google in late 2011 and is now used by most internet sites to establish a secure session, including:

Amazon, Bing, Dropbox, Facebook, Flickr, GitHub, Instagram, LinkedIn, MSN, Netflix, Pinterest, PirateBay, Quora, Snapchat, SoundCloud, Spotify, StackOverflow, Tumblr, Twitter, Uber, Vimeo, Vine, Yahoo, Yelp, YouTube, Wikipedia, Wordpress, . . .

MIT OpenCourseWare
<http://ocw.mit.edu>

18.783 Elliptic Curves

Spring 2015

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.