Elliptic curves and surfaces and their Mordell-Weil groups

EC@UC, 17 May 2014

Noam D. Elkies, Harvard University

Motivation: families of elliptic curves

$$E = E_t : Y^2 + a_1XY + a_3Y = X^3 + a_2X^2 + a_4X + a_6$$

depending on parameter(s) t (i.e. the coefficients a_i are functions of t).

Examples:

• $(a_1, a_2, a_3, a_4, a_6) = (1 - t - t^2,\ t^2 + t^3,\ t^2 + t^3,\ 0,\ 0)$:
The general elliptic curve with a 7-torsion point, here at $(X, Y) = (0, 0)$, depending on $t \in \mathbf{P}^1 \cong X_1(7)$

• $(a_1, a_2, a_3, a_4, a_6) = (t_0 + t_1 + t_2,\ 0,\ t_0t_1t_2,\ 0,\ 0)$:
[cont'd]

Motivating examples cont'd:

• $(a_1, a_2, a_3, a_4, a_6) = (t_0 + t_1 + t_2,\ 0,\ t_0 t_1 t_2,\ 0,\ 0)$:
The general (E, T, P) where E is an elliptic curve with a 3-torsion point T, here at $(X, Y) = (0, 0)$, and P is an arbitrary point on E, here at $(-t_1 t_2, t_1 t_2^2)$; depending on $(t_0 : t_1 : t_2) \in \mathbf{P}^2$. [Cyclic permutations yield $(E, T, P \pm T)$, while $t_1 \leftrightarrow t_2$ yields $(E, T, -P)$.]

• $(a_1, a_2, a_3, a_4, a_6) = (a,\ b,\ ab,\ 0, 0)$ where

$$a = (8t - 1)(32t + 7), \quad b = 8(t + 1)(15t - 8)(31t - 7):$$

$(E, T, P_1, P_2, P_3, P_4)$ where T is 4-torsion and the P_i are independent for all but finitely many $t \in \mathbf{P}^1(\mathbf{Q})$ [NDE 2006].

P_i can be taken to have X-coordinates $-15(t + 1)(31t - 7)(32t + 7)/4$,
$(8t - 1)(15t - 8)(31t - 7)(32t + 7)$, $-(t + 1)(8t - 1)(15t - 8)(32t + 7)$, and
$-4(t + 1)(2t + 5)(15t - 8)(32t + 7)$.

The parameter t varies over a base variety B of dimension $d > 0$; for many applications we want B to have plenty of rational points (ideally $\mathbf{P}^d$ or some other rational variety).

The family of elliptic curves E_t can also be regarded as a single elliptic curve over the function field of B (since this function field contains the coefficients a_i), and also as a variety $\mathcal{E}$ of dimension $d+1$ together wih a map $\pi : \mathcal{E} \to B$, as we explain next.

The interaction of these three points of view (family of curves E_t, curve E over the function field, and variety $\mathcal{E}$) is our main theme here.

The equation defining E_t can be regarded as the equation for a (hyper)surface in the product of B with (X, Y) space. This is our variety $\mathcal{E}$ of dimension $d+1$. It comes with a map $\pi : \mathcal{E} \to B$ (forget the X, Y coordinates) such that $E_t = \pi^{-1}(t)$ for all $t \in B$.

A rational point P on the generic curve E_t in the family is then tantamount to a point on the curve over the function field of B, and to a *section* of π, i.e. a map $s_P : B \to \mathcal{E}$ such that $\pi \circ s_P = \mathrm{id}_B$. The group law on E_t gives a group structure to the set of sections, giving the *Mordell-Weil group* of $\mathcal{E}$.

[Yes, with more effort we can give a more precise picture of $\mathcal{E}$, π, and the sections s_P, and we'll soon do this at least for $B = \mathbf{P}^1$; but for now we're just doing birational algebraic geometry for motivation.]

Algebraic geometry in dimension d gets hard quickly as d increases ("zero, one, two, many"). Since we want to use the geometry of $\mathcal{E}$, which has dimension $d+1$, we'll concentrate on the case $d=1$, when $\mathcal{E}$ is a surface (dimension $d+1=2$), and usually $B=\mathbf{P}^1$. [If $d>1$ but B is rational then we can still use this special case via a a nonconstant $\iota:\mathbf{P}^1\to B$, for which there are many choices.]

In the case of surfaces we can also use intersection theory: s_P always has *codimension* 1, but here it also has *dimension* 1 so we can form $s_P\cdot s_Q$, and even $s_P\cdot s_P$. This lets us recover the Mordell-Weil group from standard invariants of $\mathcal{E}$ together with the elliptic fibration π (Shioda-Tate).

Applications include:

• Silverman's specialization theorem: if K is a number field (think $K = \mathbf{Q}$), and $B = \mathbf{P}^1(K)$ with coordinate t, then the specialization homomorphism $E(K(t)) \to E_t(K)$ is injective for all but finitely many $t \in K$ (and the exceptions can be listed effectively). In particular, E_t/K almost always has the same rank as $E/K(t)$ (and at least the same torsion).

• [Šafarevič-Tate] for a finite field k, there exist curves $E/k(t)$ of arbitrarily high rank.

• [Shioda, NDE] In some cases the height pairing on $E(k(t))$ yields dense lattice packings of spheres.

The elliptic surface $\mathcal{E}$

We'll almost always take $B = \mathbf{P}^1$. To get a projective variety over the $(t_1 : t_0)$ line, regard the coefficients a_i of E ($i = 1, 2, 3, 4, 6$) as homogeneous polynomials in (t_0, t_1) of degrees ic, with c minimal. So a nonzero section is a solution of the equation of E in homog. rational functions X, Y of degree $2c, 3c$ (and the X, Y plane is really the projective $(X : Y : 1)$ plane associated to $O(2c) \oplus O(3c) \oplus O(0)$).

E_t is a smooth elliptic curve if $\Delta(t) \neq 0$. It turns out that $\deg \Delta = 12c$. So once $c > 0$ there's at least one singular fiber. Usually (as when Δ has distinct roots) the surface is still smooth. But sometimes we need to blow up once or more to reach a smooth surface. Do it minimally, as described by Tate. We'll use $\mathcal{E}$ to mean the resulting smooth projective surface, which still has a map $\pi : \mathcal{E} \to \mathbf{P}^1$ whose generic fiber E_t is an elliptic curve.

Standard picture/cartoon of an elliptic surface:

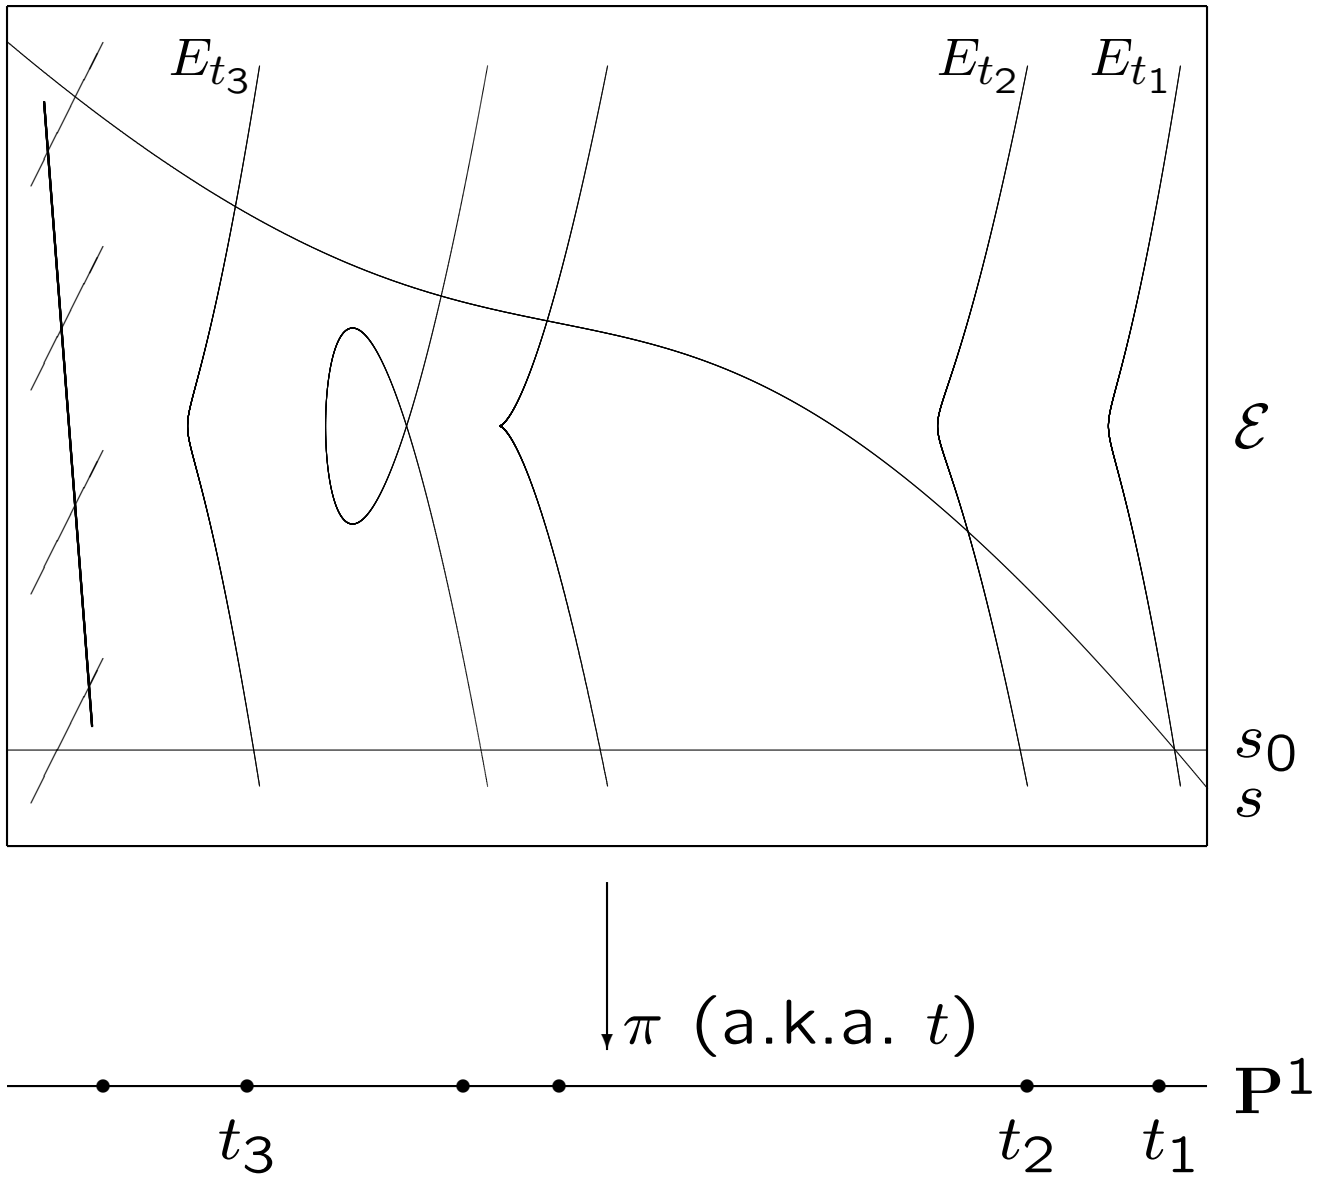

The larger c is, the more complicated $\mathcal{E}$ gets:

$c = 0$: constant curve $\mathbf{P}^1 \times E_0$

$c = 1$: rational elliptic surface (birationally $\mathbf{P}^2$ if $K = \overline{K}$)

$c = 2$: elliptic K3 surface

$c \geq 3$: "honestly elliptic" surface

For $c = 1$ and $c = 2$ there may be many essentially different choices of π. Example with $c = 2$: the surface

$$E : Y^2 = X^3 + \alpha t^4 X + (\beta_+ t^7 + \beta t^6 + \beta_- t^5)$$

is birationally

$$Y_1^2 = X_1^3 + \alpha X_1 + (\beta_+ t + \beta + \beta_- t^{-1})$$

where $(X, Y) = (t^2 X_1, t^3 Y_1)$, and this is also an elliptic surface over the X_1-line.

The Néron-Severi group NS(V)

The *Néron-Severi group* of any smooth variety V is the group of divisors on V modulo algebraic equivalence. A divisor D is a formal $\mathbf{Z}$-linear combination $\sum_i n_i D_i$ of subvarieties D_i of codimension 1. If $f : V \dashrightarrow \mathbf{P}^1$ is a rational function on V then its divisor $(f) = (f)_0 - (f)_\infty$ is said to be "linearly equivalent to zero"; such divisors form a subgroup of Div(V) because $(f) + (g) = (fg)$. This special case of "algebraic equivalence" is usually all we need (in general, replace $V \dashrightarrow \mathbf{P}^1$ by $V \dashrightarrow C$ for any curve C, and $0, \infty$ by any two points of C). It is known that NS(V) is a finitely-generated abelian group (Néron-Severi theorem).

Intersection pairing on NS(V) when dim(V) = 2

When dim $V = 2$, the Néron-Severi group carries a bilinear intersection pairing, characterized by the property that if D, D' are divisors with no common component then $D \cdot D'$ is the number of intersections of D with D', counted with multiplicity. The pairing is symmetric, and has signature $(1, \rho - 1)$ (Hodge index theorem; $\rho =$ rank of NS(V) is the Picard number).

Examples: if $V = \mathbf{P}^2$ then NS(V) = $\mathbf{Z}l$ where $l =$ class of a line, with $l^2 = 1$. If $V = \mathbf{P}^1 \times \mathbf{P}^1$ then NS(V) = $\mathbf{Z}l_1 + \mathbf{Z}l_2$ wih $l_i^2 = 0$ and $l_1 \cdot l_2 = 1$. (As usual D^2 means $D \cdot D$.)

And if $V = \mathcal{E}$...

Néron-Severi group of an elliptic surface

Various sources of divisors on an ellliptic surface $\mathcal{E}$: The fiber class f (NB all E_t are linearly equivalent); components of reducible fibers, as described by Tate's algorithm; and sections s_P. In fact that's all we need. More precisely [Tate, Shioda]: once $c > 0$ the map $P \mapsto s_P$ is a group isomorphism from the Mordell-Weil group to the quotient of NS($\mathcal{E}$) by the "trivial subgroup" of NS($\mathcal{E}$) generated by s_0, f, and components of reducible fibers.

It follows that the Mordell-Weil rank equals $\rho - 2 - \sum_t (n_t - 1)$, where n_t = # components of E_t (so $n_t - 1 = 0$ with finitely many exceptions; NB it's $n_t - 1$, not n_t, because each E_t, whether reducible or not, is in the fiber class f). $\Rightarrow$ Mordell-Weil rank $\leq 10c - 2$ (or $10c$ in positive characteristic).

This also yields the *height pairing* on MW/torsion, which is a canonical pos.-def. pairing with values in $\mathbf{Q}$. The divisors s_0, f generate a subgroup of NS($\mathcal{E}$) with intersection pairing $\left(\begin{smallmatrix} ? & 1 \\ 1 & 0 \end{smallmatrix}\right)$ where the "?" is s_0^2 (known to be $-c$); hence discriminant -1 and signature $(1,1)$. So NS($\mathcal{E}$) is the orthogonal direct sum of $\mathbf{Z}s_0 \oplus \mathbf{Z}f$ with some negative-definite lattice $L\langle -1\rangle$ of rank $\rho - 2$. So L is positive-definite, and it turns out it's always even: $D \cdot D \equiv 0 \bmod 2$ for all divisors D orthogonal to both s_0 and f.

By Tate-Shioda, the Mordell-Weil group is the quotient of L by the sublattice R coming from reducible-fiber components. We identify MW/torsion with the orthogonal projection of L to $(R \otimes \mathbf{Q})^\perp$, and this gives us the height pairing.

What does R look like?

Each reducible fiber's contribution to L is a root lattice, which may be $A_{\nu-1}$ (as for a fiber of Kodaira type I_ν), $D_{4+\nu}$ (for I_ν^*), or E_6, E_7, E_8 (for IV*, III*, II* respectively). These root sublattices are pairwise orthogonal, which is how Tate and Shioda knew that their contributions to NS($\mathcal{E}$); and once $c > 1$ they account for all the norm-2 vectors in L (so R is the root sublattice of L).

What does $\widehat{h}$ look like?

For any point P, the canonical height $\widehat{h}(P)$ [= pairing of P with itself] is then $-(s_P - s_0)^2 + O(1)$, with the $O(1)$ coming from the projection of s_P to $R \otimes \mathbf{Q}$ (only finitely many possibilities at each t). But

$$-(s_P - s_0)^2 = 2s_0 \cdot s_P - s_0^2 - s_P^2 = 2s_0 \cdot s_P + 2c.$$

So $\widehat{h}(P) = 2s_0 \cdot s_P + O(1)$, and $2s_0 \cdot s_P$ is the degree of the denominator of $x(P)$. So this fits well with Tate's approach to $\widehat{h}$ as well as Néron's. (Néron: sum of local contributions; Tate: quadratic form within $O(1)$ of log of usual height.)

Some connections and consequences

1. Silverman's specialization theorem

If the ground field K is "global" (e.g. $\mathbf{Q}$), and the Mordell-Weil group has rank r, then the specialization E_{t_0} has rank $\geq r$ for all sufficiently high(?) $t_0 \in K$.

More precisely: let $P_1, \ldots, P_r$ be independent elements in $E(K(t))$. For all t_0 where $\Delta(t_0) \neq 0$ we get an elliptic curve E_{t_0} and a specialization homomorphism $\sigma_{t_0} : E(K(t)) \to E_{t_0}$. We claim the $\sigma_{t_0}(P_i)$ eventually remain independent.

[cont'd: r points $P_i \in E(K(t))$ independent; fix $t_0 \in K$, get $\sigma_{t_0} : E(K(t)) \to E_{t_0}$; we claim the $\sigma_{t_0}(P_i)$ eventually remain independent.]

Combine two key facts:

First, (in both E and E_t) independence holds iff the matrix $\langle P_i, P_j \rangle$ is positive definite, where

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

(or maybe 1/2 that, depending on convention).

Second, σ_{t_0} multiplies $\widehat{h}$ by $h(t_0)$ to within $O(1)$. [This "h" is the logarithmic height $h(m/n) = \log \max(|m|, |n|)$.] More precisely: given P, we have

$$\widehat{h}(\sigma_{t_0}(P)) = h(t_0)\widehat{h}(P) + O_P(1)$$

for all $t_0 \in K$.

[Silverman specialization, cont'd]

So, we start with a positive-definite matrix $M = (\langle P_i, P_j \rangle)_{i,j=1}^r$, multiply by a scalar $h(t_0)$, and change by $O(1)$. [Only $(r^2+r)/2 =$ finitely many choices of P to use in formula $\widehat{h}(\sigma_{t_0}(P)) = h(t_0)\widehat{h}(P) + O_P(1)$.] That's $h(t_0)(M + O(h(t_0))^{-1})$, so still positive-definite for $h(t_0)$ large enough (pos.-def. matrices are open in $\mathbf{R}^{(r^2+r)/2}$).

Note that unlike many finiteness results in advanced number theory this is entirely "effective": given K, E, and P_i we can track down all the $O(1)$'s and deduce an upper bound on $h(t_0)$, etc.

[NB Not true in $\overline{K}$: for any $P = \sum_i c_i P_i$ with $s_0 \cdot s_P > 0$ there's $t_0 \in \overline{K}$ s.t. $\sigma_{t_0}(P) = 0$ (unless E_{t_0} happens to be singular).]

2. Tate & Šafarevič: Arbitrarily large rank over $\mathbf{F}_p(t)$

It is not known whether there exist nonconstant elliptic curves of unbounded rank over $\mathbf{Q}(t)$, or even $\mathbf{C}(t)$ (the records are respectively 18 [NDE 2006, complicated] and 68 [Shioda 1992, $Y^2 = X^3 + t^{360} + 1$]). But over $\mathbf{F}_q(t)$, yes, even when $q = p^f$ is prime ($f = 1$).

Tate and Šafarevič proved this by exploiting supersingular curves. Let $E_0 : Y^2 = X^3 + aX + b$ be such a curve. (e.g. take $b = 0$ when $p \equiv -1 \bmod 4$, and $a = 0$ when $p \equiv -1 \bmod 3$. If $p = 2$ use $Y^2 + Y = X^3$.) Then consider its quadratic twist $D(t)Y^2 = X^3 + aX + b$, i.e. $Y^2 = X^3 + D^2aX + D^3b$. This becomes isomorphic with E over the quadratic extension $u^2 = D(t)$. [For $p = 2$: $Y^2 + Y = X^3 + Q(t)$, $u^2 + u = Q(t)$.]

So, let C be the hyperelliptic curve $u^2 = D(t)$ [or use the curve $u^2 + u = Q(t)$ for $p = 2$]. The surface $\mathcal{E}$ is birationally $(C \times E_0)/\{1, \iota\}$ where $\iota =$ hyperelliptic $\times(-1)$. For suitable C (namely C with supersingular Jacobian), $C \times E_0$ has large Picard number, which is mostly inherited by its quotient $\mathcal{E}$ [maps from C to E_0 that commute with ι, which comes to $\mathrm{Hom}(J(C), E_0)$]. Example: $D(t) = q^{t^e} - q$ (or $u^2 + u = t^{2^e} + 1$).

The Picard number grows as Cp^e. Most of $\mathrm{NS}(\mathcal{E})$ is not defined over $\mathbf{F}_p$, only $\mathbf{F}_{p^{2e}}$. But (a common tool) the part that is defined over $\mathbf{F}_p$ is the subgroup fixed by $\mathrm{Gal}(\mathbf{F}_{p^{2e}}/\mathbf{F}_p)$. Tate & Šafarevič [1967] show that this group's rank still grows as $C'p^e/e \to \infty$. (Ulmer [2002] constructed $E/\mathbf{F}_p$ with nonconstant j_0 and $r \to \infty$ by finding another way to make $\mathcal{E}$ "supersingular".)

3. Dense lattice packings of spheres (Shioda, NDE)

Recall: positive-definite quadratic form on $\mathbf{Z}^r \iff$ lattice $L \subset \mathbf{R}^r \iff$ lattice packing of balls (traditionally "spheres") of radius $\frac{1}{2} N_{\min}(L)^{1/2}$. [$N_{\min}(L) =$ minimal nonzero norm $\langle v, v \rangle$.]

For $r > 8$ (except $r = 24$) we don't now how large a fraction of $\mathbf{R}^r$ these balls can cover; i.e. we don't know

$$\max_L \big((N_{\min}(L))^r / \operatorname{disc}(L) \big).$$

We have upper bounds (inequalities) and lower bounds (constructions) but they rarely match.

For some r the best lattices we know are Mordell-Weil lattices of surfaces like those of Tate and Šafarevič, but over $\overline{\mathbf{F}}_p$.

Examples: $y^2 + y = x^3 + t^{2^e+1}$ over $\overline{\mathbf{F}}_2$ yields lattice of rank $2^{e+1} = 4$, 8, 16, 32 that tied previous records, and 64, 128, . . . , 1024(, 2048) that set new records. Also $y^2+y = x^3+t^{13}$ gives the Leech lattice in $\mathbf{R}^{24}$ (there are 196560 solutions $(x(t), y(t))$ with $\deg(x) = 8$).

Likewise if $p \equiv -1 \bmod 6$, can use $y^2 + y = x^3 + t^{p^e+1} + 1$ to get rank $2(p^e - 1)$ [Shioda].

π. Elliptic K3 surfaces

Sorry, not nearly enough time to say much here, even if the whole hour aimed towards that (cf. arXiv:0709.2908). However, I might at least say a bit more about K3's $\mathcal{X}$ with more than one elliptic fibration. Each fibration yields an even pos.-def. lattice L, but $L \oplus \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right) \cong \mathrm{NS}(\mathcal{X})\langle -1\rangle$. So, L in the same "genus" of quadratic forms (in particular, all have the same discriminant). Now back to our example of

$$E : Y^2 = X^3 + \alpha t^4 X + (\beta_+ t^7 + \beta t^6 + \beta_- t^5)$$

This has reducible fibers of type II^* at $t = 0$ and $t = \infty$; while X/t^2 has one reducible fiber at ∞, of type I^*_{12}, and also has a 2-torsion point. The corresponding L are E_8^2 and D_{16}^+, the two even unimodular lattices of rank 16. [See Coxeter diagram on blackboard.] It is often useful to start from one elliptic fibration that's easy to compute, and then switch fibrations to reach the one we want (with desired torsion and rank). But that's a longer story. . .

T H E E N D