2-DESCENT THROUGH THE AGES

Sir Peter Swinnerton-Dyer

The main object of this note, which expands an expository lecture given at the conference, is to provide the reader with an account of the process of 2-descent on elliptic curves defined over $\mathbf{Q}$ which have the form

$$\Gamma : y^2 = (x - c_1)(x - c_2)(x - c_3)$$

— that is, elliptic curves all of whose 2-division points are rational. I have also included a description of the algorithm of Cassels [1] for 4-descents. My intention is to provide the reader with the tools which he may need for applications, in a way which requires minimal effort on his part. I have therefore not included proofs, except in Appendix 2 which contains a proof/algorithm the full details of which may be needed for some applications. Instead, I have provided the necessary references.

This note describes the processes over $\mathbf{Q}$. But the statements of the theory over an arbitrary algebraic number field are not very different, except that the analogues of certain explicit results relating to the prime 2 are not known. On the other hand, some of the proofs are much harder.

We can clearly take the c_i to be integers. Let $\mathcal{B}$, the set of bad primes, be any finite set of primes containing 2, ∞ and all the odd primes dividing $(c_1 - c_2)(c_1 - c_3)(c_2 - c_3)$; thus $\mathcal{B}$ contains the primes of bad reduction for Γ. If $\mathcal{B}$ also contains some primes of good reduction, that is harmless.

The basic version of 2-descent, which goes back to Fermat, is as follows. (Good places to find proofs of the results that follow are Silverman [5] or Husemöller [3].) To any rational point (x, y) on Γ there correspond rational m_1, m_2, m_3 with $m_1 m_2 m_3 = m^2 \neq 0$ such that the three equations

$$m_i y_i^2 = x - c_i \quad \text{for} \quad i = 1, 2, 3 \tag{1}$$

are simultaneously soluble. We can multiply the m_i by non-zero squares, so that for example we can require them to be square-free integers; indeed one should really think of them as elements of $\mathbf{Q}^*/\mathbf{Q}^{*2}$, with a suitable interpretation of the equations which involve them. Denote by $\mathcal{C}(\mathbf{m})$ the curve given by the three equations (1), where $\mathbf{m} = (m_1, m_2, m_3)$. Looking for solutions of Γ is the same as looking for quadruples x, y_1, y_2, y_3 which satisfy (1) for some $\mathbf{m}$. For this purpose we need only consider the finitely many $\mathbf{m}$ for which the m_i are units at all primes outside $\mathcal{B}$; for if any m_i is divisible to an odd power by some prime p not in $\mathcal{B}$ then Γ is already insoluble in $\mathbf{Q}_p$.

One question of interest is the effect of *twisting* on the arithmetic properties of the curve Γ. If b is a nonzero rational, the twist of Γ by b is defined to be the curve

$$\Gamma_b : y^2 = (x - bc_1)(x - bc_2)(x - bc_3),$$

where we can regard b as an element of $\mathbf{Q}^*/\mathbf{Q}^{*2}$. The curve Γ_b is often written in the alternative form

$$v^2 = b(u - c_1)(u - c_2)(u - c_3).$$

The analogue of (1) for Γ_b is

$$m_i y_i^2 = x - bc_i \quad \text{for} \quad i = 1, 2, 3;$$

we shall call the curve given by these three equations $\mathcal{C}_b(\mathbf{m})$. It is often natural to compare $\mathcal{C}(\mathbf{m})$ and $\mathcal{C}_b(\mathbf{m})$ for the same $\mathbf{m}$.

Provided one treats the m_i as elements of $\mathbf{Q}^*/\mathbf{Q}^{*2}$, the triples $\mathbf{m}$ form an abelian group under componentwise multiplication:

$$\mathbf{m}' \times \mathbf{m}'' \mapsto \mathbf{m}'\mathbf{m}'' = (m_1' m_1'', m_2' m_2'', m_3' m_3'').$$

The $\mathbf{m}$ for which $\mathcal{C}(\mathbf{m})$ is everywhere locally soluble form a finite subgroup, called the 2-*Selmer group*. This is computable, and it contains the group of those $\mathbf{m}$ for which $\mathcal{C}(\mathbf{m})$ is actually soluble in $\mathbf{Q}$. This smaller group is $\Gamma(\mathbf{Q})/2\Gamma(\mathbf{Q})$, where $\Gamma(\mathbf{Q})$, the group of rational points on Γ, is the *Mordell-Weil* group of Γ. The quotient of the 2-Selmer group by this smaller group is ${}_2\text{Ш}$, the group of those elements of the *Tate-Safarevic group* which are killed by 2. One of the key conjectures in the subject is that the order of ${}_2\text{Ш}$ is a square.

The process of going from the curve Γ to the set of curves $\mathcal{C}(\mathbf{m})$, or the finite subset which is the 2-Selmer group, is called a 2-*descent*, or sometimes a *first descent*, and the curves $\mathcal{C}(\mathbf{m})$ themselves are called 2-*coverings*. The reason for this terminology is that there is a commutative diagram

$$\begin{array}{ccc} \Gamma & \longrightarrow & \Gamma \\ \| & \nearrow & \\ \mathcal{C}(\mathbf{m}) & & \end{array} \tag{2}$$

in which the left hand map is biregular (but defined over $\mathbf{C}$ rather than $\mathbf{Q}$), the top map is multiplication by 2 and the diagonal map is given by $y = my_1y_2y_3$. A 2-covering which is everywhere locally soluble, and therefore in the 2-Selmer group, can also be written in the form

$$\eta^2 = f(\xi) \quad \text{where} \quad f(\xi) = a\xi^4 + b\xi^3 + c\xi^2 + d\xi + e,$$

and many 2-coverings do arise in this way; but a 2-covering which is not in the 2-Selmer group cannot always be put into this form.

We now put this process into more modern language. In what follows, italic capitals will always denote vector spaces over $\mathbf{F}_2$, the finite field of two elements, and each of p and q will be either a finite prime or ∞. Write

$$Y_p = \mathbf{Q}_p^*/\mathbf{Q}_p^{*2}, \quad Y_{\mathcal{B}} = \oplus_{p\in\mathcal{B}} Y_p.$$

Let V_p denote the vector space of all triples (μ_1, μ_2, μ_3) with each μ_i in Y_p and $\mu_1\mu_2\mu_3 = 1$; and write $V_{\mathcal{B}} = \oplus_{p\in\mathcal{B}} V_p$. This is the best way to introduce these spaces, because it preserves symmetry; but the reader should note that the prevailing custom in the literature is to define V_p as $Y_p \times Y_p$, which is isomorphic to the V_p defined above but not in a canonical way. Next, write $X_{\mathcal{B}} = \mathfrak{o}_{\mathcal{B}}^*/\mathfrak{o}_{\mathcal{B}}^{*2}$ where $\mathfrak{o}_{\mathcal{B}}^*$ is the group of nonzero rationals which are units outside $\mathcal{B}$; and let $U_{\mathcal{B}}$ be the image in $V_{\mathcal{B}}$ of the group of triples (m_1, m_2, m_3) such that the m_i are in $X_{\mathcal{B}}$ and $m_1m_2m_3 = 1$. It is known that the map $X_{\mathcal{B}} \to Y_{\mathcal{B}}$ is an embedding and $\dim U_{\mathcal{B}} = \frac{1}{2}\dim V_{\mathcal{B}}$; both these depend on the requirement that $\mathcal{B}$ contains 2 and ∞. Finally, if (x, y) is a point of Γ defined over $\mathbf{Q}_p$ other than a 2-division point then the product of the three components in the triple $(x - c_1, x - c_2, x - c_3)$ is y^2 which is in $\mathbf{Q}_p^{*2}$; so this triple has a natural image in V_p. We can supply the images of the 2-division points by continuity; for example the image of $(c_1, 0)$ is

$$((c_1 - c_2)(c_1 - c_3), c_1 - c_2, c_1 - c_3), \tag{3}$$

and the image of the point at infinity is the trivial triple $(1, 1, 1)$, which is also the product of the three triples like (3). Thus we obtain a map $\Gamma(\mathbf{Q}_p) \to V_p$. This map, which is called the *Kummer map*, is a homomorphism. We denote its image by W_p; clearly W_p is the set of those triples $\mathbf{m}$ for which (1) is soluble in $\mathbf{Q}_p$. It is sometimes useful to have explicit descriptions of the W_p, so these are given in Appendix 1. The 2-Selmer group of Γ can now be identified with $U_{\mathcal{B}} \cap W_{\mathcal{B}}$ where $W_{\mathcal{B}} = \oplus_{p\in\mathcal{B}} W_p$; for as was noted above, (1) is soluble at every prime outside $\mathcal{B}$ if and only if the elements of $\mathbf{m}$ are in $X_{\mathcal{B}}$.

Over the years, many people must have noticed that

$$\dim W_{\mathcal{B}} = \dim U_{\mathcal{B}} = \tfrac{1}{2} \dim V_{\mathcal{B}}. \tag{4}$$

The next major step, which explains and may well have been inspired by this relation, was taken by Tate. He introduced the bilinear form e_p on $V_p \times V_p$, defined by

$$e_p(\mathbf{m}', \mathbf{m}'') = (m_1', m_1'')_p (m_2', m_2'')_p (m_3', m_3'')_p.$$

Here $(u, v)_p$ is the multiplicative Hilbert symbol with values in $\{\pm 1\}$, defined by

$$(u, v)_p = \begin{cases} 1 & \text{if } ux^2 + vy^2 = 1 \text{ is soluble in } \mathbf{Q}_p, \\ -1 & \text{otherwise.} \end{cases}$$

The Hilbert symbol is symmetric and multiplicative in each argument:

$$(u, v)_p = (v, u)_p \quad \text{and} \quad (u_1 u_2, v)_p = (u_1, v)_p (u_2, v)_p.$$

Effectively it is a replacement for the quadratic residue symbol, with the advantage that it treats the primes 2 and ∞ in just the same way as any other prime. Its other key property is the Hilbert product formula

$$\prod_p (u, v)_p = 1,$$

where the product is taken over all p including ∞; the left hand side is meaningful because $(u, v)_p = 1$ whenever p is an odd prime at which u and v are units.

The bilinear form e_p is non-degenerate and alternating on $V_p \times V_p$; we use it to define $e_{\mathcal{B}} = \prod_{p \in \mathcal{B}} e_p$, which is a non-degenerate alternating bilinear form on $V_{\mathcal{B}} \times V_{\mathcal{B}}$. (For a bilinear form with values in $\{\pm 1\}$, "symmetric" and "skew-symmetric" are the same and they each mean that $e(\mathbf{m}', \mathbf{m}'') = e(\mathbf{m}'', \mathbf{m}')$; "alternating" means that also $e(\mathbf{m}, \mathbf{m}) = 1$.) It is known from class field theory that $U_{\mathcal{B}}$ is a maximal isotropic subspace of $V_{\mathcal{B}}$. Tate showed that W_p is a maximal isotropic subspace of V_p, and therefore $W_{\mathcal{B}}$ is a maximal isotropic subspace of $V_{\mathcal{B}}$. (The proof of this, which is difficult, can be found in Milne [4].) This explains (4); and it also shows that the 2-Selmer group of Γ can be identified with both the left and the right kernel of the restriction of $e_{\mathcal{B}}$ to $U_{\mathcal{B}} \times W_{\mathcal{B}}$.

For both aesthetic and practical reasons, one would like to show that this restriction is symmetric or skew-symmetric — these two properties being the same. But to make such a statement meaningful we need an isomorphism between $U_{\mathcal{B}}$ and $W_{\mathcal{B}}$; and though they have the same structure as vector spaces it is not obvious that there is a natural isomorphism between them. The way round this obstacle was first shown in [2]. It requires the construction inside each V_p of a maximal isotropic subspace K_p such that $V_{\mathcal{B}} = U_{\mathcal{B}} \oplus K_{\mathcal{B}}$ where $K_{\mathcal{B}} = \oplus_{p \in \mathcal{B}} K_p$. Assuming that such spaces K_p can be constructed, let $t_{\mathcal{B}} :\ V_{\mathcal{B}} \to U_{\mathcal{B}}$ be the projection along $K_{\mathcal{B}}$ and write

$$U'_{\mathcal{B}} = U_{\mathcal{B}} \cap (W_{\mathcal{B}} + K_{\mathcal{B}}), \quad W'_{\mathcal{B}} = W_{\mathcal{B}}/(W_{\mathcal{B}} \cap K_{\mathcal{B}}) = \bigoplus\nolimits_{p \in \mathcal{B}} W'_p$$

where $W'_p = W_p/(W_p \cap K_p)$. The map $t_{\mathcal{B}}$ induces an isomorphism

$$\tau_{\mathcal{B}} :\ W'_{\mathcal{B}} \to U'_{\mathcal{B}},$$

and the bilinear function $e_{\mathcal{B}}$ induces a bilinear function

$$e'_{\mathcal{B}} :\ U'_{\mathcal{B}} \times W'_{\mathcal{B}} \to \{\pm 1\}.$$

The bilinear functions $U'_{\mathcal{B}} \times U'_{\mathcal{B}} \to \{\pm 1\}$ and $W'_{\mathcal{B}} \times W'_{\mathcal{B}} \to \{\pm 1\}$ defined respectively by

$$\theta^{\flat}_{\mathcal{B}} : u'_1 \times u'_2 \mapsto e'_{\mathcal{B}}(u'_1, \tau_{\mathcal{B}}^{-1}(u'_2)) \quad \text{and} \quad \theta^{\sharp}_{\mathcal{B}} : w'_1 \times w'_2 \mapsto e'_{\mathcal{B}}(\tau_{\mathcal{B}} w'_1, w'_2) \qquad (5)$$

are symmetric. (For the proof, see [2] or [8].) Here the images of $w'_1 \times w'_2$ under the second map and of $\tau_{\mathcal{B}} w'_1 \times \tau_{\mathcal{B}} w'_2$ under the first map are the same. The 2-Selmer group of Γ is isomorphic to both the left and the right kernel of $e'_{\mathcal{B}}$, and hence also to the kernels of the two maps (5).

There is considerable freedom in choosing the K_p, and this raises three obvious questions:

- Is there a canonical choice of the K_p?
- How small can we make U' and W'?
- Can we ensure that the functions (5) are not merely symmetric but alternating?

These questions were first raised and also to a large extent answered in [6]; proofs of the assertions which follow can be found there. The motive for ensuring that the functions (5) are alternating is that it implies that the ranks of these functions are even; this means that their coranks, which are equal to the dimension of the 2-Selmer group, are congruent mod 2 to $\dim U'_{\mathcal{B}}$ and $\dim W'_{\mathcal{B}}$.

The answer to the first question appears to be negative, though there is little freedom in the optimum choice of the K_p — particularly if one wishes to obtain not merely Lemma 1 but Theorem 1. Since $U'_{\mathcal{B}} \supset U_{\mathcal{B}} \cap W_{\mathcal{B}}$, the best possible answer to the second question would be that we can achieve $U'_{\mathcal{B}} = U_{\mathcal{B}} \cap W_{\mathcal{B}}$; we shall do this by satisfying the stronger requirement

$$W_{\mathcal{B}} = (U_{\mathcal{B}} \cap W_{\mathcal{B}}) \oplus (K_{\mathcal{B}} \cap W_{\mathcal{B}}). \tag{6}$$

For suppose that (6) holds; then $W_{\mathcal{B}} + K_{\mathcal{B}} = (U_{\mathcal{B}} \cap W_{\mathcal{B}}) + K_{\mathcal{B}}$ and it follows immediately that

$$U'_{\mathcal{B}} = U_{\mathcal{B}} \cap (W_{\mathcal{B}} + K_{\mathcal{B}}) = U_{\mathcal{B}} \cap W_{\mathcal{B}}. \tag{7}$$

The motivation for (6) is that we want to make $W_{\mathcal{B}} \cap K_{\mathcal{B}}$ as large as possible — that is, to choose $K_{\mathcal{B}}$ so that as much of it as possible is contained in $W_{\mathcal{B}}$. But because $K_{\mathcal{B}}$ must be complementary to $U_{\mathcal{B}}$, only the part of $W_{\mathcal{B}}$ which is complementary to $W_{\mathcal{B}} \cap U_{\mathcal{B}}$ is available for this purpose.

Since the 2-Selmer group $U_{\mathcal{B}} \cap W_{\mathcal{B}}$ is identified with the left and right kernels of each of the functions (5), if (7) holds then these functions are trivial and therefore alternating. The formal statement of all this is as follows.

Lemma 1 *We can choose maximal isotropic subspaces $K_p \subset V_p$ for each p in $\mathcal{B}$ so that $V_{\mathcal{B}} = U_{\mathcal{B}} \oplus K_{\mathcal{B}}$. We can further ensure that*

$$W_{\mathcal{B}} = (U_{\mathcal{B}} \cap W_{\mathcal{B}}) \oplus (K_{\mathcal{B}} \cap W_{\mathcal{B}}),$$

which implies $U'_{\mathcal{B}} = U_{\mathcal{B}} \cap W_{\mathcal{B}}$. If so, the functions $\theta^\flat_{\mathcal{B}}$ and $\theta^\sharp_{\mathcal{B}}$ defined in (5) *are trivial.*

For some applications it is convenient to have an explicit description of the construction of the K_p; this is given in Appendix 2. But the other properties of the K_p chosen in this way are not at all obvious. Hence it is advantageous to consider other recipes for choosing the K_p, for which (6) does not hold but we can still prove that the functions (5) are alternating.

For this purpose we write $\mathcal{B}$ as the disjoint union of $\mathcal{B}'$ and $\mathcal{B}''$, where we shall always suppose that 2 and ∞ are both in $\mathcal{B}'$. For any odd prime p we denote by T_p the subset of V_p consisting of those triples (μ_1, μ_2, μ_3) with $\mu_1\mu_2\mu_3 = 1$ for which each μ_i is in $\mathfrak{o}_p^*/\mathfrak{o}_p^{*2}$ — that is, each μ_i is the image of a p-adic unit. The main point of the following theorem is that for p in $\mathcal{B}''$ it enables us to replace the complicated inductive definition of K_p used in the proof of Lemma 1 by the much simpler choice $K_p = T_p$. How one chooses $\mathcal{B}''$ depends on the particular application which one has in mind.

Theorem 1 *Let $\mathcal{B}$ be the disjoint union of $\mathcal{B}' \supset \{2, \infty\}$ and $\mathcal{B}''$. We can construct maximal isotropic subspaces $K_p \subset V_p$ such that $V_{\mathcal{B}} = U_{\mathcal{B}} \oplus K_{\mathcal{B}}$,*

$$W_{\mathcal{B}'} = (U_{\mathcal{B}'} \cap W_{\mathcal{B}'}) \oplus (K_{\mathcal{B}'} \cap W_{\mathcal{B}'}) \tag{8}$$

and $K_v = T_v$ for all v in $\mathcal{B}''$; and (8) *implies that $U'_{\mathcal{B}'} = U_{\mathcal{B}'} \cap W_{\mathcal{B}'}$. Moreover*

$$U'_{\mathcal{B}} = \jmath_* U'_{\mathcal{B}'} \oplus \tau_{\mathcal{B}} W'_{\mathcal{B}''} = \jmath_* U'_{\mathcal{B}'} \oplus \left(\oplus_{p \in \mathcal{B}''} \tau_B W'_p \right), \tag{9}$$

and the restriction of $\theta^\flat_{\mathcal{B}}$ to $\jmath_* U'_{\mathcal{B}'} \times \jmath_* U'_{\mathcal{B}'}$ is trivial.

If $\mathcal{B}'$ also contains all the odd primes p such that the $v_p(c_i - c_j)$ are not all congruent mod 2, *then we can choose the K_p for p in $\mathcal{B}'$ so that also $\theta^\flat_{\mathcal{B}}$ is alternating on $U'_{\mathcal{B}}$.*

The appearance of $\jmath_* U'_{\mathcal{B}'}$ in and just after (9) calls for some explanation. Let u be any element of $U_{\mathcal{B}'}$; then u is in $U_{\mathcal{B}}$. Moreover, for p in $\mathcal{B}''$ the image of u in V_p is in $T_p = K_p$ and therefore in $K_p + W_p$; hence u is in $U'_{\mathcal{B}}$. In this way we define a map $U'_{\mathcal{B}'} \to U'_{\mathcal{B}}$ which is clearly an injection and which we denote by $\jmath_*$.

Lemma 1 is the special case of Theorem 1 in which $\mathcal{B}' = \mathcal{B}$ and $\mathcal{B}''$ is empty. But the proof of Lemma 1 is a necessary step (and indeed the most substantial step) in the proof of Theorem 1. Indeed, to prove Theorem 1 we construct the K_p for p in $\mathcal{B}'$ according to the recipe in Appendix 2; for the final sentence of the theorem we need the particular version of the recipe which involves the functions ϕ_i.

The main application of Theorem 1 is to twisted curves Γ_b, where we can clearly take b to be an integer. Let $\mathcal{S}$ denote the set of bad primes for Γ itself — that is, $2, \infty$ and the odd primes dividing $(c_1 - c_2)(c_1 - c_3)(c_2 - c_3)$; and let $\mathcal{B} \supset \mathcal{S}$ be the set of bad primes for Γ_b. If we are to apply any part of Theorem 1, $\mathcal{B}$ must also contain all the odd primes dividing b; and such applications are much simpler when b is a unit at every prime of $\mathcal{S}$. (We can always arrange this by treating Γ_b as the twist of Γ_c by b/c, where c is the largest divisor of b which is a unit outside $\mathcal{S}$.) To describe the effect of twisting, we shall denote by d_b the dimension of the 2-Selmer group of

Γ_b regarded as a vector space over $\mathbf{F}_2$; we write $d = d_1$ for the dimension of the 2-Selmer group of Γ itself. It is now possible to prove results about $d_b - d$, the change in the dimension of the 2-Selmer group as one goes from Γ to Γ_b. There is reason to expect that statements about the parities of d and d_b will be simpler and much easier to prove than statements about their actual values. The two major statements known about d_b are Lemma 2 and Theorem 2; Lemma 2 is an easy consequence of the last sentence of Theorem 1, and Theorem 2 is an easy consequence of Lemma 3 below.

Lemma 2 *If b is in $\mathfrak{o}_p^*$ for every $p \in \mathcal{S}$, then $d_b \equiv \dim(U_{\mathcal{S}} \cap W_{\mathcal{S}}) \bmod 2$ where $W_{\mathcal{S}} = \oplus_{p \in \mathcal{S}} W_p$ and the W_p must be defined with respect to Γ_b and not with respect to Γ. Thus $d_b \bmod 2$ only depends on the classes of b in the k_p^*/k_p^{*2} for p in $\mathcal{S}$.*

To prove Lemma 3 we need to take $\mathcal{B}' = \mathcal{S} \setminus \{p\}$; thus the last sentence of Theorem 1 is not applicable though the rest of that theorem is.

Lemma 3 *Let p be an odd prime in $\mathcal{S}$ such that*

$$v_p(c_1 - c_2) > 0, \quad v_p(c_1 - c_3) = v_p(c_2 - c_3) = 0.$$

Let b in k^* be such that b is in k_q^{*2} for all q in $\mathcal{S}$ other than p and b is a quadratic non-residue at p. Then d and d_b have opposite parities.

It is not hard to prove the analogue of Lemma 3 for the case $p = \infty$, though the proof falls outside the machinery described in this note. The combination of this result and Lemma 3 yields Theorem 2. (The analogue of Lemma 3 for $p = 2$ can be confidently asserted, on the basis of a large amount of numerical evidence, and the proof of it probably requires no new ideas. But even the statement involves so extensive a separation of cases that it is unlikely soon to appear in print.)

Theorem 2 *Let b', b'' in k^* be such that b'/b'' is a unit at all $p \in \mathcal{S}$ and $b'/b'' \equiv 1 \bmod 8$. Let $\mathcal{S}^*$ be the set of $p \in \mathcal{S}$ for which b'/b'' is not in k_p^{*2}. Let $\mathcal{S}^{**}$ consist of the finite odd p in $\mathcal{S}^*$ for which the $v_p(c_i - c_j)$ are not all equal and the smallest two of them are even, together with ∞ if $b'/b'' < 0$. Then*

$$d_{b'} - d_{b''} \equiv \#\mathcal{S}^{**} \bmod 2.$$

We can define a 4-*covering* and a 4-*descent* (sometimes called a *second descent*) by extension of the diagram (2). Let $\mathcal{C}$ be a 2-covering of Γ; then a 4-covering of Γ above this 2-covering is a curve $\mathcal{D}$ which fits into the commutative diagram

$$\begin{array}{ccccc} \Gamma & \longrightarrow & \Gamma & \longrightarrow & \Gamma \\ \| & & \| & \nearrow & \\ \mathcal{D} & \longrightarrow & \mathcal{C} & & \end{array}$$

in which the vertical maps are biregular (but defined over $\mathbf{C}$ rather than $\mathbf{Q}$) and each upper map is multiplication by 2. If $\mathcal{C}$ is everywhere locally soluble, we say that it *admits a second descent* if we can find such a $\mathcal{D}$ which is everywhere locally soluble. If $\mathcal{C}$ is actually soluble in $\mathbf{Q}$, then it certainly admits a second descent; thus carrying out a second descent is a way of replacing the 2-Selmer group by a hopefully smaller group which however still contains $\Gamma(\mathbf{Q})/2\Gamma(\mathbf{Q})$. A second descent may therefore refine the information about the Mordell-Weil group which is obtained from the 2-descent.

In its classical form, the process of 4-descent was constructive but it was arithmetically unattractive, largely because it involved a field extension. But Cassels [1] has shown how to determine which elements of the 2-Selmer group do admit a second descent, while working entirely in $\mathbf{Q}$. He constructs an alternating bilinear form g on the 2-Selmer group, whose kernel consists of exactly those elements which admit a second descent. Let $\mathcal{S}$ again be the set of bad primes for Γ, with $\mathcal{S} \supset \{2, \infty\}$, and let $\mathbf{m}'$ and $\mathbf{m}''$ be two triples in $U_{\mathcal{S}}$ which represent elements of the 2-Selmer group of Γ. If i, j, k is any permutation of $1, 2, 3$ we denote by $\mathcal{C}_i(\mathbf{m}')$ the conic

$$m'_j y_j^2 - m'_k y_k^2 = (c_k - c_j) y_0^2. \tag{10}$$

In view of (1) there is a map $\mathcal{C}(\mathbf{m}') \to \mathcal{C}_i(\mathbf{m}')$; so $\mathcal{C}_i(\mathbf{m}')$ is everywhere locally soluble. Because $\mathcal{C}_i(\mathbf{m}')$ is a conic, this implies that it is soluble in $\mathbf{Q}$; so choose a rational point P_i on $\mathcal{C}_i(\mathbf{m}')$ and let $\mathsf{L}_i(y_0, y_j, y_k) = 0$ be the equation of the tangent to $\mathcal{C}_i(\mathbf{m}')$ at P_i. By abuse of language, we can treat L_i as a homogeneous linear form in y_0, y_j, y_k; strictly speaking, it is only defined up to multiplication by an element of $\mathbf{Q}^*$, but it will not matter which multiple we choose. For each p in $\mathcal{S}$, choose a p-adic point Q_p on the affine curve $\mathcal{C}(\mathbf{m}')$. Then g is defined by

$$g(\mathbf{m}', \mathbf{m}'') = \prod_{p \in \mathcal{S}} \prod_i (\mathsf{L}_i(\mathsf{Q}_p), \mathbf{m}''_i)_p$$

where the bracket on the right is as usual the Hilbert symbol.

APPENDIX 1 — Explicit description of the W_p

The main purpose of this Appendix is to give an explicit description of the W_p. The calculations are sometimes simplified by using the fact that W_p is isotropic and contains the three triples like (3); thus if $\mathbf{m}$ is in W_p then the three results like

$$(c_1 - c_2, m_3)_p = (c_1 - c_3, m_2)_p$$

all hold. The case $p = \infty$, which is trivial, is Lemma 5. The case when p is odd, the simplest proof of which can be found in [6], is Lemma 5. The results for the case $p = 2$ are much more complicated; they can be found in [7] but are not reproduced here.

Lemma 4 *After renumbering, suppose that $c_1 > c_2 > c_3$. Then W_∞ consists of the classes of $(1,1,1)$ and $(-1,-1,1)$.*

In Lemma 5 and Theorem 3, $a_1 \sim a_2$ will mean that a_1/a_2 is in k_p^{*2}.

Lemma 5 *Let p be an odd prime.*

If p divides all the $c_i - c_j$ to the same even power, then $W_p = (\mathfrak{o}_p^*/\mathfrak{o}_p^{*2})^2$. If p divides all the $c_i - c_j$ to the same odd power, then W_p consists of the classes of $(1,1,1)$ and the three triples like (3).

Now suppose that p does not divide all the $c_i - c_j$ to the same power. After renumbering, let

$$v_p(c_1 - c_2) > v_p(c_1 - c_3) = v_p(c_2 - c_3).$$

Denote by η the class of $c_1 - c_2$, by ϵ the class of $c_1 - c_3$ and $c_2 - c_3$, and by ν the class of quadratic non-residues $\bmod p$.

If $v(\epsilon)$ is odd then W_p consists of the classes of

$$(1,1,1),\ (\eta\epsilon, \eta, \epsilon),\ (-\eta, -\eta\epsilon, \epsilon),\ (-\epsilon, -\epsilon, 1).$$

If $v(\eta)$ is odd and $v(\epsilon)$ even then W_p consists of the classes of

$$(1,1,1),\ (\eta\epsilon, \eta, \epsilon),\ (\nu, \nu, 1),\ (\nu\eta\epsilon, \nu\eta, \epsilon).$$

If $v(\eta)$ and $v(\epsilon)$ are both even and $\epsilon \sim \nu$ then W_p consists of the classes of

$$(1,1,1),\ (\nu,\nu,1),\ (\nu,1,\nu),\ (1,\nu,\nu).$$

If $v(\eta)$ and $v(\epsilon)$ are both even and $\epsilon \sim 1$ then W_p consists of the classes of

$$(1,1,1),\ (\nu,\nu,1),\ (p,p,1),\ (p\nu,p\nu,1).$$

A number of people have proved results of the form: let p be in $\mathcal{S}$ and assume that $\mathcal{C}(\mathbf{m})$ is locally soluble at all primes other than perhaps p; then provided that certain local conditions on Γ hold, $\mathcal{C}(\mathbf{m})$ is also locally soluble at p. The best approach to this kind of result is as follows. For any permutation i,j,k of $1,2,3$ let $\mathcal{C}_k(\mathbf{m})$ denote the conic

$$m_i y_i^2 - m_j y_j^2 = (c_j - c_i) y_0^2,$$

this being essentially the same as the notation of (10). The existence of a map $\mathcal{C}(\mathbf{m}) \to \mathcal{C}_k(\mathbf{m})$ implies that $\mathcal{C}_k(\mathbf{m})$ is also locally soluble everywhere except possibly at p. Since $\mathcal{C}_k(\mathbf{m})$ is a conic, it follows that $\mathcal{C}_k(\mathbf{m})$ is also locally soluble at p — a condition which is equivalent to

$$(m_i(c_j - c_i), m_k)_p = 1. \tag{11}$$

Hence $\mathcal{C}(\mathbf{m})$ is locally soluble at p provided that this is implied by the local solubility of the three $\mathcal{C}_k(\mathbf{m}')$ at p — that is, by the three conditions like (11). The question is under what local conditions on Γ at p this holds. Such results can be read off from the description of W_p; but in fact we can decide this question without knowing W_p. For we do know that the order of W_p is 2, 4 or 8 according as p is ∞, odd or 2. It is therefore enough to count the set of triples $\mathbf{m}$ which satisfy the three equations like (11); for this set contains W_p, so that it is equal to W_p if and only if it has the same order as W_p. Even when $p = 2$, this calculation is trivial to program.

The conclusions for $p = \infty$ and p odd are given in the following theorem. Those for $p = 2$ are too complicated to justify explicit statement.

Theorem 3 *Suppose that $\mathcal{C}(\mathbf{m})$ is everywhere locally soluble except possibly at one prime p which is in $\mathcal{S}$. If $p = \infty$ then $\mathcal{C}(\mathbf{m})$ is also locally soluble at p. If p is odd then $\mathcal{C}(\mathbf{m})$ is also locally soluble at p except perhaps when $c_i - c_k \sim c_j - c_k \sim 1$ for some permutation i,j,k of $1,2,3$.*

APPENDIX 2 — Construction of the K_p

In this Appendix we show how to construct the K_p. We do in fact prove a more general result, but this is only because otherwise we would be forced into a needlessly complicated notation. The reader will see that (subject to the introduction of the temporarily mysterious functions ϕ_i) the hypotheses of Lemma 6 mimic the structure described in the main body of the text. I give here only that part of the proof which is really an algorithm for the construction; a complete proof can be found in [6].

Lemma 6 *Let the V_i be n vector spaces over $\mathbf{F}_2$, each equipped with a non-degenerate additive alternating bilinear form ψ_i with values in $\mathbf{F}_2$. Denote by ψ the sum of the ψ_i, which is a non-degenerate bilinear form on $V = \oplus V_i$. For each i let W_i be maximal isotropic in V_i, and let U be maximal isotropic in V with respect to ψ. Then there exist maximal isotropic subspaces $K_i \subset V_i$ such that $V = U \oplus K$ and*

$$W = (U \cap W) \oplus (K \cap W) \tag{12}$$

where $W = \oplus W_i$ and $K = \oplus K_i$. Moreover $U \cap (W + K) = U \cap W$.

Suppose also that there are functions ϕ_i on V_i with values in $\mathbf{F}_2$ which satisfy

$$\phi_i(\xi + \eta) = \phi_i(\xi) + \phi_i(\eta) + \psi_i(\xi, \eta) \tag{13}$$

for any ξ, η in V_i, and let ϕ on V be the sum of the ϕ_i. Assume that ϕ is trivial on U and ϕ_i is trivial on W_i. Then we can further ensure that ϕ_i is trivial on K_i and therefore ϕ is trivial on K.

Proof If any V_i has dimension greater than 2, we can decompose it as a direct sum of mutually orthogonal subspaces of dimension 2, on each of which the restriction of the bilinear form ψ_i is non-degenerate and each of which meets W_i in a subspace of dimension 1. This only reduces our freedom to choose the K_i, and the triviality of ϕ_i on the old K_i will follow from its triviality on the new and smaller K_i by (13). Thus we can assume that every V_i has dimension 2 and every W_i has dimension 1. We proceed by induction on n, the case $n = 0$ being trivial.

We shall assume that the ϕ_i exist, noting in the appropriate place how to modify the argument to prove the first part of the lemma without using the existence of the ϕ_i. If we regard W_n as a subspace of V, either $W_n \subset U$ or

W_n is not contained in U and therefore meets it only in the origin. In each of these cases, we shall choose an α_i in V_i with $\phi_i(\alpha_i) = 0$ and use it to generate K_i. After reordering, we can assume that either W_n is not contained in U or every W_i is contained in U and therefore $W \subset U$.

Since U is isotropic it cannot contain V_i; so if $W_n \subset U$ and therefore $W_i \subset U$ for each i, then each V_i contains just two elements which do not lie in U. Denote them by α_i' and α_i'', and let β_i be the nontrivial element of W_i; thus $\alpha_i'' = \alpha_i' + \beta_i$. Since $\phi_i(\beta_i) = 0$ it follows from (13) and the non-degeneracy of ψ_i on V_i that

$$\phi_i(\alpha_i') + \phi_i(\alpha_i'') = \psi_i(\alpha_i', \beta_i) = 1;$$

choose α_i to be whichever of α_i' and α_i'' satisfies $\phi_i(\alpha_i) = 0$. (If we do not assume the existence of the ϕ_i then we can take α_i to be either of α_i' and α_i''.) Let K_i be the vector space generated by α_i; thus

$$W_i = U \cap W_i = (U \cap W_i) \oplus (K_i \cap W_i)$$

for each i, which implies (12). Moreover $U \supset W$ and therefore $U = W$ because U and W have the same dimension. So

$$V = \oplus V_i = \oplus (W_i \oplus K_i) = W \oplus K = U \oplus K.$$

If U does not contain W_n, then the non-trivial element of W_n is not in U. Denote this element by α_n, so that $\phi_n(\alpha_n) = 0$ by hypothesis. Let K_n be the vector space generated by α_n; thus $K_n = W_n$ and

$$W_n = (U \cap W_n) \oplus (K_n \cap W_n). \tag{14}$$

The construction now proceeds by induction on n. Write

$$V^- = V_1 \oplus \ldots \oplus V_{n-1}, \quad U^- = V^- \cap (U \oplus W_n). \tag{15}$$

It is straightforward to show that U^- is maximal isotropic in V^-. For the pair U^-, V^- we must replace the question whether $U \supset W$ by the question whether $U \oplus W_n$ contains $W^- = W_1 \oplus \ldots \oplus W_{n-1}$. By the induction hypothesis for the pair $U^- \subset V^-$, there exist K_i maximal isotropic in V_i for each $i < n$ such that if $K^- = (K_1 \oplus \ldots \oplus K_{n-1})$ then $V^- = U^- \oplus K^-$ and

$$W^- = (U^- \cap W^-) \oplus (K^- \cap W^-). \tag{16}$$

The need to check the remaining details of the argument can be circumvented by an appeal to Cassels' Axiom: all vector space theorems are trivial.

When we apply Lemma 6 to the construction of the K_p for p in $\mathcal{B}'$ and the proof of Theorem 1, we replace i by p and ψ_i by e_p; but note that we have chosen to write e_p multiplicatively and ψ_i additively. For $\mathbf{m}$ in V_p we take $\phi_p(\mathbf{m})$ to be any one of the expressions

$$(m_i(c_i - c_j)(c_i - c_k), m_j(c_j - c_i)(c_j - c_k))_p,$$

whose values are easily shown to be equal. The significance of ϕ_p is as follows. The antipodal involution $(x, y) \mapsto (x, -y)$ on Γ induces an involution on each 2-covering $\mathcal{C}(\mathbf{m})$; in the notation of (1) this involution reverses the signs of y_1, y_2, y_3. The quotient of $\mathcal{C}(\mathbf{m})$ by this involution is a smooth projective curve $\mathcal{D}(\mathbf{m})$ of genus 0, which is given by

$$(c_2 - c_3)m_1 y_1^2 + (c_3 - c_1)m_2 y_2^2 + (c_1 - c_2)m_3 y_3^2 = 0; \tag{17}$$

and $\phi_p(\mathbf{m})$ is just the class $[\mathcal{D}(\mathbf{m})]$ as an element of Br k_p.

REFERENCES

[1] J.W.S.Cassels, Second descent for elliptic curves, J. reine angew. Math. 494(1998), 101-127.

[2] J-L.Colliot-Thélène, A.N.Skorobogatov and Sir Peter Swinnerton-Dyer, Hasse principle for pencils of curves of genus one whose Jacobians have rational 2-division points, Invent. Math. 134(1998), 579-650.

[3] D.Husemöller, Elliptic Curves, Graduate Texts in Mathematics, 111 (Springer, 1987).

[4] J.S.Milne, Arithmetic Duality Theorems (Academic Press, Boston, 1986).

[5] J.H.Silverman, The Arithmetic of Elliptic Curves, Graduate Texts in Mathematics, 106 (Springer, 1986).

[6] A.N.Skorobogatov and Sir Peter Swinnerton-Dyer, A further refinement of 2-descent on elliptic curves, I (to appear).

[7] Sir Peter Swinnerton-Dyer, Rational points on certain intersections of two quadrics, in *Abelian varieties* (ed. W.Barth, K.Hulek and H.Lange) (de Gruyter, Berlin, 1995), 273-292.

[8] Sir Peter Swinnerton-Dyer, Some applications of Schinzel's hypothesis to diophantine equations, in *Number theory in progress* (ed. K.Győry, H.Iwaniec and J.Urbanowicz) (Berlin, 1999), 503-530.