1.17 26/07/10 - The Néron-Tate Canonical Height

Last time, we discussed the height function $h : E(\mathbb{Q}) \to [0, \infty)$ on an elliptic curve $E/\mathbb{Q}$ and used this to reduce the proof of the Mordell Theorem (1.15.8) to the weaker 1.15.9. Today, I want to spend a little bit more time talking about the theory of heights before we move on the proof of the weak Mordell Theorem.

1.17.A The Approximate Parallelogram Law for the Naive Height

For this, we will need one more result about heights. This is

Proposition 1.17.1. Let $E : y^2 = x^3 + ax^2 + bx + c$ be an elliptic curve over $\mathbb{Q}$ with $a, b, c \in \mathbb{Z}$. Then for all $P, Q \in E(\mathbb{Q})$, we have that

$$h(P+Q) + h(P-Q) = 2h(P) + 2h(Q) + O(1), \tag{1.17.2}$$

where the constant implicit the $O(1)$ depends on a, b and c but not on P, Q.

At this point, it's helpful to introduce the "big O" notation from complexity theory. The above proposition is just saying that there is a constant $C_5 \in \mathbb{R}_{\geq 0}$, depending only on a, b, c, such that for all $P, Q \in E(\mathbb{Q})$, we have that

$$|h(P+Q) + h(P-Q) - 2h(P) - 2h(Q)| \leq C_5.$$

In general, we will use this notation to mean that there is a constant (we will make to sure to emphasize what exactly it depends on) such that the quantity supposed to be $O(1)$ is bounded in absolute value by that constant.

Proof. The proof of 1.17.2 is very similar to that of 1.16.6 and is left as Exercise 2.5.5. ∎

Remark 1.17.3. What does the identity 1.17.2–except for the extra $O(1)$ term–remind you of? It should remind you of both the polynomial identity $(p+q)^2 + (p-q)^2 = 2p^2 + 2q^2$ in $\mathbb{Z}[p, q]$, and of the identity

$$\left\|\vec{P} + \vec{Q}\right\|^2 + \left\|\vec{P} - \vec{Q}\right\|^2 = 2\left\|\vec{P}\right\|^2 + 2\left\|\vec{Q}\right\|^2$$

for vectors $\vec{P}, \vec{Q}$ in an inner product space, which says that the sum of the squares of the four sides of a parallelogram is the sum of the squares of the its diagonals. For this reason, we will refer to 1.17.2 as the (approximate) parallelogram law.

Let's now derive some consequences of 1.17.1.

Corollary 1.17.4. Let $E/\mathbb{Q}$ be an elliptic curve as in 1.17.1.

(a) Fix a $Q \in E(\mathbb{Q})$. For each $P \in E(\mathbb{Q})$, we have that $h(P+Q) \leq 2h(P) + O(1)$, where the constant in $O(1)$ depends on a, b, c, and Q, but not P.
(b) Fix an $m \in \mathbb{Z}$. For each $P \in E(\mathbb{Q})$, we have that $h(mP) = m^2 h(P) + O(1)$, where the constant in $O(1)$ depends on a, b, c, and m but not P.

Proof.

(a) Follows immediately from 1.17.1, since $h(P-Q) \geq 0$.
(b) Since $h(P) = h(-P)$, it suffices to prove this for $m \in \mathbb{Z}_{\geq 0}$. We proceed by induction, with the cases $m = 0, 1$ being clear. If $m \geq 1$, then applying 1.17.2 to the pair of points (mP, P) gives us that
$$h((m+1)P) + h((m-1)P) = 2h(mP) + 2h(P) + O(1),$$
where the absolute $O(1)$ depends only on a, b, c. Now using the inductive hypothesis to write
$$h((m-1)P) = (m-1)^2 h(P) + O(1) \text{ and } h(mP) = m^2 h(P) + O(1),$$

this time with the $O(1)$s depending on a, b, c, and m but not P, and putting everything together gives us

$$h((m+1)P) = \left(2m^2 + 2 - (m-1)^2\right) + O(1)$$

as needed.

■

You might start to see the advantage of the big O notation in this proof.

Remark 1.17.5.

(a) Note that 1.17.4(a) is just a restatement of 1.16.6(b); the reason I point this out is because I want to emphasize that it is a consequence of the approximate parallelogram law (1.17.1).
(b) Similarly, the condition 1.16.6(c) is one half of the statement of 1.17.4(b). It says in particular that, in general, point doubling on elliptic curve *quadruples* the number of digits of the numerator and denominator; in particular, this applies to the Bachet formula of 1.1.6. This fulfils a promise made in 1.1.6.

1.17.B The Néron-Tate Canonical Height

The above theory says that the (naive) height function h behaves approximately like a quadratic function of the point, except for these annoying $O(1)$s floating around. It is at least theoretically interesting to be able to get an actual quadratic function out of the height function. For that: how do we get rid of the $O(1)$s? *Can* we get rid of them?

It was an insight of John Tate (1925 - 2019) that we indeed can. To motivate this result, suppose that $E/\mathbb{Q}$ is an elliptic curve as above and that $\hat{h} : E(\mathbb{Q}) \to [0, \infty)$ is a function that

(a) differs from h by a bounded amount, i.e., is such that for all $P \in E(\mathbb{Q})$, we have that $\hat{h}(P) = h(P) + O(1)$, where the constant in $O(1)$ depends only on a, b, c and not P, and
(b) is genuinely quadratic, i.e., satisfies that for all $m \in \mathbb{Z}$ and $P \in E(\mathbb{Q})$ we have that $\hat{h}(mP) = m^2\hat{h}(P)$.

Tate's observation was that these two conditions are enough to pin down such an $\hat{h}$ uniquely–if it exists. Indeed, suppose that $\hat{h}$ is such a function. Then (a) is saying that there is a $C \in \mathbb{R}_{\geq 0}$ such that for all $P \in E(\mathbb{Q})$ we have that

$$\left|\hat{h}(P) - h(P)\right| \leq C.$$

If this is true for all P, this is in particular true if we replace P by $2^N P$ for any $N \in \mathbb{Z}_{\geq 0}$. Doing this, dividing throughout by 4^N, and using property (b) then gives us that for all $P \in E(\mathbb{Q})$ and $N \in \mathbb{Z}_{\geq 0}$ we have that

$$\left|\hat{h}(P) - \frac{1}{4^N} h(2^N P)\right| \leq \frac{C}{4^N}.$$

At this point, it doesn't matter what C is at all. If such an $\hat{h}$ exists, then we may take the limit as $N \to \infty$ to conclude that

$$\hat{h}(P) = \lim_{N \to \infty} \frac{1}{4^N} h(2^N P).$$

Tate observed that this can be taken as a definition of $\hat{h}$ that does actually do what we want.

Lemma/Definition 1.17.6 (The Néron-Tate Canonical Height). Let $E : y^2 = x^3 + ax^2 + bx + c$ be an elliptic curve over $\mathbb{Q}$ with $a, b, c \in \mathbb{Z}$. For each $P \in E(\mathbb{Q})$, the limit

$$\hat{h}(P) = \lim_{N \to \infty} \frac{1}{4^N} h(2^N P) \tag{1.17.7}$$

exists. The resulting function $\hat{h} : E(\mathbb{Q}) \to [0, \infty)$ is called the *Néron-Tate canonical height on the elliptic curve* E.

Remark 1.17.8.

(a) The other person whose name is attached to the quanitity $\hat{h}$ was the French mathematician André Néron (1922-1985), who had discovered this quantity earlier in terms of a more complicated expression in terms of local height functions from arithmetic intersection theory. It was Tate was realized that Néron's definition could be made much simpler via the limiting process (1.17.7).

(b) Our definition 1.17.7 is the original one due to Tate, and is consistent with the one used by Sage and the LMFDB, although it differs from the definition in [8, Prop. VIII.9.1] by a factor of 2: Silverman's height is half of ours. This won't really matter for what follows.

Proof. It suffices to show that the sequence $4^{-N}h(2^N P)$ is Cauchy. For this, note first that 1.17.4(b) implies that there is a $C \in \mathbb{R}_{\geq 0}$ such that for all $P \in E(\mathbb{Q})$ we have that

$$|h(2P) - 4h(P)| \leq C.$$

Now suppose we are given $P \in E(\mathbb{Q})$ and integers $N \geq M \geq 0$. Then we estimate

$$\begin{aligned}\left|\frac{1}{4^N}h(2^N P) - \frac{1}{4^M}h(2^M P)\right| &= \left|\sum_{n=M}^{N-1}\left(\frac{1}{4^{n+1}}h(2^{n+1}P) - \frac{1}{4^n}h(2^n P)\right)\right|. \\ &\leq \sum_{n=M}^{N-1}\frac{1}{4^{n+1}}\left|h(2^{n+1}P) - 4h(2^n P)\right| \\ &\leq C \cdot \sum_{n=M}^{N-1}\frac{1}{4^{n+1}} \leq C \cdot \sum_{n=M}^{\infty}\frac{1}{4^{n+1}} = \frac{C}{3}\cdot\frac{1}{4^M}.\end{aligned}$$

This proves that the sequence $N \mapsto 4^{-N}h(2^N P)$ is Cauchy, and hence that $\hat{h}(P)$ exists. ∎

This canonical height has all the properties we would like it to.

Theorem 1.17.9. Let $E : y^2 = x^3 + ax^2 + bx + c$ be an elliptic curve over $\mathbb{Q}$ with $a, b, c \in \mathbb{Z}$, and let $\hat{h} : E(\mathbb{Q}) \to [0, \infty)$ be the Néron-Tate canonical height on E as defined in 1.17.6. Then

(a) For each $P \in E(\mathbb{Q})$, we have that $\hat{h}(P) = h(P) + O(1)$, where the constant implicit in the $O(1)$ depends on a, b, c but not P. In particular, for each $C_1 \in \mathbb{R}_{\geq 0}$, the set $\{P \in E(\mathbb{Q}) : \hat{h}(P) \leq C_1\}$ is finite.

(b) For each $m \in \mathbb{Z}$ and $P \in E(\mathbb{Q})$ we have that $\hat{h}(mP) = m^2\hat{h}(P)$.

The properties (a) and (b) characterize the function $\hat{h}$ uniquely. This function $\hat{h}$ further satisfies the following properties.

(c) For any $P \in E(\mathbb{Q})$ we have that $\hat{h}(P) = 0$ iff $P \in \operatorname{Tors} E(\mathbb{Q})$.

(d) For all $P, Q \in E(\mathbb{Q})$, we have that $\hat{h}(P+Q) + \hat{h}(P-Q) = 2\hat{h}(P) + 2\hat{h}(Q)$. In other words, the *exact* parallelogram law holds for $\hat{h}$.

(e) The map $\langle -, - \rangle : E(\mathbb{Q}) \times E(\mathbb{Q}) \to \mathbb{R}$ given by

$$\langle P, Q \rangle := \hat{h}(P+Q) - \hat{h}(P) - \hat{h}(Q)$$

is a bilinear pairing.

Proof.

(a) Taking $M = 0$ in the inequality used in the proof of 1.17.6, we see that for all $P \in E(\mathbb{Q})$ and $N \in \mathbb{Z}_{\geq 0}$ we have that

$$\left|\frac{1}{4^N}h(2^N P) - h(P)\right| \leq \frac{C}{3}.$$

Taking the limit as $N \to \infty$ shows us that for all $P \in E(\mathbb{Q})$ we have $\left|\hat{h}(P) - h(P)\right| \leq C/3$ as needed. The second part of (a) follows immediately from the first and 1.16.6(a).

(b) The statement of 1.17.4(b) says that for any $m \in \mathbb{Z}$, there is a constant $C' \in \mathbb{R}_{\geq 0}$ (depending only on a, b, c, m) such that for all $P \in E(\mathbb{Q})$ we have that

$$\left|h(mP) - m^2 h(P)\right| \leq C'.$$

Replacing P by $2^N P$ for $N \geq 0$ and dividing throughout by 4^N tells us that for all N we have

$$\left|\frac{1}{4^N} h(2^N(mP)) - m^2 \cdot \frac{1}{4^N} h(2^N P)\right| \leq \frac{C'}{4^N}.$$

Taking the limit as $N \to \infty$ shows that $\hat{h}(mP) = m^2 \hat{h}(P)$ as needed.

We see in this way how the limiting procedure gets rid of all of the $O(1)$s automatically. The fact that (a) and (b) characterize $\hat{h}$ uniquely was discussed already at the beginning of §1.17.B.

(c) Note first that it follows immediately from 1.17.7 that $\hat{h}(\mathscr{O}) = 0$ (no matter what the convention for $h(\mathscr{O})$ is, as long as $h(\mathscr{O}) \in [0, \infty)$!). Next, if $P \in E(\mathbb{Q})$ is a point of order $m \in \mathbb{Z}_{\geq 1}$, then we get frm (b) that

$$0 = \hat{h}(\mathscr{O}) = \hat{h}(mP) = m^2 \hat{h}(P),$$

which shows $\hat{h}(P) = 0$. Conversely, if $\hat{h}(P) = 0$, then again (b) shows that for each $m \in \mathbb{Z}_{\geq 1}$ we have $\hat{h}(mP) = 0$. Then taking $C_1 = 0$ in (a), we conclude that the set $\{mP : m \in \mathbb{Z}_{\geq 1}\}$ is finite, which can only happen if $P \in \operatorname{Tors} E(\mathbb{Q})$.

(d) This follows from 1.17.1 the same way that (b) followed from 1.17.4(b); the details are left to the reader as a(n) (easy) exercise (1.17.10).

(e) Surprisingly, this is a consequence of the parallelogram law (d); since this is fun to figure out, I leave it as Exercise 2.5.8.

■

Exercise 1.17.10. Show that 1.17.9(d) follows from 1.17.1 the same way that 1.17.9(b) follows from 1.17.4(b).

Remark 1.17.11.

(a) Note that the 1.17.9(c) was the promised relationship between torsion points and heights on the elliptic curves.

(b) The results 1.17.9(c) and (e) (see also 2.5.6) combined tell us that the pairing $\langle -, - \rangle$ of 1.17.9(e) gives rise to a nondegenerate bilinear pairing on the Mordell-Weil lattice

$$\operatorname{MW}(E/\mathbb{Q}) := E(\mathbb{Q})/\operatorname{Tors} E(\mathbb{Q}),$$

which (thanks to 1.15.8) is a finitely generated torsion-free abelian group and hence (as we shall see below) is a finitely generated free abelian group. This bilinear pairing further extends to a (positive-definite) inner product on the finite-dimensional real vector space $V := \mathbb{R} \otimes_{\mathbb{Z}} E(\mathbb{Q}) \supset \mathbb{R} \otimes_{\mathbb{Z}} \operatorname{MW}(E/\mathbb{Q})$ (see 2.5.9), known as the *Néron-Tate height pairing*. Now $\operatorname{MW}(E/\mathbb{Q})$ sits inside of the inner product space V as a full-rank lattice, and hence has a real invariant associated to it, namely its covolume. Concretely, if $r = \operatorname{rank} E = \dim_{\mathbb{R}} V$ and we pick an integral basis $P_1, \ldots, P_r$ of $\operatorname{MW}(E/\mathbb{Q})$, then this covolume is just the determinant of the matrix $[\langle P_i, P_j \rangle]_{1 \leq i,j \leq r}$, which is independent of the choice of basis (technically, this is the *square* of the covolume). This squared covolume has a name: it is called the *regulator* of the elliptic curve $E/\mathbb{Q}$, and is denoted either $\operatorname{Reg}(E)$ or $R(E/\mathbb{Q})$. This is an invariant of E which measures its arithmetic complexity. For instance, if $\operatorname{rank} E = 1$, then the regulator equals the height of *the* generator of $\operatorname{MW}(E/\mathbb{Q})$ (which is unique up sign), and hence measures how "big" the rational solutions on E are. As another example, if the rank of E is say 2, and we have one "small" generator of $\operatorname{MW}(E/\mathbb{Q})$, then the regulator dictates how big the other generator has to be. Therefore, if we can find an independent way of calculating the regulator (as can be done, conditional on the strong Birch and Swinnerton-Dyer conjecture; see [8, Conjecture C.16.5(b)]), then this gives us an estimate of how far we have to search to get generators of the Mordell-Weil group. As a final application, we will see later how the height pairing can be used to test the independence of points on an elliptic curve.

That's all I have to say about heights for now, but we will return to them after we haven seen the proof of the Mordell Theorem.