Computing a Lower Bound for the Canonical Height on Elliptic Curves over Totally Real Number Fields

Thotsaphon Thongjunthug

University of Warwick
Coventry, UK

8th Algorithmic Number Theory Symposium
Banff, Canada
19 May 2008

Outline

Elliptic Curves

Let K be a number field. An elliptic curve E over K is the set of all (x, y) satisfying the Weierstrass equation

$$E: \quad y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$$

for some $a_i \in K$, with non-zero discriminant.

For any field $L \supseteq K$, define the set of all L-points of E as

$$E(L) = \{(x, y) \in L \times L : (x, y) \in E\} \cup \{O\}$$

where O denotes the point at infinity.

The set $E(L)$ is an abelian group under "addition", with O as the identity

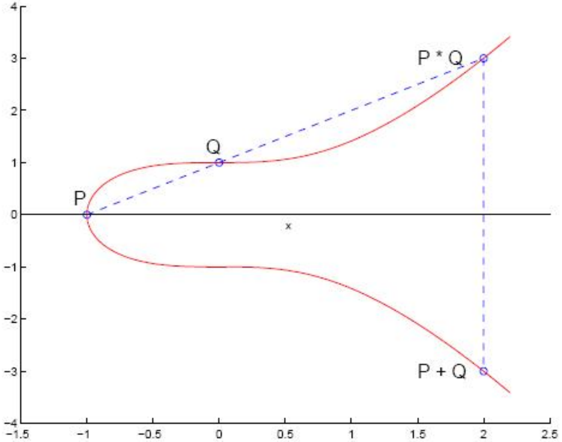
P * Q
Q
P
P + Q
x
4
3
2
1
0
−1
−2
−3
−4
−1.5
−1
−0.5
0
0.5
1
1.5
2
2.5

Moreover,

Theorem (Mordell–Weil)

Let K be a number field. The group $E(K)$ is finitely generated.

Equivalently,

$$E(K) \cong T \times \mathbb{Z}^s$$

where the torsion subgroup T of $E(K)$ is finite, and the rank s of $E(K)$ is non-negative.

Thus every point $P \in E(K)$ is a linear combination of points in the T, and a Mordell–Weil basis $\{P_1, \ldots, P_s\}$ of $E(K)$. In contrast to T, determining a Mordell–Weil basis is much harder.

The Problem

In general, the task of explicit computation of a Mordell–Weil basis consists of:

1. An m-descent (for some $m \geq 2$) is used to determine $P_1, \ldots, P_s$, a basis for $E(K)/mE(K)$.
2. A lower bound $\lambda > 0$ for the canonical height $\hat{h}(P)$ is determined. This together with the geometry of numbers yields an upper bound for the index

$$n = [E(K)/T : \langle P_1, \ldots, P_s \rangle].$$

3. A sieving procedure is used to deduce a Mordell–Weil basis.

In Step 2, we wish to have the upper bound for n as small as possible. This can be achieved if we have a larger value of λ (Siksek 1995).

In the past, a number of algorithms for computing such lower bound have been proposed. This includes:

- Hindry and Silverman (1988): Works for any number field K, model-independent, but rather theoretical.
- Cremona and Siksek (2006): Works for $K = \mathbb{Q}$. Recently known to be the sharpest one for such K.

This work is mainly a generalisation of Cremona and Siksek's algorithm. In particular, I aim to extend their algorithm to any elliptic curves over totally real number fields.

Points of Good Reduction

Suppose K is a totally real number field of degree $r = [K : \mathbb{Q}]$. Let E be an elliptic curve over K given by an integral Weierstrass model, and $\Delta = \mathrm{disc}(E)$. Define a map

$$\phi : E(K) \to \prod_{v \in S} E^{(v)}(K_v)$$

where

$$S = \{\infty_1, \ldots, \infty_r\} \cup \{\mathfrak{p} : \mathfrak{p} \mid \Delta\}$$

in such a way that ϕ maps each point $P \in E(K)$ to its corresponding point on each real embedding $E^1, \ldots, E^r$, and on each minimal model of E at $\mathfrak{p}$, denoted by $E^{(\mathfrak{p})}$.

We wish to estimate a lower bound for $\hat{h}(P)$, where $P \in E(K)$. Instead of working over $E(K)$ itself, we compute a lower bound of $\hat{h}(P)$ for

$$P \in E_{\mathrm{gr}}(K) := \phi^{-1}\left(\prod_{v \in S} E_0^{(v)}(K_v)\right)$$

where

$$E_0^{(v)}(K_v) = \begin{cases} \text{connected component of the identity} & \text{if } v = \infty_j \\ \text{set of points of good reduction} & \text{if } v = \mathfrak{p}. \end{cases}$$

In other words, $E_{\mathrm{gr}}(K)$ is the set of all points having good reduction on every $E^{(v)}(K_v)$.

The lower bound for the canonical height on the whole $E(K)$ can be easily deduced once the lower bound μ for the canonical height on $E_{\mathrm{gr}}(K)$ is determined.

Let c be the least common multiple of the Tamagawa indices

$$c_v = [E^{(v)}(K_v) : E_0^{(v)}(K_v)]$$

for every $v \in M_K$ (This is well-defined since $c_v = 1$ for almost all v). Then the lower bound for the canonical height of all non-torsion points in $E(K)$ is

$$\lambda = \mu / c^2.$$

Estimating the Local Heights

From the properties of the canonical height, we have

$$\hat{h}(P) = \frac{1}{r}\sum_{v\in M_K} n_v\lambda_v(P)$$
$$= \frac{1}{r}\left(\sum_{\mathfrak{p}} n_{\mathfrak{p}}\lambda_{\mathfrak{p}}(P) + \sum_{j=1}^{r}\lambda_{\infty_j}(P)\right).$$

Note that $n_v = [K_v : \mathbb{Q}_v] = [\mathbb{R} : \mathbb{R}] = 1$ for all $v = \infty_j$. The function $\lambda_v : E(K_v) \to \mathbb{R}$ is called the local height of P at v.

It then suffices to estimate a lower bound for each sum, in order to obtain a lower bound for $\hat{h}(P)$ on $E_{\mathrm{gr}}(K)$.

Non-Archimedean Local Heights

Let $k_{\mathfrak{p}}$ be the residue class field of $\mathfrak{p}$, with $c(\mathfrak{p}) = \mathrm{char}(k_{\mathfrak{p}})$. Also let $e_{\mathfrak{p}}$ be the exponent of the group $E_{\mathrm{ns}}^{(\mathfrak{p})}(k_{\mathfrak{p}}) \cong E_0^{(\mathfrak{p})}(K_{\mathfrak{p}})/E_1^{(\mathfrak{p})}(K_{\mathfrak{p}})$. Then

Proposition

Suppose $P \in E_{\mathrm{gr}}(K) \setminus \{O\}$. Then

$$\sum_{\mathfrak{p}} n_{\mathfrak{p}} \lambda_{\mathfrak{p}}(nP) \geq D_E(n) - \frac{1}{6} \log \mathcal{N}\left(\prod_{\mathfrak{p}} \mathfrak{p}^{\mathrm{ord}_{\mathfrak{p}}(\Delta/\Delta^{(\mathfrak{p})})} \right)$$

where $D_E(n) = \sum_{\substack{\mathfrak{p} \\ e_{\mathfrak{p}} \mid n}} 2(1 + \mathrm{ord}_{c(\mathfrak{p})}(n/e_{\mathfrak{p}})) \log \mathcal{N}(\mathfrak{p})$.

Moreover, if $e_{\mathfrak{p}} \mid n$, then $\mathcal{N}(\mathfrak{p}) \leq (n+1)^{\max\{2,[K:\mathbb{Q}]\}}$ (i.e. the sum for D_E is finite).

Archimedean Local Heights

Let

$$\alpha_j^{-3} = \inf_{P \in E_0^j(\mathbb{R})} \left\{ \frac{\max\{|f(P)|_{\infty_j}, |g(P)|_{\infty_j}\}}{\max\{1, |x(P)|_{\infty_j}\}^4} \right\}$$

where

$$\begin{aligned} f(P) &= 4x(P)^3 + b_2 x(P)^2 + 2b_4 x(P) + b_6 \\ g(P) &= x(P)^4 - b_4 x(P)^2 - 2b_6 x(P) - b_8 \end{aligned}$$

and $b_2, b_4, b_6, b_8 \in K$ are usual constants associated to E.

Lemma

If $P \in E_0^j(\mathbb{R}) \setminus \{O\}$, then

$$\lambda_{\infty_j}(P) \geq \log \max\{1, |x(P)|_{\infty_j}\} - \log \alpha_j.$$

The number α_j can be efficiently computed (Cremona, Prickett, Siksek 2006).

A Bound for Multiples of Points of Good Reduction

We wish to show whether a given $\mu > 0$ satisfies $\hat{h}(P) > \mu$ for all non-torsion $P \in E_{\mathrm{gr}}(K)$. This involves the approximation of a bound for $x(nP)$, which we derive from our previous estimate on local heights.

Let

$$B_n(\mu) = \exp\left(rn^2\mu - D_E(n) + \frac{1}{6}\log \mathcal{N}\left(\prod_{\mathfrak{p}} \mathfrak{p}^{\mathrm{ord}_{\mathfrak{p}}(\Delta/\Delta^{(\mathfrak{p})})}\right) + \sum_{j=1}^{r} \log \alpha_j \right).$$

Proposition

If $B_n(\mu) < 1$, then $\hat{h}(P) > \mu$ for all non-torsion $P \in E_{\mathrm{gr}}(K)$.
If $B_n(\mu) \geq 1$ then for all non-torsion $P \in E_{\mathrm{gr}}(K)$ with $\hat{h}(P) \leq \mu$, we have

$$|x(nP)|_{\infty_j} \leq B_n(\mu)$$

for all $j = 1, \ldots, r$.

Note that μ may still be a lower bound for $\hat{h}(P)$ even $B_n(\mu) \geq 1$. In this case, we shall prove this by solving the inequalities involving $x(nP)$ on each real embedding E^j.

Solving Inequalities on Real Embeddings

For $j = 1, \ldots, r$, the previous proposition says that every non-torsion point $P \in E_{\mathrm{gr}}(K)$ with $\hat{h}(P) \leq \mu$ must satisfy $|x(nP)|_j \leq B_n(\mu)$. This means we need to consider s elliptic curves over $\mathbb{R}$, say

$$E^j : \quad y^2 + \sigma_j(a_1)xy + \sigma_j(a_3)y = x^3 + \sigma_j(a_2)x^2 + \sigma_j(a_4)x + \sigma_j(a_6)$$

where $\sigma_j : K \to \mathbb{R}$ are the real embeddings of K. In particular, we need to consider the system of inequalities involving $x(\sigma_j(nP))$ on each $E_0^j(\mathbb{R})$.

To do this, we use an application of elliptic logarithm, which is an isomorphism

$$\varphi : E_0(\mathbb{R}) \to \mathbb{R}/\mathbb{Z} \cong [0, 1).$$

To solve the inequalities, first we fix a real embedding E^j at a time. Let $\varphi_j : E_0^j(\mathbb{R}) \to [0,1)$ be the corresponding elliptic logarithm map.

Suppose $P \in E_0^j(\mathbb{R})$ such that $|x(nP)| \leq B_n(\mu)$ for every $n > 0$. Then we have $\varphi_j(nP) \in \mathcal{S}^j(-B_n(\mu), B_n(\mu))$ where $\mathcal{S}^j : \mathbb{R} \times \mathbb{R} \to [0,1)$ yields a subinterval of $[0,1)$.

Since φ_j is an isomorphism, we have $\varphi_j(nP) = n\varphi_j(P) \pmod 1$. Hence

$$\varphi_j(P) \in \mathcal{S}_n^j(-B_n(\mu), B_n(\mu))$$

for every n, where

$$S_n^j(-B_n(\mu), B_n(\mu)) = \bigcup_{t=0}^{n-1} \left(\frac{t}{n} + \frac{1}{n}\mathcal{S}^j(-B_n(\mu), B_n(\mu)) \right).$$

The Algorithm

To check if $\mu > 0$ is a lower bound for $\hat{h}(P)$ on $E_{\mathrm{gr}}(K)$:

1. Start with a given initial guess $\mu > 0$ and $k \in \mathbb{Z}^+$.
2. For $n = 1, \ldots, k$, compute $B_n(\mu)$.
3. If $B_n(\mu) < 1$ for some n, then $\hat{h}(P) > \mu$ for all non-torsion $P \in E_{\mathrm{gr}}(K) \implies$ Done.
4. Otherwise, choose a real embedding E^j. Compute $\bigcap_{n=1}^{k} \mathcal{S}_n^j(-B_n(\mu), B_n(\mu))$.
5. If the intersection is empty, we conclude that $\hat{h}(P) > \mu$ for all non-torsion $P \in E_{\mathrm{gr}}(K) \implies$ Done.
6. If not, repeat (4)–(6) with a different E^j.

If for all E^j the intersections are not empty, we fail to show that μ is a lower bound for $\hat{h}(P)$.

Example I

Let E be the elliptic curve over $K = \mathbb{Q}(\sqrt{10})$ given by

$$E : \quad y^2 = f(x) = x^3 + 125.$$

Note that K has class number 2. The decomposition of the discriminant Δ of E is $\langle \Delta \rangle = \mathfrak{p}_1^{12}\mathfrak{p}_2^3\mathfrak{p}_3^3\mathfrak{p}_4^8$, where

$\mathfrak{p}_1 = \langle 5, \sqrt{10} \rangle$, $\mathfrak{p}_2 = \langle 3, 4+\sqrt{10} \rangle$, $\mathfrak{p}_3 = \langle 3, 2+\sqrt{10} \rangle$, $\mathfrak{p}_4 = \langle 2, \sqrt{10} \rangle$.

Indeed E is minimal everywhere except at $\mathfrak{p}_1$.
By substituting

$$x = (\sqrt{10})^2 x', \quad y = (\sqrt{10})^3 y'$$

we have a new elliptic curve $E' : {y'}^2 = {x'}^3 + 1/8$.

Hence E' is minimal at $\mathfrak{p}_1$ and elsewhere, except at all prime ideals dividing 2. Thus we let $E^{(\mathfrak{p}_1)} = E'$ and $E^{(\mathfrak{p})} = E$ for any $\mathfrak{p} \neq \mathfrak{p}_1$. Our program shows that

$$\hat{h}(P) > 0.2859$$

for every non-torsion $P \in E_{\mathrm{gr}}(K)$.

The Tamagawa indices at $\mathfrak{p}_1, \mathfrak{p}_2, \mathfrak{p}_3, \mathfrak{p}_4$ are 1, 2, 2, and 1 respectively. Also since $c_{\infty_1} = c_{\infty_2} = 1$, then $c = 2$. Hence for any non-torsion point $P \in E(K)$, we have

$$\hat{h}(P) > 0.2859/(2^2) = 0.0714.$$

Observe that the point $P = (5, 5\sqrt{10}) \in E(K)$ is non-torsion. Assume $E(K)$ has rank 1. Then by Siksek's theorem, we have

$$n = [E(K) : \langle P \rangle] \leq 3.0229.$$

Example II

Let E be the elliptic curve over $K = \mathbb{Q}(\sqrt{7})$ given by

$$E : y^2 + (3 + 3\sqrt{7})xy + y = x^3 + (26 + 4\sqrt{7})x^2 + x$$

By computing the discriminant Δ of E, we have $\langle \Delta \rangle = \mathfrak{p}_1\mathfrak{p}_2\mathfrak{p}_3$, where

$$\mathfrak{p}_1 = \langle 4219, 1083 + \sqrt{7} \rangle, \quad \mathfrak{p}_2 = \langle 4657, 35443 + \sqrt{7} \rangle,$$
$$\mathfrak{p}_3 = \langle 12799, 5358 + \sqrt{7} \rangle.$$

Thus E is already a globally minimal model.
The algorithm shows that

$$\hat{h}(P) > 0.1415$$

for every non-torsion point $P \in E_{\mathrm{gr}}(K)$.

The Tamagawa indices at $\mathfrak{p}_1, \mathfrak{p}_2, \mathfrak{p}_3$ are all 1. In addition, $c_{\infty_1} = c_{\infty_2} = 2$. Hence $c = 2$. This gives us

$$\hat{h}(P) > 0.1415/2^2 = 0.0353$$

for all non-torsion points $P \in E(K)$.

Finally, let

$$P_1 = (0,0), \quad P_2 = (1, \sqrt{7}).$$

Then $P_1, P_2 \in E(K)$ and are non-torsion. Assume that $E(K)$ has rank 2, then by Siksek's theorem we have

$$n = [E(K) : \langle P_1, P_2 \rangle] \leq 35.2450.$$