25.1 Overview of Mordell's theorem

In the last lecture we proved that the torsion subgroup of the rational points on an elliptic curve $E/\mathbb{Q}$ is finite. In this lecture we will prove a special case of Mordell's theorem, which states that $E(\mathbb{Q})$ is finitely generated. By the structure theorem for finitely generated abelian groups, this implies

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

where $\mathbb{Z}^r$ is a free abelian group of rank r, and T is the (necessarily finite) torsion subgroup.[1] Thus Mordell's theorem provides an alternative proof that T is finite, but unlike our earlier proof, it does not provide an explicit method for computing T. Indeed, Mordell's theorem is notably *ineffective*; it does not give us a way to compute a set of generators for $E(\mathbb{Q})$, or even to determine the rank r. It is a major open question as to whether there exists an algorithm to compute r; it is also not known whether r can be uniformly bounded.[2]

Mordell's theorem was generalized to number fields (finite extensions of $\mathbb{Q}$) and to abelian varieties (recall that elliptic curves are abelian varieties of dimension one) by André Weil and is often called the Mordell-Weil theorem. All known proofs of Mordell's theorem (and its generalizations) essentially amount to two proving two things:

(a) $E(\mathbb{Q})/2E(\mathbb{Q})$ is a finite group.

(b) For any fixed $Q \in E(\mathbb{Q})$, the *height* of $2P + Q$ is greater than the height of P for all but finitely many P.

We note that there is nothing special about 2 here, any integer $n > 1$ works.

We will explain what (b) means in a moment, but let us first note that we really do need some sort of (b); it is not enough to just prove (a). To see why, consider the additive abelian group $\mathbb{Q}$. the quotient $\mathbb{Q}/2\mathbb{Q}$ is certainly finite (it is the trivial group), but $\mathbb{Q}$ is not finitely generated. To see this, note that for any finite $S \subseteq \mathbb{Q}$, we can pick a prime p such that under the canonical embedding $\mathbb{Q} \subseteq \mathbb{Q}_p$ we have $S \subseteq \mathbb{Z}_p$, and therefore $\langle S\rangle \subseteq \mathbb{Z}_p$, but we never have $\mathbb{Q} \not\subseteq \mathbb{Z}_p$.

The *height* of a projective point $P = (x : y : z)$ with $x, y, z \in \mathbb{Z}$ sharing no common factor is defined as

$$H(P) := \max(|x|, |y|, |z|),$$

where $|\ |$ is the usual archimedean absolute value on $\mathbb{Q}$. The height $H(P)$ is a positive integer that is independent of the representation of the representation of P, and for any bound B, the set

$$\{P \in E(\mathbb{Q}) : H(P) \le B\}$$

is finite, since it cannot possibly have more than $(2B + 1)^3$ elements. We will actually use a slightly more precise notion of height, the *canonical height*, which we will define later.

Now let us suppose that we have proved (a) and (b), and see why this implies that $E(\mathbb{Q})$ is finitely generated. Since $E(\mathbb{Q})/2E(\mathbb{Q})$ is finite, for any sufficiently large B the finite set $S = \{P \in E(\mathbb{Q}) : H(P) \le B\}$ must contain a set of representatives for $E(\mathbb{Q})/2E(\mathbb{Q})$, and

[1] Any finitely generated *abelian* torsion group must be finite; this does not hold for nonabelian groups.
[2] Most number theorists think not, but there are some notable dissenters.

we can pick B so that (b) holds for all $Q \in S$ and $P \notin S$. If S does not generate $E(\mathbb{Q})$, then there is a point $P_0 \in E(\mathbb{Q}) - \langle S\rangle$ of minimal height $H(P_0)$. Since S contains a set of representatives for $E(\mathbb{Q})/2E(\mathbb{Q})$, we can write P_0 in the form

$$P_0 = 2P + Q,$$

for some $Q \in S$ and $P \in E(\mathbb{Q})$. Since $P_0 \notin \langle S\rangle$, we must have $P \notin \langle S\rangle$, but (b) implies $H(P) < H(P_0)$, contradicting the minimality of $H(P_0)$. So the set $E(\mathbb{Q}) - \langle S\rangle$ must be empty and S is a finite set of generators for $E(\mathbb{Q})$.

We should note that this argument does not yield an algorithm to compute S because we do not have an effective bound on B (we know B exists, but not how big it is).

25.2 Elliptic curves with a rational point of order 2

In order to simplify the presentation, we will restrict our attention to elliptic curves $E/\mathbb{Q}$ that have a rational point of order 2 (to prove the general case one can work over a cubic extension of $\mathbb{Q}$ for which this is true). In short Weierstrass form any point of order 2 is an affine point of the form $(x_0, 0)$. After replacing x with $x + x_0$ we obtain an equation for E of the form

$$E\colon y^2 = x(x^2 + ax + b),$$

on which $P = (0, 0)$ is a point of order two. Since E is not singular, the cubic on the RHS has no repeated roots, which implies

$$b \neq 0, \qquad a^2 - 4b \neq 0.$$

The algebraic equations for the group law on curves of this form are slightly different than for curves in short Weierstrass form; the formula for the inverse of a point is the same, we simply negate the y-coordinate, but the formulas for addition and doubling are slightly different. To add two affine points $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$ with $x_1 \neq x_2$, as in Lecture 23 we consider the line L through P_1 and P_2 with equation

$$L\colon (y - y_1) = \lambda(x - x_1),$$

where $\lambda = (y_2 - y_1)/(x_2 - x_1)$. Solving for y and plugging into equation for E, we have

$$\begin{aligned} \lambda^2 x^2 &= x(x^2 + ax + b) \\ 0 &= x^3 + (a - \lambda^2)x^2 + \cdots \end{aligned}$$

The x-coordinate x_3 of the third point in the intersection $L \cap E$ is a root of the cubic on the RHS, as are x_1 and x_2, and the sum $x_1 + x_2 + x_3$ must be equal to the negation of the quadratic coefficient. Thus

$$\begin{aligned} x_3 &= \lambda^2 - a - x_1 - x_2, \\ y_3 &= \lambda(x_1 - x_3) - y_1, \end{aligned}$$

where we computed y_3 by plugging x_3 into the equation for L and negating the result. The doubling formula for $P_1 = P_2$ is the same, except now $\lambda = (3x^2 + 2ax + b)/(2y)$.

25.3 2-isogenies

In order to prove that $E(\mathbb{Q})/2E(\mathbb{Q})$ is finite, we need to understand the image of the multiplication-by-2 map [2]. We could use the doubling formula derived above to do this, but it turns out to be simpler to decompose [2] as a composition of two isogenies

$$[2] = \hat{\varphi} \circ \varphi,$$

where $\varphi\colon E \to E'$ and $\hat{\varphi}\colon E' \to E$ for some elliptic curve E' that we will determine. The kernel of φ will be $\{O, P\}$, where $P = (0,0)$ is our rational point of order 2. Similarly, the kernel of $\hat{\varphi}$ will be $\{O', P'\}$, where O' is the distinguished point on E' and P' is a rational point of order 2 on E'.

Recall from Lecture 24 that for any isogeny $\varphi\colon E \to E'$ we have an injective map

$$\ker\varphi \to \operatorname{Aut}(\overline{\mathbb{Q}}(E)/\varphi^*(\overline{\mathbb{Q}}(E')))$$

defined by $P \mapsto \tau_P^*$, where τ_P is the translation-by-P morphism. In our present situation there is only one non-trivial point in the kernel of φ, the point $P = (0,0)$, and it is rational, so we can work over $\mathbb{Q}$. We can determine both E' and the morphism φ by computing $\varphi^*(\mathbb{Q}(E'))$ as the fixed field of the automorphism $\tau_P^*\colon \mathbb{Q}(E) \to \mathbb{Q}(E)$.

Remark 25.1. This strategy applies in general to any separable isogeny with a cyclic kernel (a *cyclic isogeny*), all we need is a point P that generates the kernel.

For an affine point $Q = (x, y)$ not equal to $P = (0,0)$ the x-coordinate of $\tau_P(Q) = P+Q$ is $\lambda^2 - a - x$, where $\lambda = y/x$ is the slope of line throught P and Q. Using the curve equation for E, we can simplify this to

$$\lambda^2 - a - x = \frac{y^2 - ax^2 - x^3}{x^2} = \frac{bx}{x^2} = \frac{b}{x}.$$

The y-coordinate of $\tau_P(Q)$ is then $\lambda(0 - b/x) - 0 = -by/x^2$. Thus for $Q \notin \{O, P\}$ the map τ_P is given by

$$(x, y) \mapsto (b/x, -by/x^2).$$

To compute the fixed field of τ_P^*, note that if we regard the slope $\lambda = y/x$ as a function in $\mathbb{Q}(E)$, then composition with τ_P merely changes its sign. Thus

$$\tau_P^*(\lambda^2) = \left(\frac{-by/x^2}{b/x}\right)^2 = \left(\frac{-y}{x}\right)^2 = \lambda^2.$$

We also note that the point $Q + \tau_P(Q)$ is fixed by τ_P, hence the sum of the y-coordinates of Q and $\tau_P(Q)$ is fixed by τ_P (when represented as affine points $(x : y : 1)$). Thus

$$\tau_P^*(y - by/x^2) = \tau_P^*\left(\frac{x^2y - by}{x^2}\right) = \frac{(b/x)^2(-by/x^2) - b(-by/x^2)}{(b/x)^2} = y - by/x^2.$$

Note that $\lambda^2 = y^2/x^2 = x(x^2 + ax + b)/x^2 = x + a + b/x$, so let us define

$$X = x + a + b/x \qquad \text{and} \qquad Y = y(1 - b/x^2)$$

Then $\mathbb{Q}(X, Y)$ is a subfield of $E(\mathbb{Q}) = \mathbb{Q}(x, y)$ fixed by τ_P^*, hence a subfield of $\varphi^*(\mathbb{Q}(E'))$, and we claim that it is a subfield of index 2. To see this, note that

$$x = (X + Y\sqrt{X} - a)/2 \qquad \text{and} \qquad y = x\sqrt{X},$$

thus $[\mathbb{Q}(E):\mathbb{Q}(X,Y)] \le 2$ and $[\mathbb{Q}(E):\mathbb{Q}(X,Y)] \ne 1$ because $\mathbb{Q}(E)$ contains $x/y = \sqrt{X}$ and $\mathbb{Q}(X,Y)$ does not. We also know that $[\mathbb{Q}(E):\varphi^*(\mathbb{Q}(E'))] \ge 2$, since $\ker\varphi \subseteq \mathbb{Q}(E)$ has order 2 and injects into $\operatorname{Aut}(\mathbb{Q}(E)/\varphi^*(\mathbb{Q}(E)))$, therefore $\varphi^*(\mathbb{Q}(E')) = \mathbb{Q}(X,Y)$.

Since φ^* is a field embedding, we have $\mathbb{Q}(E') \simeq \mathbb{Q}(X,Y)$. We now know the function field of E'; to compute an equation for E' we just need a relation between X and Y.

$$\begin{aligned} Y^2 &= y^2(1-b/x^2)^2 \\ &= x(x^2+ax+b)(1-2b/x^2-b^2/x^4) \\ &= X(x^2-2b+b^2/x^2) \\ &= X\big((x+b/x)^2-4b\big) \\ &= X\big((X-a)^2-4b\big) \\ &= X(X^2-2aX+a^2-4b). \end{aligned}$$

Let us now define $A = -2a$ and $B = a^2-4b$. Then the equation

$$Y^2 = X(X^2+AX+B)$$

has the same form as that of E, and since $B = a^2-4b \ne 0$ and $A^2-4B = 16b \ne 0$, it defines an elliptic curve E' with distinguished point $O' = (0:1:0)$, and the affine point $P' = (0,0)$ has order 2. The 2-isogeny $\varphi\colon E \to E'$ sends O to O' and each affine point (x,y) on E to $(X,Y) = (x+a+b/x,\ y(1-b/x^2))$ on E'.

Since E' has the same form has E, we can repeat the process above to compute the 2-isogeny $\hat\varphi\colon E' \to E$ that sends O' to O and (X,Y) to $(X+A+B/X, Y(1-B/X^2))$. One can then verify that

$$[2] = \hat\varphi \circ \varphi,$$

by composing $\hat\varphi$ and φ and comparing the result to the doubling formula on E.

But we can see this more directly by noting that $\ker(\hat\varphi\circ\varphi) = E[2]$ and

$$\deg(\hat\varphi\circ\varphi) = \deg\hat\varphi\deg\varphi = 2\cdot 2 = 4 = \#E[2] = \#\ker(\hat\varphi\circ\varphi).$$

Thus the injective homomorphism $E[2] \to \operatorname{Aut}(\overline{\mathbb{Q}}(E)/(\hat\varphi\circ\varphi)^*(\overline{\mathbb{Q}}(E)))$ is an isomorphism, and the same holds for $\operatorname{Aut}(\overline{\mathbb{Q}}(E)/[2]^*\overline{\mathbb{Q}}(E))$. Since we are in characteristic zero, both extensions are separable, and it follows from Galois theory that there is a unique subfield of $\overline{\mathbb{Q}}(E)$ fixed by the automorphism group $\{\tau_P^* : P \in E[2]\}$. Thus the function field embeddings $(\hat\varphi\circ\varphi)^*$ and $[2]^*$ are equal, and the corresponding morphisms must be equal (by the functorial equivalence of smooth projective curves and their function fields).

Remark 25.2. The construction and argument above applies quite generally. Given any finite subgroup H of $E(\bar k)$ there is a unique elliptic curve E' and separable isogeny $E \to E'$ with H as its kernel; see [2, Prop. III.4.12].

25.4 The weak Mordell-Weil theorem

We are now ready to prove that $E(\mathbb{Q})/2E(\mathbb{Q})$ is finite (in the case that $E(\mathbb{Q})$ has a rational point of order 2). This is a special case of what is known as the weak Mordell-Weil theorem, which says that $E(k)/nE(k)$ is finite, for any positive integer n and any number field k. Our strategy is to prove that $E(\mathbb{Q})/\varphi(E(\mathbb{Q}))$ is finite, where $\varphi\colon E \to E'$ is the 2-isogeny from the previous section. This will also show that $E'(\mathbb{Q})/\hat\varphi(E(\mathbb{Q}))$ is finite, and it will follow that $E/2E(\mathbb{Q})$ is finite.

We begin by characterizing the image of φ in $E'(\mathbb{Q})$.

Lemma 25.3. *An affine point $(X, Y) \in E'(\mathbb{Q})$ lies in the image of φ if and only if either $X \in \mathbb{Q}^{\times 2}$, or $X = 0$ and $a^2 - 4b \in \mathbb{Q}^{\times 2}$.*

Proof. Suppose $(X, Y) = \varphi(x, y)$. T If $X \neq 0$ then $X = (y/x)^2 \in \mathbb{Q}^{\times 2}$. If $X = 0$ then $x(x^2 + ax + b) = 0$, and $x \neq 0$ (since $\varphi(0, 0) = O'$), so $x^2 + ax + b = 0$ has a rational solution, which implies $a^2 - 4b \in \mathbb{Q}^{\times 2}$.

Conversely, if $X \in \mathbb{Q}^{\times 2}$ then $x = (X + Y\sqrt{X} - a)/2$ and $y = x\sqrt{X}$ gives a point $(x, y) \in E(\mathbb{Q})$ for which $\varphi(x, y) = (X, Y)$, and if $X = 0$ and $a^2 - 4b \in \mathbb{Q}^{\times 2}$, then $x^2 + ax + b$ has a nonzero rational root x for which $\varphi(x, 0) = (0, 0) = (X, Y)$. $\square$

Now let us define the map $\pi\colon E'(\mathbb{Q}) \to \mathbb{Q}^{\times}/\mathbb{Q}^{\times 2}$ by

$$(X, Y) \mapsto \begin{cases} X & \text{if } X \neq 0, \\ a^2 - 4b & \text{if } X = 0, \end{cases}$$

and let $\pi(O') = 1$.

Lemma 25.4. *The map $\pi\colon E'(\mathbb{Q}) \to \mathbb{Q}^{\times}/\mathbb{Q}^{\times 2}$ is a group homomorphism.*

Proof. By definition, $\pi(O') = 1$, so π preserves the identity element and behaves correctly on sums involving O'. and since $\pi(P) = \pi(-P)$ and the square classes of X and $1/X$ are the same, π preserves inverses. We just need to verify $\pi(P + Q) = \pi(P)\pi(Q)$ for affine points P, Q that are not inverses.

So let P and Q be affine points whose sum is an affine point R, let $Y = \ell X + m$ be the line L containing P and Q (the line L is not vertical because $P + Q = R \neq O'$). Plugging the equation for Y given by L into the equation for E' gives

$$\begin{aligned} (\ell X + m)^2 &= X(X^2 + AX + B) \\ 0 &= X^3 + (A - \ell^2)x^2 + (B - \ell m)x - m^2. \end{aligned}$$

The X-coordinates X_1, X_2, X_3 of P, Q, R are all roots of the cubic on the RHS, hence their product is equal to m^2, the negation of the constant term. Thus $X_1X_2X_3$ is a square, which means that $\pi(P)\pi(Q)\pi(P + Q) = 1$, and therefore $\pi(P)\pi(Q) = 1/\pi(P + Q) = \pi(P + Q)$, since $\pi(P + Q)$ and $1/\pi(P + Q)$ are in the same square-class of $\mathbb{Q}^{\times}$. $\square$

Lemma 25.5. *The image of $\pi\colon E'(\mathbb{Q}) \to \mathbb{Q}^{\times}/\mathbb{Q}^{\times 2}$ is finite.*

Proof. Let (X, Y) be an affine point in $E'(\mathbb{Q})$ with $X \neq 0$, and let $r \in \mathbb{Z}$ be a square-free integer representative of the square-class $\pi(X, Y)$. We will show that r must divide B, which clearly implies that $\operatorname{im} \pi$ is finite. The equation $Y^2 = X(X + aX + B)$ for E' implies that X and $X + aX + B$ lie in the same square-class, thus

$$\begin{aligned} X^2 + AX + B &= rs^2 \\ X &= rt^2, \end{aligned}$$

for some $s, t \in \mathbb{Q}^{\times}$. Let us write $t = \ell/m$ with $\ell, m \in \mathbb{Z}$ relatively prime. Plugging $X = rt^2$ into the first equation gives

$$\begin{aligned} r^2t^4 + Art^2 + B &= rs^2 \\ r^2\ell^4/m^4 + Ar\ell^2/m^2 + B &= rs^2 \\ r^2\ell^4 + Ar\ell^2m^2 + Bm^4 &= rm^4s^2, \end{aligned}$$

and since the LHS is an integer, so is the RHS. Let p be any prime dividing r. Then p must divide Bm^4, since it divides every other term. If p divides m then p^3 must divide $r^2\ell^4$, since it divides every other term, but then p divides ℓ, since r is squarefree, which is impossible because ℓ and m are relatively prime. So p does not divide m and therefore must divide B. This holds for every prime divisor of the squarefree integer r, so r divides B as claimed. □

Corollary 25.6. *$E'(\mathbb{Q})/\varphi(E(\mathbb{Q}))$ and $E(\mathbb{Q})/\hat{\varphi}(E(\mathbb{Q}))$ are finite.*

Proof. Lemma 25.3 implies that $\ker\pi = \varphi(E(\mathbb{Q}))$, thus $E'(\mathbb{Q})/\varphi(E(\mathbb{Q})) \simeq \operatorname{im}\pi$ is finite, and this remains true if we replace E with E' and φ with $\hat{\varphi}$. □

Corollary 25.7. *$E(\mathbb{Q})/2E(\mathbb{Q})$ is finite.*

Proof. The fact that $[2] = \hat{\varphi}\circ\varphi$ implies that each $\hat{\varphi}(E'(\mathbb{Q}))$-coset in $E(\mathbb{Q})$ can be partitioned into $2E(\mathbb{Q})$-cosets. Two points P and Q in the same $\hat{\varphi}(E'(\mathbb{Q}))$-coset lie in the same $2E(\mathbb{Q})$-coset if and only if $(P-Q) \in 2E(\mathbb{Q}) = (\hat{\varphi}\circ\varphi)(E(\mathbb{Q}))$, equivalently, $\hat{\varphi}^{-1}(P-Q) \in \varphi(E(\mathbb{Q})$. Thus the number of $2E(\mathbb{Q})$-cosets in each $\hat{\varphi}(E'(\mathbb{Q}))$-coset is precisely $E'(\mathbb{Q})/\varphi(E(\mathbb{Q}))$, thus

$$\#E(\mathbb{Q})/2E(\mathbb{Q}) = \#E(\mathbb{Q})/\hat{\varphi}(E'(\mathbb{Q}))\ \#E'(\mathbb{Q})/\varphi(E(\mathbb{Q}))$$

is finite. □

Remark 25.8. The only place in our work above where we really used the fact that we are working over $\mathbb{Q}$, as opposed to a general number field, is in the proof of Lemma 25.5. Specifically, we used the fact that the ring of integers $\mathbb{Z}$ of $\mathbb{Q}$ is a UFD, and that its unit group $\mathbb{Z}^\times$ is finite. Neither is true of the ring of integers $\mathcal{O}_k$ of a number field k, in general, but there are analogous facts that one can use; specifically, $\mathcal{O}_k$ is a Dedekind domain, hence ideals can be unique factored into prime ideals, the class number of $\mathcal{O}_k$ is finite, and $\mathcal{O}_k^\times$ is finitely generated. We also assumed that E has a rational point of order 2, but after a base extension to a number field we can assume this without loss of generality.

25.5 Height functions

Let k be any number field. Recall from Lecture 6 that (up to equivalence) the absolute values of k consist of non-archimedean absolute values, one for each prime ideal $\mathfrak{p}$ of the ring of integers $\mathcal{O}_k$ (these are the *finite places* of k), and archimedean absolute values, one for each embedding of k into $\mathbb{R}$ and one for each conjugate pair of embeddings of k into $\mathbb{C}$ (these are the *infinite places* of k). Let $\mathcal{P}_k$ denote the set of (finite and infinite) places of k.

For each place $p \in \mathcal{P}_k$ we want to normalize the associated absolute value $|\ |_p$ so that

(a) The product formula $\prod_{p\in\mathcal{P}_k} |x|_p = 1$ holds for all $x \in k^\times$.

(b) For any number field $k' \subseteq k$ and any place p of k' we have $\prod_{q|p} |x|_q = |x|_p$, where $q|p$ means that the restriction of $|\ |_q$ to k' is equivalent to $|\ |_p$.

Both requirements are satisfied by using the standard normalization for $\mathbb{Q}$, with

$$|x|_p = p^{-v_p(x)}$$

for $p < \infty$ and $|x|_\infty = |x|$, and then for each $q \in \mathcal{P}_k$ with $q|p$ defining

$$|x|_q = |N_{k_q/\mathbb{Q}_p}(x)|_p^{1/[k:\mathbb{Q}]},$$

where k_q and $\mathbb{Q}_p$ denote the completions of k at q and $\mathbb{Q}$ at p, respectively.[3]

Definition 25.9. The (absolute) *height* of a projective point $P = (x_0 : \cdots : x_n) \in \mathbb{P}^n(\overline{\mathbb{Q}})$ is

$$H(P) := \prod_{p \in \mathcal{P}_k} \max_i |x_i|_p,$$

where $k = \mathbb{Q}(x_0, \ldots, x_n)$. For any $\lambda \in \overline{\mathbb{Q}}^\times$, if we let $k = \mathbb{Q}(x_0, \ldots, x_n\lambda)$, then

$$\prod_{p \in \mathcal{P}_k} \max_i |\lambda x_i|_p = \prod_{p \in \mathcal{P}_k} \max_i(|\lambda|_p |x_i|_p) = \prod_{p \in \mathcal{P}_k} |\lambda|_p \prod_p \max_i |x_i| = \prod_{p \in \mathcal{P}_k} \max_i |x_i|,$$

thus $H(P)$ is well defined (it does not depend on a particular choice of $x_0, \ldots, x_n$).

For $k = \mathbb{Q}$ we can write $P = (x_0 : \cdots : x_n)$ with the $x_i \in \mathbb{Z}$ having no common factor. Then $\max |x_i|_p = 1$ for $p < \infty$ and $H(P) = \max_i |x_i|_\infty$; this agrees with the definition we gave earlier.

Lemma 25.10. *For all $P = (x_0 : \cdots : x_n) \in \mathbb{P}^n(\overline{\mathbb{Q}})$ we have $H(P) \geq 1$.*

Proof. Pick a nonzero x_j and let $k = \mathbb{Q}(x_0, \ldots, x_n)$. Then

$$H(P) = \prod_{p \in \mathcal{P}_k} \max_i |x_i|_p \ \geq \prod_{p \in \mathcal{P}_k} |x_j|_p = 1. \qquad \square$$

Definition 25.11. The *logarithmic height* of $P \in \mathbb{P}^n(\overline{\mathbb{Q}})$ is the nonnegative real number

$$h(P) := \log H(P).$$

We now consider how the height of a point changes when we apply a morphism to it. We will show that there for any fixed morphism $\phi\colon \mathbb{P}^m \to \mathbb{P}^n$ there are constants c and d (depending on ϕ) such that for any point $P \in \mathbb{P}^m(\overline{\mathbb{Q}})$ we have

$$dh(P) - c \ \leq \ h(\phi(P)) \ \leq \ dh(P) + c.$$

This can be written more succinctly write as

$$h(\phi(P)) = dh(P) + O(1),$$

where the $O(1)$ term indicates a bounded real function of P (the function $h(\phi(P)) - dh(P)$).

We first prove the upper bound; this is easy.

Lemma 25.12. *Let k be a number field and let $\phi\colon \mathbb{P}^n \to \mathbb{P}^m$ be a morphism $(\phi_0\colon \cdots : \phi_n)$ defined by homogeneous polynomials $\phi_i \in k[x_0, \ldots, x_n]$ of degree d. There is a constant c such that*

$$h(\phi(P)) \leq dh(P) + c$$

for all $P \in \mathbb{P}^n(\bar{k})$.

[3] The correctness of this definition relies on some standard results from algebraic number theory that we will not prove here; the details are not important, all we need to know is that a normalization satisfying both (a) and (b) exists, see [1, p. 9] or [2, pp. 225-227] for a more detailed exposition.

Proof. Let $c = N\prod_p \max_j |c_j|_p$, where c_j ranges over coefficients that appear in any ϕ_i, and N bounds the number of monomials appearing in any ϕ_i. If $P = (a_0 : \ldots : a_n)$ and $k = \mathbb{Q}(a_0, \ldots, a_n)$, then

$$H(\phi(P)) = \prod_{p \in P_k} \max_i |\phi_i(P)|_p \;\leq\; \prod_{p\in P_k} \max_{i,j} |c_j a_i^d|_p = cH(P)^d,$$

by the multiplicativity of $|\ |_p$ and the triangle inequality. The lemma follows. □

We now make a few remarks about the morphism $\phi\colon \mathbb{P}^n \to \mathbb{P}^m$ appearing in the lemma. Morphisms with domain $\mathbb{P}^n$ are tightly constrained, more so than projective morphisms in general, because the ideal of $\mathbb{P}^n$ (as a variety), is trivial; this means that the polynomials defining ϕ are essentially unique up to scaling. This has several consequences.

- The polynomials ϕ_i defining ϕ cannot have a common zero in $\mathbb{P}^n(\bar{k})$; otherwise there would be a point at which ϕ is not defined. This requirement is not explicitly stated because it is implied by the definition of a morphism as a regular map.
- The image of ϕ in $\mathbb{P}^m$ is either a point (in which case $d = 0$), or a subvariety of dimension n; if this were not the case then the polynomials defining ϕ would have a common zero in $\mathbb{P}^n(\bar{k})$. The fact that $\operatorname{im}\phi$ is a variety follows from the fact that projective varieties are complete (so every morphism is a closed map). In particular, if ϕ is non-constant then we must have $m \geq n$.
- If ϕ is non-constant, then $d = [k(\mathbb{P}^n) : \phi^*(k(\operatorname{im}\phi))]$ is equal to the degree of the ϕ_i. In particular, if $d = 1$ then ϕ is a bijection from $\mathbb{P}^n$ to its image. Note that this agrees with out definition of the degree of a morphism of curves.

Corollary 25.13. *It ϕ is any automorphism of $\mathbb{P}^n$, then*

$$h(\phi(P)) = h(P) + O(1). \tag{1}$$

Proof. We must have $d = 1$, and we can apply Lemma 25.12 to ϕ^{-1} as well. □

The corollary achieves our goal in the case $d = 1$ and $m = n$. If $d = 1$ and $m > n$, after applying a suitable automorphism to $\mathbb{P}^m$ we can assume that $\operatorname{im}\phi$ is the linear subvariety of $\mathbb{P}^m$ defined by $x_{n+1} = x_{n+2} = \cdots = x_{m+1} = 0$, and it is clear that the orthogonal projection $(x_0 : \cdots : x_m) \mapsto (x_0 : \cdots : x_n)$ does not change the height of any point in this subvariety. It follows that (2) holds whenever $d = 1$, whether $m = n$ or not.

We now prove the general case

Theorem 25.14. *Let k be a number field and let $\phi\colon \mathbb{P}^n \to \mathbb{P}^m$ be a morphism $(\phi_0\colon \cdots : \phi_n)$ defined by homogeneous polynomials $\phi_i \in k[x_0, \ldots, x_n]$ of degree d. Then*

$$h(\phi(P)) = dh(P) + O(1). \tag{2}$$

Proof. If $d = 0$ then ϕ is constant and the theorem holds trivially, so we assume $d > 0$. We will decompose ϕ as the composition of four morphisms: a morphism $\psi\colon \mathbb{P}^n \to \mathbb{P}^N$, an automorphism of $\mathbb{P}^N$, an orthogonal projection $\mathbb{P}^N \to \mathbb{P}^n \subseteq \mathbb{P}^m$, and an automorphism of $\mathbb{P}^m$. All but the morphism ψ change the logarithmic height of a point P by at most an additive constant that does not depend on P, and we will show that $h(\psi(P)) = dh(P)$.

The map $\psi = (\psi_0\colon \cdots : \psi_N)$ is defined as follows. We let $N = \binom{n+d}{d} - 1$, and take $\psi_0, \ldots, \psi_N$ to be the distinct monomials of degree d in the variables $x_0, \ldots, x_n$, in some order. Clearly the ψ_N have no common zero in $\mathbb{P}^n(\overline{\mathbb{Q}})$, so ψ defines a morphism $\mathbb{P}^n \to \mathbb{P}^N$. Let $P = (a_0\colon \cdots : a_n)$ be any point in $\mathbb{P}^n$, and let $k = \mathbb{Q}(a_0, \ldots, a_n)$. For each $p \in \mathcal{P}_k$,

$$\max_i |\psi_i(P)|_p = \max_j |a_j^d|_p = \max_j |a_j|_p^d = (\max_j |a_j|_p)^d,$$

and it follows that

$$H(\psi(P)) = \prod_{p\in\mathcal{P}_k} \max_i |\psi_i(P)|_p = \prod_{p\in\mathcal{P}_k} (\max_j |a_j|_p)^d = H(P)^d.$$

Thus $h(\psi(P)) = dh(P)$ as claimed. We now note that each ϕ_i is a linear combination of the ψ_j, thus ϕ induces an automorphism $\hat{\phi}\colon \mathbb{P}^N \to \mathbb{P}^N$, and after applying a second automorphism of $\mathbb{P}^N$ we may assume that the image of $\hat{\phi}\circ\psi$ in $\mathbb{P}^N$ is the variety defined by $x_{n+1} = \cdots = x_N = 0$. Taking the orthogonal projection from $\mathbb{P}^N$ to $\mathbb{P}^n$ embedded in $\mathbb{P}^m$ as the locus of $x_{n+1} = \cdots = x_m = 0$ does not change the height of any point, and we may then apply an automorphism of $\mathbb{P}^m$ to map this embedded copy of $\mathbb{P}^n$ to $\operatorname{im}\phi$. □

Remark 25.15. For an alternative proof of Theorem 25.14 using the Nullstellensatz, see [2, VIII.5.6].

Lemma 25.16. *Let $k/\mathbb{Q}$ be a finite Galois extension. Then $h(P^\sigma) = h(P)$ for all $P \in \mathbb{P}^n(k)$ and $\sigma \in \mathrm{Gal}(k/\mathbb{Q})$.*

Proof. The action of σ permutes $\mathcal{P}_k$, so if $P = (x_0 : \cdots : x_n)$ with $x_i \in k$, then

$$H(P^\sigma) = \prod_{p\in\mathcal{P}_k} \max_i |x_i^\sigma|_p = \prod_{p^\sigma\in\mathcal{P}_k} \max_i |x_i^\sigma|_{p^\sigma} = \prod_{p^\sigma\in\mathcal{P}_k} \max_i |x_i|_p = \prod_{p\in\mathcal{P}_k} \max_i |x_i|_p = H(P).$$

□

Remark 25.17. Lemma 25.16 also holds for $k = \overline{\mathbb{Q}}$.

Theorem 25.18 (Northcott)**.** *For any positive integers B, d, and n, the set*

$$\{P \in \mathbb{P}^n(k) : h(P) \le B \text{ and } [k:\mathbb{Q}] \le d\}$$

is finite.

Proof. Let $P = (x_0 : \cdots : x_n) \in \mathbb{P}^n(k)$ with $[k : \mathbb{Q}] \le d$. We can view each x_i as a point $P_i = (x_i : 1)$ in $\mathbb{P}^1(k)$, and we have

$$H(P) = \prod_{p\in\mathcal{P}_k} \max |x_i|_p \ge \max_i \prod_{p\in\mathcal{P}} \max_i(|x_i|_p, 1) = \max_i H(P_i).$$

Thus it suffices to consider the case $n = 1$, and we may assume $P = (x : 1)$ and $k = \mathbb{Q}(x)$.

Without loss of generality we may replace k by its Galois closure, so let $k/\mathbb{Q}$ be Galois with $\mathrm{Gal}(k/\mathbb{Q}) = \{\sigma_1, \ldots \sigma_d\}$. The point $Q = (x^{\sigma_1} : \cdots : x^{\sigma_d}) \in \mathbb{P}^{d-1}(k)$ is fixed by $\mathrm{Gal}(k/\mathbb{Q})$, hence by $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, so $Q \in \mathbb{P}^{d-1}(\mathbb{Q})$. By Lemma 25.16, $h(Q) = h(P)$, so we have reduced to the case $k = \mathbb{Q}$, and by the argument above we can also assume $n = 1$.

The set $\{P \in \mathbb{P}^1(\mathbb{Q}) : h(P) \le B\}$ is clearly finite; each P can be represented as a pair of relatively prime integers of which only finitely many have absolute value at most e^B. □

25.6 Canonical height functions on elliptic curves

Theorem 25.19 (Tate)**.** *Let S be a set and let $r > 1$ a real number. Let $\phi\colon X \to X$ and $h\colon X \to \mathbb{R}$ be functions such that $h \circ \phi = rh + O(1)$, and let*

$$\hat{h}_\phi(x) := \lim_{n\to\infty} \frac{1}{r^n} h(\phi^n(x)).$$

Then $\hat{h}_\phi$ is the unique function $S \to \mathbb{R}$ for which

(i) $\hat{h}_\phi = h + O(1)$;

(ii) $\hat{h}_\phi \circ \phi = r\hat{h}_\phi$.

Proof. Choose c so that $|\frac{1}{r}h(\phi(x)) - h(x)| \le \frac{c}{r}$ for all $x \in S$. For all $n > 1$ we have

$$\left|\frac{1}{r^n}h(\phi^n(x)) - \frac{1}{r^{n-1}}h(\phi^{n-1}(x))\right| = \frac{1}{r^{n-1}}\left|\frac{1}{r}h(\phi(\phi^{n-1}(x)) - h(\phi^{n-1}(x))\right| \le \frac{c}{r^{n-1}},$$

thus for all $x \in S$ the sequence $\frac{1}{r^n}h(\phi^n(x))$ converges, so $\hat{h}_\phi$ is well defined.

For all $x \in S$ we have

$$|\hat{h}_\phi(x) - h(x)| \le \sum_{n=1}^{\infty}\left|\frac{1}{r^n}h(\phi^n(x)) - \frac{1}{r^{n-1}}h(\phi^{n-1}(x))\right| \le \sum_{n=1}^{\infty}\frac{c}{r^n} = \frac{c}{r-1},$$

so (i) holds. Property (ii) is clear, and for uniqueness we note that if $f = h + O(1)$ and $f \circ \phi = rf$ then applying the construction above with h replaced by f yields $\hat{f}_\phi = \hat{h}_\phi$, but it is also clear that $\hat{f}_\phi = f$, so $f = \hat{h}_\phi$. ∎

We now want to apply Theorem 25.19 to the set $S = E(\overline{\mathbb{Q}})$ with $\phi = [2]$ the multiplication-by-2 map and $r = 4$, It might seem natural to let h be the height function on the projective plane $\mathbb{P}^2$ containing our elliptic curve E, but as E is a one-dimensional variety, it is better to work with $\mathbb{P}^1$, so we will use the image of E under the projection $\mathbb{P}^2 \to \mathbb{P}^1$ defined by $(x : y : z) \mapsto (x : z)$.

To understand how [2] operates on $\pi(E)$, we recall the formula to double an affine point $P = (x_1 : y_1 : 1)$ with $y_1 \ne 0$ computes the x-coordinate of $2P = (x_3 : y_3 : 1)$ via $x_3 = \lambda^2 - 2x_1$, with

$$\lambda^2 = \left(\frac{3x_1^2 + a_4}{2y_1}\right)^2 = \frac{9x_1^4 + 6a_4x_1^2 + a_4^2}{4y^2} = \frac{9x_1^4 + 6a_4x_1^2 + a_4^2}{4x_1^3 + 4a_4x_1 + 4a_6},$$

where we have used the curve equation $y^2 = x^3 + a_4x + a_6$ to get a formula that only depends on x_1. We then have

$$x_3 = \frac{9x_1^4 + 6a_4x_1^2 + a_4^2}{4x_1^3 + 4a_4x_1 + a_6} - 2x_1 = \frac{x_1^4 + 2a_4x_1^2 - 8a_6x_1 + a_4^2}{4x_1^3 + 4a_4x_1 + a_6}.$$

Putting this in projective form, we now define the map $\phi\colon \mathbb{P}^1 \to \mathbb{P}^1$ by

$$\phi(x : z) = (x^4 + 2a_4x^2z^2 - 8a_6xz^3 + a_4^2z^4 : 4x^3z + 4a_4xz^3 + a_6z^4).$$

The fact that $4a_4^3 + 27a_6^2 \ne 0$ ensures that the polynomials defining ϕ have no common zero in $\mathbb{P}^1(\overline{\mathbb{Q}})$, thus $\phi\colon \mathbb{P}^1 \to \mathbb{P}^1$ is a morphism of degree 4, and Theorem 25.14 implies that

$$h(\phi(P)) = 4h(P) + O(1).$$

Definition 25.20. Let E be an elliptic curve over a number field k. The *canonical height*

$$\hat{h}\colon E(\bar{k}) \to \mathbb{R}$$

is the function $\hat{h} = \hat{h}_\phi \circ \pi$, where $\hat{h}_\phi$ is the function given by Theorem 25.19, with $\phi\colon \mathbb{P}^1 \to \mathbb{P}^1$ as above and h the absolute height on $\mathbb{P}^1$. It satisfies $\hat{h}(2P) = 4\hat{h}(P)$ for all $P \in E(\mathbb{Q})$.

Theorem 25.21. *Let E be an elliptic curve over a number field k. For any bound B the set $\{P \in E(k) : \hat{h}(P) \leq B\}$ is finite.*

Proof. This follows immediately from Northcott's theorem and Theorem 25.19 part (i). ☐

Theorem 25.22 (Parallelogram Law)**.** *Let $\hat{h}$ be the canonical height function of an elliptic curve E over a number field k. Then for all $P, Q \in E(\bar{k})$ we have*

$$\hat{h}(P+Q) + \hat{h}(P-Q) = 2\hat{h}(P) + 2\hat{h}(Q)$$

Proof. This is a straight-forward but tedious calculation that we omit; see [2, VIII.6.2]. ☐

25.7 Proof of the Mordell's Theorem

With all the pieces in place we now complete the proof of Mordell's theorem for an elliptic curve $E/\mathbb{Q}$ with a rational point of order 2.

Theorem 25.23. *Let $E/\mathbb{Q}$ be an elliptic curve with a rational point of order 2. Then $E(\mathbb{Q})$ is finitely generated.*

Proof. By the weak Mordell-Weil theorem that we proved in §25.4 for this case we know that $E(\mathbb{Q})/2E(\mathbb{Q})$ is finite. So let us choose a bound B such that the set

$$S\colon = \{P \in E(\mathbb{Q}) : \hat{h}(P) \leq B\}$$

contains a set S_0 of representatives for $E(\mathbb{Q})/2E(\mathbb{Q})$. We claim that S generates $E(\mathbb{Q})$.

Suppose for the sake of obtaining a contradiction that this is not the case. Then there is a point $Q \in E(\mathbb{Q}) - \langle S \rangle$ of minimal height $\hat{h}(Q)$; the fact that every set of bounded height is finite implies that $\hat{h}$ takes on discrete values, so such a Q exists. There is then a point $P \in S_0 \subset S$ such that $Q = P + 2R$ for some $R \in E(\mathbb{Q})$. Since $Q \notin \langle S \rangle$, we must have $R \notin \langle S \rangle$, so $\hat{h}(R) \geq \hat{h}(Q)$, by the minimality of $\hat{h}(Q)$. By the parallelogram law,

$$\begin{aligned} 2\hat{h}(P) &= \hat{h}(Q+P) + \hat{h}(Q-P) - 2\hat{h}(Q) \\ &\geq 0 + \hat{h}(2R) - 2\hat{h}(Q) \\ &= 4\hat{h}(R) - 2\hat{h}(Q) \\ &\geq 2\hat{h}(Q) \end{aligned}$$

So $\hat{h}(Q) \leq \hat{h}(P) \leq B$ and therefore $Q \in S$, a contradiction. ☐

References

[1] J-P. Serre, *Lectures on the Mordell-Weil theorem*, 3rd edition, Springer Fachmedien Wiesbaden, 1997.

[2] J. H. Silverman, *The arithmetic of elliptic curves*, Springer, 2009.