1.19 26/07/14 - Weak Mordell for Split Two Torsion

The first order of business today is to wrap up the proof of the weak Mordell theorem (1.15.9) in the case of full rational two-torsion.

Recall the set-up: $E/\mathbb{Q}$ is an elliptic curve in short Weierstrass form $y^2 = f(x)$, where $f(x) \in \mathbb{Z}[x]$ is a monic cubic polynomial that splits completely into linear factors over $\mathbb{Q}$ as $f(x) = \prod_{i=1}^{3}(x - e_i)$ for pairwise distinct $e_i \in \mathbb{Z}$. Again, we let $T_i = (e_i, 0)$ for $i = 1, 2, 3$ be the rational two-torsion points, and note that the elliptic discriminant is $\Delta = 16\prod_{1\leq i<j\leq 3}(e_i - e_j)^2$. Finally, note that T_i is the unique point on $E(\mathbb{Q})$ with x-coordinate e_i.

Continuing the discussion from last time (§1.18.B), we let $L := \mathbb{Q}[\beta] := \mathbb{Q}[x]/(f(x))$. As noted before, L is *not* a field; rather, it is a split algebra in the sense that the Chinese Remainder Theorem tells us that there is a $\mathbb{Q}$-algebra isomorphism

$$L \xrightarrow{\sim} \mathbb{Q} \times \mathbb{Q} \times \mathbb{Q}, \quad g(\beta) \mapsto (g(e_1), g(e_2), g(e_3)),$$

with the inverse map given by taking a triple (a_1, a_2, a_3) of rational numbers to the class $g(\beta)$ of the unique quadratic polynomial $g(x)$ satisfying $g(e_i) = a_i$ for $i = 1, 2, 3$; an explicit formula for g can be given by Lagrange interpolation, but we won't need it.[6]

From the above isomorphism we get also that $L^\times/(L^\times)^2 \xrightarrow{\sim} \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 \times \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 \times \mathbb{Q}^\times/(\mathbb{Q}^\times)^2$, and as before we consider the map

$$\alpha : E(\mathbb{Q}) \to L^\times/(L^\times)^2 \xrightarrow{\sim} \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 \times \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 \times \mathbb{Q}^\times/(\mathbb{Q}^\times)^2$$

given by taking $\mathscr{O}$ to $[1] \leftrightarrow ([1], [1], [1])$ and $P = (x, y) \in E(\mathbb{Q}) \smallsetminus E[2](\mathbb{Q})$ to $[x - \beta] \leftrightarrow ([x - e_1], [x - e_2], [x - e_3])$. We need to figure out what the right values of $\alpha(T_i)$ for $i = 1, 2, 3$ should be to make the same theorem as last time work in this setting. We will see what they need to be in

Theorem 1.19.1. In the above setup, consider the map

$$\alpha : E(\mathbb{Q}) \to \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 \times \mathbb{Q}^\times/(\mathbb{Q}^\times)^2 \times \mathbb{Q}^\times/(\mathbb{Q}^\times)^2$$

given by

$$P \mapsto \begin{cases} ([1], [1], [1]), & \text{if } P = \mathscr{O}, \\ ([x - e_1], [x - e_2], [x - e_3]), & \text{if } P = (x, y) \in E(\mathbb{Q}) \smallsetminus E[2](\mathbb{Q}), \\ ([(e_1 - e_2)(e_1 - e_3)], [e_1 - e_2], [e_1 - e_3]), & \text{if } P = T_1, \\ ([e_2 - e_1], [(e_2 - e_1)(e_2 - e_3)], [e_2 - e_3]), & \text{if } P = T_2, \text{ and} \\ ([e_3 - e_1], [e_3 - e_2], [(e_3 - e_1)(e_3 - e_2)]), & \text{if } P = T_3. \end{cases}$$

(a) The map α is a group homomorphism.
(b) We have that $\ker \alpha = 2E(\mathbb{Q})$.
(c) The image $\operatorname{Im} \alpha$ is finite.

We will see in the proof below where the right choices of the values of the $\alpha(T_i)$ for $i = 1, 2, 3$ come from.

Proof.

(a) Of course, we invoke 1.15.5. As in the proof of 1.18.1, we are reduced to showing that if $P_1, P_2, P_3 \in E(\mathbb{Q}) \smallsetminus \{\mathscr{O}\}$ are collinear (as usual with multiplicities), then $\alpha(P_1) \cdot \alpha(P_2) \cdot \alpha(P_3) = ([1], [1], [1])$. First suppose that none of the P_i are two-torsion points, which is the same as saying that the x-coordinates x_i of the P_i satisfy $x_i \neq e_j$ for $i, j = 1, 2, 3$. The line ℓ joining the P_i is not vertical, and hence has equation $\ell : y = \lambda x + \nu$ for some $\lambda, \nu \in \mathbb{Q}$. Then x_i are the roots of the polynomial

$$g(x) := f(x) - (\lambda x + \nu)^2.$$

[6] This is the reason the Lagrange interpolating polynomial shows up in the proof of this result found in [13, §2.8].

We wish to show for $j = 1, 2, 3$ that the product $\prod_{i=1}^{3}(x_i - e_j)$ is a nonzero square, but indeed this product is

$$\prod_{i=1}^{3}(x_i - e_j) = -g(e_j) = (\lambda e_j + \nu)^2.$$

Since $T_j \notin \ell(\mathbb{Q})$, we have that $\lambda e_j + \nu \neq 0$, and so we are done.
Suppose now that some of the P_i coincide with the T_js. If all three did, then we must have (up to permutation) that $P_i = T_i$ for $i = 1, 2, 3$. This gives us the condition $\alpha(T_1)\alpha(T_2)\alpha(T_3) = ([1], [1], [1])$, and this is already enough to tell us what the "patches" should be; this is where the values of the $\alpha(T_i)$ for $i = 1, 2, 3$ come from. Let's now check that this works in the remaining cases too.
If exactly two of the P_is were T_js, say P_1 and P_2, then we could not have $P_1 = P_2$, since this would force $P_3 = \mathscr{O}$. Therefore, P_1 and P_2 must be distinct; up to relabelling, this again forces $P_i = T_i$ for $i = 1, 2, 3$, and we are back in the case above.
It only remains to deal with the case in which exactly one of the P_is coincides with the T_js, say $P_1 = T_1$ and $P_2, P_3 \in E(\mathbb{Q}) \smallsetminus E[2](\mathbb{Q})$, which means again that $x_i \neq e_j$ for $i = 2, 3$ and $j = 1, 2, 3$. In this case, if the line $L : y = \lambda x + \nu$ joining the P_i were to be horizontal (i.e., $\lambda = 0$), then $T_1 \in L(\mathbb{Q})$ would force $\nu = 0$ as well, and again we'd have $P_i = T_i$ for $i = 1, 2, 3$ (up to relabelling); since we are assuming we are not in this case, we must have that $\lambda \neq 0$. The fact that $T_1 \in \ell(\mathbb{Q})$ implies that $\nu = -\lambda e_1$. We conclude from this that the polynomial $g(x) = f(x) - (\lambda x + \nu)^2$ factors as $g(x) = (x - e_1)h(x)$, where

$$h(x) := (x - e_2)(x - e_3) - \lambda^2(x - e_1) = (x - x_2)(x - x_3).$$

Now let us compute the product $\alpha(P_1)\alpha(P_2)\alpha(P_3)$. In the second factor, this is precisely (the class of)

$$(e_1 - e_2)(x_2 - e_2)(x_3 - e_2) = (e_1 - e_2) \cdot h(e_2) = \lambda^2(e_1 - e_2)^2,$$

which is indeed a nonzero square. Similarly, the last factor is $\lambda^2(e_1 - e_3)^2$. Finally, the first factor is exactly

$$(e_1 - e_2)(e_1 - e_3) \cdot (x_2 - e_1)(x_3 - e_1) = h(e_1)^2.$$

Again, we know that $h(e_1) \neq 0$, and we are done.

(b) This part of the proof is not just similar to that of 1.18.1(b); it is identical. Indeed, the point is that the proof there did not really use that L is a number field. Let's breeze through the proof again. As before, $\ker\alpha \subset 2E(\mathbb{Q})$ is clear. If $P_1 = (x_1, y_1) \in \ker\alpha \smallsetminus \{\mathscr{O}\}$, then there are $u, v, w \in \mathbb{Q}$ with $x_1 - \beta = (u\beta^2 + v\beta + w)^2$ and $u \neq 0$. As before, we let $t := (v/u) - a$, where $a = -(e_1 + e_2 + e_3)$, and we let λ, ν be determined by $(t - \beta)(u\beta^2 + \nu\beta + w) = \lambda\beta + \nu$. The polynomial $(x - x_1)(x - t)^2 + (\lambda x + \nu)^2$ is a monic cubic polynomial satisfied by $\beta \in L$, but this implies that this must be equal to $f(x)$ by properties of L. The rest of the proof is identical.[7]

(c) This part of the proof is even easier than last time. Again, the goal is to find a finite set S of (rational!) primes such that if $\mathbb{Q}(S, 2) \subset \mathbb{Q}^\times/(\mathbb{Q}^\times)^2$ is the subgroup generated by -1 and the $p \in S$, then $\operatorname{Im}\alpha \subset \mathbb{Q}(S, 2) \times \mathbb{Q}(S, 2) \times \mathbb{Q}(S, 2)$. We claim that we can take S to be the set of primes dividing Δ.
Since this is clear for $\alpha(T)$ for $T \in E[2](\mathbb{Q})$, we only need to show that if $P = (x, y) \in E(\mathbb{Q}) \smallsetminus E[2](\mathbb{Q})$ and p is a prime such that for some j we have $v_p(x - e_j) \equiv 1 \pmod 2$, then $p \mid \Delta$. Without loss of generality, suppose that $v_p(x - e_1) \equiv 1 \pmod 2$. Then the equation $y^2 = \prod_{j=1}^{3}(x - e_j)$ forces that $v_p(x - e_j) \equiv 1 \pmod 2$ for at least one of $j = 2, 3$. Without loss of generality, suppose this true for $j = 2$.
As in the proof of 1.18.1, we now clear denominators using 1.14.8 to write $P = (m/e^2, n/e^3)$ for some $m, n, e \in \mathbb{Z}$ with $e \geq 1$ and $\gcd(m, e) = \gcd(n, e) = 1$. Then $P \in E(\mathbb{Q})$ translates to

$$n^2 = (m - e^2e_1)(m - e^2e_2)(m - e^2e_3).$$

The condition $v_p(m - e^2e_j) \equiv 1 \pmod 2$ for $j = 1, 2$ now implies that $p \mid m - e^2e_j$ for $j = 1, 2$, which forces that $p \mid e^2(e_1 - e_2)$. But again $p \nmid e$ because $p \mid m - e^2e_1$ and $\gcd(m, e) = 1$. Therefore, we conclude that $p \mid e_1 - e_2 \mid \Delta$ as needed.

■

[7] If you wish, you can phrase this proof without using properties of L and rather using the Lagrange interpolating polynomial to translate directly between $E(\mathbb{Q})$ and $(\mathbb{Q}^\times/(\mathbb{Q}^\times)^2)^3$ without referring to L at all, as is done in [13, §2.8].

This finishes the proof of 1.19.1 and hence 1.15.9 and 1.15.8 in the case of full rational two-torsion.

Remark 1.19.2.

(a) Note that in the proof of 1.19.1 the image of α is contained in the "hyperplane" in $(\mathbb{Q}^\times/(\mathbb{Q}^\times)^2)^3$ consisting of triples for which the product of the coordinates is trivial (i.e., a square). This allows us to express α (up to non-canonical choices) as a map to $(\mathbb{Q}^\times/(\mathbb{Q}^\times)^2)^2$ instead (as is done in [8, §X.1]). I don't like this formulation because it hides the symmetry between the e_i evident in the statement in 1.19.1.

(b) One might rightfully ask where the maps α in the 1.18.1 and 1.19.1 *really* come from. The answer to this question is that they come from Galois cohomology. This is beyond the scope of this course, but the interested reader can find more about this in [8, Chapters VIII, X].

(c) Technically, to finish the proof of 1.15.9 and hence 1.15.8, we still need to deal with one more case: the case in which $f(x)$ has exactly one rational root. This can be done very similarly to 1.18.1 above (as is done in [1, §15]); I leave this to the interested reader. We will follow a different approach–namely that of descent via two-isogeny, as in [1, §14] and [2, §3.4-3.5]–to give a proof of this case (and a second proof of the case of full two-torsion) next time.

(d) To what extent does the proof of 1.19.1 (and even 1.18.1) rely on the fact that our base field is $\mathbb{Q}$? We only really needed this fact in the statement (c) (of both theorems). Therefore, (i) the proofs are basically correct as stated for any number field K in place of $\mathbb{Q}$, and (ii) we have proven something quite general in parts (a) and (b). For instance, we have shown already the result of 1.19.3.

Corollary 1.19.3. Let K be a field of characteristic not two, $e_1, e_2, e_3 \in K$ be pairwise distinct, and let E/K be the elliptic curve with Weierstrass equation $y^2 = \prod_{i=1}^3 (x - e_i)$.

(a) A point $P = (x, y) \in E(K) \smallsetminus E[2](K)$ lies in $2E(K)$ iff each of $x - e_1, x - e_2, x - e_3$ are squares in K.

(b) The point $T_1 = (e_1, 0)$ lies in $2E(K)$ iff $e_1 - e_2$ and $e_1 - e_3$ are squares in K. Similar results hold for the points $T_i = (e_i, 0)$ for $i = 2, 3$.

Proof. This follows from the proof of parts (a) and (b) of 1.19.1 above. ∎

Corollary 1.19.4. Let K be a field of characteristic other than 2 such that -1 is not a square in K (e.g., $K = \mathbb{Q}$ or $K = \mathbb{F}_p$ for p a prime with $p \equiv 3 \pmod 4$). If E/K is an elliptic curve, then $E[4](K) \ncong_{\mathsf{Ab}} \mathbb{Z}/4 \oplus \mathbb{Z}/4$, i.e., it is not possible for all the 4-torsion on E to be defined over K.

Proof. Suppose that K is a field of characteristic other than two, and E/K is an elliptic curve such that $E[4](K) \cong_{\mathsf{Ab}} \mathbb{Z}/4 \oplus \mathbb{Z}/4$. In particular, $E[2](K) \cong_{\mathsf{Ab}} \mathbb{Z}/2 \oplus \mathbb{Z}/2$, and hence E has fully split two-torsion over K. This means that E can be put into Weierstrass form $y^2 = f(x)$ where $f(x) = (x - e_1)(x - e_2)(x - e_3)$ for some pairwise distinct $e_i \in K$. Since $E[4](K) \cong_{\mathsf{Ab}} \mathbb{Z}/4 \oplus \mathbb{Z}/4$, at least two of the points $T_i = (e_i, 0)$ for $i = 1, 2, 3$ belong to $2E(K)$; without loss of generality, suppose these are T_1 and T_2. Since $T_1 \in 2E(K)$, 1.19.3(b) tells us that $e_1 - e_2$ is a square in K. Since $T_2 \in 2E(K)$, 1.19.3(b) tells us that $e_2 - e_1$ is a square in K. Since $e_1 \neq e_2$, these facts combined tell us that -1 is a square in K. ∎

Unimportant Remark 1.19.5. The argument of 1.19.4 can be generalized to prove that if K is a field, E/K an elliptic curve, and $n \in \mathbb{Z}_{\geq 1}$ such that $E[n](K) \cong_{\mathsf{Ab}} \mathbb{Z}/n \oplus \mathbb{Z}/n$, then in fact K contains a primitive n^{th} root of unity. This shows in particular that if $C \subset \mathbb{P}^2_{\mathbb{R}}$ is a smooth cubic curve, then not all of its 9 complex inflection points can be real (which can also be proven by other means; see 2.2.9(e) and 2.4.12(b)). The cleanest modern approach to proving this involves using the Weil pairing. We will not explain what that is, but the interested reader can find this topic in [8, §III.8].