Hyperelliptic Curves

Benjamin Smith

Team **GRACE**

INRIA Saclay–Île-de-France

Laboratoire d'Informatique de l'École polytechnique (LIX)
smith@lix.polytechnique.fr

ECC Summer School 2015

Disclaimer

So far you've seen elliptic curves
from both a low-level, implementation point of view
and a high-level, theoretical point of view.

I'll try to take a "middlebrow" point of view.

(I can't promise we'll have the same idea of where "middle" is, though.)

We work over a perfect field $\Bbbk$.

Perfect?!

- Every irred. poly. over $\Bbbk$ has distinct roots in $\overline{\Bbbk}$
- *Equivalently*: Either $\mathrm{char}(\Bbbk) = 0$, or $\mathrm{char}(\Bbbk) = p$ and the Frobenius $\alpha \mapsto \alpha^p$ is an automorphism.

1. Finite fields: $\Bbbk = \mathbb{F}_q$ *(what we're really interested in)*
2. Characteristic 0: $\Bbbk = \mathbb{Q}, \mathbb{Q}(\sqrt{13}), \mathbb{Q}(t), \mathbb{Q}_p, \mathbb{R}, \mathbb{C}, \ldots$
3. ...But not (e.g.) $\Bbbk = \mathbb{F}_q(t)$
   (because then weird stuff happens with $t^{1/p}$, etc.)

Something (a point, a set, a curve, a function) is *defined over* $\mathbb{k}$ if it is fixed by $\mathrm{Gal}(\overline{\mathbb{k}}/\mathbb{k})$.

If X is a thing, then $X(\mathbb{k})$ denotes its elements/points defined over $\mathbb{k}$.

If $\mathbb{k} = \mathbb{F}_q$, then the objects defined over $\mathbb{F}_q$ are those fixed by/commuting with the q-power Frobenius.

From elliptic to hyperelliptic curves

We've considered cryptosystems built from elliptic curves.

But what's so special about **elliptic** *curves?*

Today: $\mathcal{X}$ denotes an *algebraic curve* over $\mathbb{k}$.

Examples:

- $\mathcal{X} = \mathbb{P}^1 =$ a line
- $\mathcal{X} =$ an elliptic curve $\mathcal{E} : y^2 = x^3 + Ax + B$
- $\mathcal{X} : y^2 = f(x)$ with $\deg f > 4$ (hyperelliptic curves)
- ...More generally, a plane curve $\mathcal{X} : F(x, y) = 0$ in $\mathbb{A}^2$

Hyperelliptic Curves

$$\mathcal{X} : y^2 = f(x) = x^d + \cdots$$

with f squarefree, of degree $d > 4$.
(NB: $d = 1, 2 \implies$ conics; $d = 3, 4 \implies$ elliptic.)

Hyperelliptic involution:

$$\iota : (x, y) \longmapsto (x, -y) \ .$$

d odd $\implies$ one point ∞ at infinity.
d even $\implies$ two points ∞_+, ∞_- at infinity.

Key: $P \mapsto x(P)$ defines a double cover $\mathcal{X} \to \mathcal{X}/\langle\iota\rangle \cong \mathbb{P}^1$.

The function field

If $\mathcal{X} : F(x, y) = 0$ is a plane curve over $\mathbb{k}$, then its function field is

$$\mathbb{k}(\mathcal{X}) = \mathbb{k}(x)[y]/(F(x, y)) .$$

Its elements are rational fractions in x and y, modulo the curve equation $F(x, y) = 0$.

For more general curves: $\mathbb{k}(\mathcal{X}) :=$ fraction field of the coordinate ring.

Zeroes and Poles

Rational functions on $\mathcal{X}$ have *poles* and *zeroes*:

The **zeroes** of f are the points P on $\mathcal{X}$ where $f(P) = 0$.

The **poles** of f are the points P on $\mathcal{X}$ where $f(P) = \infty$.

Note: (zeroes and poles can occur with multiplicity > 1.)

Theorem

If f is a nonzero function in $\overline{\mathbb{k}}(\mathcal{X})$, then

1. *f has only finitely many zeroes and poles, and*
2. *counted with multiplicity, $\#\mathrm{zeroes}(f) = \#\mathrm{poles}(f)$.*

Orders of vanishing

Let f be a nonzero function on $\mathcal{X}$.

We define $\mathrm{ord}_P(f)$ to be the *order of vanishing* of f at P:

- $\mathrm{ord}_P(f) := n$ if f has a zero of multiplicity n at P
- $\mathrm{ord}_P(f) := -n$ if f has a pole of multiplicity n at P
- $\mathrm{ord}_P(f) := 0$ otherwise.

Useful rules:

- $\mathrm{ord}_P(fg) = \mathrm{ord}_P(f) + \mathrm{ord}_P(g)$ for all f, g, P
- $\mathrm{ord}_P(f/g) = \mathrm{ord}_P(f) - \mathrm{ord}_P(g)$ for all f, g, P
- $\mathrm{ord}_P(\alpha) = 0$ for all constants $\alpha \neq 0$ in $\overline{\mathbb{k}}$
- $\mathrm{ord}_P(\sum_i \alpha_i x^{a_i} y^{b_i}) = n$
  if the curve $\sum_i \alpha_i x^{a_i} y^{b_i} = 0$ intersects $\mathcal{X}$ n times at P

Principal divisors

Each function $f \neq 0$ on $\mathcal{X}$ has an associated *principal divisor*: that is, a **formal** sum

$$\operatorname{div}(f) = \sum_{P \in \mathcal{X}(\overline{\mathbb{F}}_q)} \operatorname{ord}_P(f)(P) .$$

1. $\operatorname{div}(f) = 0$ if and only if f is constant (in $\overline{\mathbb{k}}_q \setminus \{0\}$);
2. $\operatorname{div}(fg) = \operatorname{div}(f) + \operatorname{div}(g)$ and $\operatorname{div}(f/g) = \operatorname{div}(f) - \operatorname{div}(g)$;
3. $\operatorname{div}(f) = \operatorname{div}(g) \iff f = \alpha g$ for some $\alpha \neq 0$ in $\overline{\mathbb{F}}_q$.

Functions are determined by their principal divisors, up to constant factors.

The set of principal divisors is denoted $\mathrm{Prin}(\mathcal{X})$:

$$\mathrm{Prin}(\mathcal{X}) := \left\{\mathrm{div}(f) : f \in \overline{\mathbb{k}}(\mathcal{X})\right\} \ .$$

Since $\mathrm{div}(fg) = \mathrm{div}(f) + \mathrm{div}(g)$, we see that

$\mathrm{Prin}(\mathcal{X})$ is a group.

If you like exact sequences:

$$1 \longrightarrow \overline{\mathbb{k}}^{\times} \longrightarrow \overline{\mathbb{k}}(\mathcal{X})^{\times} \longrightarrow \mathrm{Prin}(\mathcal{X}) \longrightarrow 0 \ .$$

Examples

Consider the elliptic curve $\mathcal{E} : y^2 = x^3 + 1$ over $\mathbb{F}_{13}$.

- $\mathrm{div}(x) = (0, 1) + (0, -1) - 2\infty$;
- $\mathrm{div}(y) = (-1, 0) + (4, 0) + (-3, 0) - 3\infty$;
- $\mathrm{div}(x^2/y) =$ $2(0, -1) + 2(0, 1) - (-1, 0) - (4, 0) - (-3, 0) - \infty$;
- $\mathrm{div}(\frac{x^2-y-1}{xy}) =$ $(0, -1) + (2, 3) + \infty - (0, 1) - (-3, 0) - (4, 0)$.

More generally:

If $f(x, y) = 0$ is the line through P and Q, then $\mathrm{div}(f) = P + Q + (\ominus(P \oplus Q)) - 3\infty$.

General divisors

Divisors on $\mathcal{X}$ are **formal sums** of points in $\mathcal{X}(\overline{\mathbb{k}})$ with *arbitrary* coefficients in $\mathbb{Z}$;

We define the (free abelian, infinitely generated) group

$$\mathrm{Div}(\mathcal{X}) := \Big\{ \sum_{P \in \mathcal{X}(\overline{\mathbb{F}}_q)} n_P(P) \Big\},$$

with the n_P in $\mathbb{Z}$, and only finitely many $n_P \neq 0$.

Observe that $\mathrm{Prin}(\mathcal{X}) \subset \mathrm{Div}(\mathcal{X})$.

The Picard group

The divisor group $\mathrm{Div}(\mathcal{X})$ is way too big,
and doesn't tell us anything about the geometry of $\mathcal{X}$.

We work with the quotient

$$\mathrm{Pic}(\mathcal{X}) := \mathrm{Div}(\mathcal{X})/\mathrm{Prin}(\mathcal{X}) .$$

Elements are *divisor classes*:

$$[D] = \{D + \mathrm{div}(f) : f \in \overline{\mathbb{k}}\} .$$

Degree

We have a **degree** homomorphism $\deg : \mathrm{Div}(\mathcal{X}) \to \mathbb{Z}$,

$$\deg\Big(\sum_P n_P(P)\Big) = \sum_P n_P \ .$$

Its kernel is a subgroup of $\mathrm{Div}(\mathcal{X})$, denoted $\mathrm{Div}^0(\mathcal{X})$:

$$\mathrm{Div}^0(\mathcal{X}) := \ker\deg = \{D \in \mathrm{Div}(\mathcal{X}) : \deg(D) = 0\} \subset \mathrm{Div}(\mathcal{X}) \ .$$

Every function has the same number of zeroes and poles, so

$$\mathrm{Prin}(\mathcal{X}) \subset \mathrm{Div}^0(\mathcal{X}) \quad \text{and} \quad \mathrm{Prin}(\mathcal{X})(\Bbbk) \subset \mathrm{Div}^0(\mathcal{X})(\Bbbk) \ .$$

This inclusion is strict for almost all curves:
not every divisor of degree zero is principal!

Why are they called divisors?

Idea: degree-0 divisors are "parts of functions".

Example: Consider $\mathcal{E} : y^2 = x^3 + 1$. The divisors

$$D_1 = (0, 1) - \infty \quad \text{and} \quad D_2 = (0, -1) - \infty$$

are both in $\mathrm{Div}^0(\mathcal{E})$. Neither is principal, but

$$D_1 + D_2 = \mathrm{div}(x) .$$

So we can view D_1 and D_2 as being "parts" (or even "factors") of the function x...

Degrees of divisor classes

deg is well-defined on divisor classes:

$$\begin{aligned} \deg : \mathrm{Pic}(\mathcal{X}) &\longrightarrow \mathbb{Z} \\ [D] &\longmapsto \deg(D) \end{aligned}$$

(since $\deg(\mathrm{div}(f)) = 0$ for all f).

$\Longrightarrow$ $\mathrm{Div}^0(\mathcal{X})$ splits up into divisor classes: we set

$$\begin{aligned} \mathrm{Pic}^0(\mathcal{X}) &:= \ker(\deg : \mathrm{Pic}(\mathcal{X}) \to \mathbb{Z}) \\ &= \mathrm{Div}^0(\mathcal{X})/\mathrm{Prin}(\mathcal{X}) . \end{aligned}$$

The map $D \mapsto (D - \deg(D)\infty, \deg(D))$ defines isomorphisms

$$\mathrm{Div}(\mathcal{X}) \xleftrightarrow{\cong} \mathrm{Div}^0(\mathcal{X}) \times \mathbb{Z}$$

$$\mathrm{Pic}(\mathcal{X}) \xleftrightarrow{\cong} \mathrm{Pic}^0(\mathcal{X}) \times \mathbb{Z} \ .$$

The "interesting" stuff all happens in $\mathrm{Pic}^0(\mathcal{X})$.

In fact, $\mathrm{Pic}^0(\mathcal{X})$ has the structure of an *abelian variety*: a geometric object defined by polynomial equations in projective coordinates, with a polynomial group law.

(Stop and think about what this means for a minute: in some weird universe, divisor classes are defined by tuples of coordinates, and addition of divisor classes modulo linear equivalence is defined by polynomial formulæ in those coordinates!)

Differentials

Differentials on $\mathcal{X}$ look like $g\,df$, where g and f are in $\Bbbk(\mathcal{X})$,

with $\quad g_1 df_1 = g_2 df_2 \iff \dfrac{g_2}{g_1} = \dfrac{df_1}{df_2} \quad$ ($\leftarrow$ usual derivative) .

Differentials obey the usual product rule: $d(fg) = fdg + gdf$.
Also: $d(\alpha f + \beta g) = \alpha df + \beta dg$ and $d\alpha = 0$ for α, β in $\overline{\Bbbk}$.

For example: on $\mathcal{E} : y^2 = x^3 + 1$, we have

$$2ydy = 3x^2dx$$

Differentials are not functions on $\mathcal{X}$:
they give linear functions on the tangent spaces of $\mathcal{X}$.

The space of differentials

The differentials on $\mathcal{X}$ form a one-dimensional $\overline{\mathbb{k}}(\mathcal{X})$-vector space, $\Omega(\mathcal{X})$.

That is: if we fix some differential dx, then every other differential in $\Omega(\mathcal{X})$ is equal to $f dx$ for some function f.

On the other hand:
$\Omega(\mathcal{X})$ is an infinite-dimensional $\overline{\mathbb{k}}$-vector space.

Divisors of differentials

Differentials have divisors!

First, for each point P of $\mathcal{X}$, we fix a *local parameter* t_P near P on $\mathcal{X}$: ie any function with a simple zero at P.

If ω is a differential then ω/dt_P is a function, so we set

$$\operatorname{ord}_P(\omega) := \operatorname{ord}_P(\omega/dt_P)$$

(amazingly, $\operatorname{ord}_P(\omega)$ is independent of choice of t_P) and

$$\operatorname{div}(\omega) := \sum_{P \in \mathcal{X}} \operatorname{ord}_P(\omega) \ .$$

Example on an elliptic curve

What is the divisor of dx on an elliptic curve $\mathcal{E} : y^2 = f(x)$?

At points (α, β) where $\beta \neq 0$, we can use $t_{(\alpha,\beta)} = x - \alpha$:

$$\mathrm{ord}_{(\alpha,\beta)}(dx) = \mathrm{ord}_{(\alpha,\beta)}\left(\frac{dx}{d(x-\alpha)}\right) = \mathrm{ord}_{(\alpha,\beta)}(1) = 0 \ .$$

If $\beta = 0$ then $x - \alpha$ is not a local parameter at $(\alpha, 0)$ (it has a double zero), but we can use $t_{(\alpha,0)} = y$; hence

$$\mathrm{ord}_{(\alpha,0)}(dx) = \mathrm{ord}_{(\alpha,0)}\left(\frac{dx}{dy}\right) = \mathrm{ord}_{(\alpha,0)}\left(\frac{2y}{f'(x)}\right) = 1 \ .$$

At infinity: we can take $t_\infty = x/y$, so

$$\mathrm{ord}_\infty(x) = \mathrm{ord}_\infty\left(\frac{dx}{d(x/y)}\right) = \mathrm{ord}_\infty\left(\frac{yf'(x)}{f'(x) - 2x}\right) = -3 \ .$$

Canonical divisors

$\operatorname{div}(f\omega) = \operatorname{div}(\omega) + \operatorname{div}(f)$ for all $f \in \overline{\mathbb{k}}(\mathcal{X})$, $\omega \in \Omega(\mathcal{X})$,
so the divisors of differentials on $\mathcal{X}$ are
all in the same divisor class,
which we call the *canonical* class $[K]$.

Any divisor in $[K]$ is called a *canonical divisor*.

On $\mathcal{H} : y^2 = f(x) = \prod_{i=1}^{d}(x - \alpha_i)$, we have

$$K = \operatorname{div}(dx) = \begin{cases} \sum_{i=1}^{d}(\alpha_i, 0) - 3\infty & d \text{ odd} \\ \sum_{i=1}^{d}(\alpha_i, 0) - 2(\infty_+ + \infty_-) & d \text{ even} \end{cases}$$

Nonconstant differentials with no poles

So: if $y^2 = f(x) = (x-\alpha_1)(x-\alpha_2)(x-\alpha_3)$, then

$$\mathrm{div}(dx) = (\alpha_1, 0) + (\alpha_2, 0) + (\alpha_3, 0) - 3\infty \ .$$

Notice that $\mathrm{div}(y) = \mathrm{div}(dx)$, so

$$\mathrm{div}(\frac{dx}{y}) = 0$$

—that is, the differential dx/y is a *nonconstant* differential with no poles (or zeroes!).

Regular differentials

We call differentials with no poles *regular*.

The regular differentials on $\mathcal{X}$ form a (finite-dimensional) $\Bbbk$-vector space

$$\Omega^1(\mathcal{X}) = \{\omega \in \Omega(\mathcal{X}) : \omega \text{ is regular}\} .$$

The *genus* of $\mathcal{X}$ is defined to be the dimension of $\Omega^1(\mathcal{X})$.

Genus of hyperelliptic curves

For hyperelliptic curves

$$\mathcal{X} : y^2 = f(x) = x^d + \cdots ,$$

we have

$$\Omega^1(\mathcal{X}) = \left\langle \frac{dx}{y}, \frac{xdx}{y}, \ldots, \frac{x^{\lfloor (d-1)/2-1 \rfloor}dx}{y} \right\rangle ,$$

so

$$g(\mathcal{X}) = \left\lfloor \frac{d-1}{2} \right\rfloor .$$

Explicit regular differentials

More generally, if $\mathcal{X}/\Bbbk$ is a nonsingular plane curve of genus g defined by

$$\mathcal{X} : F(x, y) = 0 \ ,$$

then its regular differentials are

$$\Omega^1(\mathcal{X}) = \left\langle \frac{x^i}{(\partial F/\partial y)(x, y)} dx \right\rangle_{i=0}^{g-1} .$$

For any curve $\mathcal{X}$, we have $\deg(K) = 2g - 2$.

Anomalous elliptic curves

Let's use differentials for something fun.
DLPs in the additive group are really fast:
they're just (modular) division.

When can we map an ECDLP instance into $(\mathbb{F}_p, +)$?

A homomorphism $\mathcal{E}(\mathbb{F}_p) \longrightarrow (\overline{\mathbb{F}}_p, +)$
can only be nontrivial if $p \mid \#\mathcal{E}(\mathbb{F}_p)$,
which (by Hasse) can only happen if $\#\mathcal{E}(\mathbb{F}_p) = p$.

We call these trace-1 curves **anomalous curves**.

Homomorphisms into the additive group

Suppose $\mathcal{E}$ is defined over $\mathbb{F}_p$, and that $\#\mathcal{E}(\mathbb{F}_p) = p$.

Several approaches to mapping $\mathcal{E}(\mathbb{F}_p)$ into $(\mathbb{F}_p, +)$
(Semaev, Smart, Araki–Satoh, Rück...)

Recall: $\dim \Omega^1(\mathcal{E}) = 1$, so $\Omega^1(\mathcal{E}) = (\mathbb{F}_p, +)$.

We will define a homomorphism

$$\mathcal{E}(\mathbb{F}_p) \longrightarrow \Omega^1(\mathcal{E}) \cong (\mathbb{F}_p, +)$$

using an additive version of the Tate pairing.

Suppose $\#\mathcal{E}(\mathbb{F}_p) = p$. If P is in $\mathcal{E}(\mathbb{F}_p)$ then $[p]P = 0$, so

$$p(P - \infty) = \mathrm{div}(f_P)$$

for some f_P in $\mathbb{F}_p(\mathcal{E})$ *(a Miller function!)*

Serre: the differential $\frac{df_P}{f_P}$ is regular at ∞.

Expand df_P/f_P at ∞ with local parameter $t = \frac{x}{y}$:

$$\frac{df_P}{f_P} = (a_0 + a_1 t + a_2 t^2 + \cdots)dt$$

df_P/f_P (and hence the a_i) depends *only* on P.

Product rule for differentials + Algebra of Miller functions $\Longrightarrow$
$P \longmapsto df_P/f_P \longmapsto a_0$ is a homomorphism $\mathcal{E}(\mathbb{F}_p) \to \Omega^1(\mathcal{E}) \to (\mathbb{F}_p, +)$!

Solving DLPs on anomalous curves

To solve a DLP instance $Q = [m]P$ on an anomalous curve $\mathcal{E}/\mathbb{F}_p$:

1. Compute $a_0(P)$ and $a_0(Q)$ using Miller loops
   Don't compute f_P, f_Q: as in pairing computation, build up the a_0 values using double-and-add loops
2. Then

$$m \equiv a_0(Q)/a_0(P) \pmod{p} .$$

The number of $\mathcal{E}(\mathbb{F}_p)$-operations is *linear in* $\log p$.

This reduction is easy to implement!
(It's an exercise for Friday afternoon.)

Into space!

Let's get back to functions on $\mathcal{X}$.

Evaluating functions at points maps us from $\mathcal{X}$ to $\mathbb{P}^1$.

Evaluating a collection $\{f_1, \ldots, f_n\}$ of functions gives us a map $P \mapsto (f_1(P) : \cdots : f_n(P) : 1)$ into $\mathbb{P}^n$.

We want to control behaviour at infinity, hence the poles of the f_i.

Riemann–Roch Spaces

A divisor $D = \sum_P n_P P$ is **effective** if all of the $n_P \geq 0$.

We define

$$L(D) := \{f \in \Bbbk(\mathcal{X}) : D + \mathrm{div}(f) \text{ is effective }\} \cup \{0\}$$

...So $L(D)$ consists of the functions whose poles are contained in D.

$L(D_1 + D_2) \supseteq L(D_1)L(D_2)$ for any effective D_1, D_2.

Note: if $\mathcal{X} = \mathbb{P}^1$, then $L(d\infty) = \{\text{polynomials of degree } \leq d\}$.

Dimension of Riemann–Roch Spaces

Fact: $L(D)$ is a finite-dimensional $\mathbb{k}$-vector space.

What is its dimension?

- If $\deg D < 0$, then $D + \mathrm{div}(f)$ can never be effective $\implies \dim L(D) = 0$.
- $L(0) = \mathbb{k}$ (functions with no poles are constant), so $\dim L(0) = 1$.
- More generally, $L(D) = ?$

The Riemann–Roch Theorem

The Riemann–Roch theorem tells us that for any D,

$$\dim L(D) - \dim L(K - D) = \deg D - g + 1 \ .$$

Recall that K is (any) canonical divisor, and

$$L(K - D) \longleftrightarrow \left\{ \omega \in \Omega^1(\mathcal{X}) : \omega = 0 \text{ on } D \right\} \ .$$

In particular, for large enough D, we have $L(K - D) = 0$ and hence $\dim L(D) = \deg D - g + 1$.

Weierstrass models of elliptic curves

Suppose $\mathcal{E}$ is an **abstract** elliptic curve over $\mathbb{k}$, and let $\mathcal{O} \in \mathcal{E}(\mathbb{k})$.

We have $K = 0$, so R–R gives $\dim L(D) = \deg D$ for effective D.

- $L(\mathcal{O}) = \mathbb{k} = \langle 1 \rangle$ (constants)
- $\dim L(2\mathcal{O}) = 2 \implies L(2\mathcal{O}) = \langle 1, x \rangle$ for some x
- $\dim L(3\mathcal{O}) = 3 \implies L(3\mathcal{O}) = \langle 1, x, y \rangle$ for some y
- $L(4\mathcal{O}) = \langle 1, x, x^2, y \rangle$
- $L(5\mathcal{O}) = \langle 1, x, x^2, y, xy \rangle$
- $L(6\mathcal{O}) = \langle 1, x, x^2, x^3, y, xy, y^2 \rangle$, but $\dim L(6\mathcal{O}) = 6$:
  so must have a nontrivial linear relation between the 7 functions

$\implies$ Weierstrass equation $y^2 + a_1 xy + a_3 y = a_0 x^3 + a_2 x^2 + a_4 x + a_6$.

$L(3\mathcal{O})$ gives us an embedding $\mathcal{E} \to \mathbb{P}^2 = \mathbb{P}(L(3\mathcal{O}))$
defined by $P \longmapsto (x(P) : y(P) : 1)$, mapping $\mathcal{O} \mapsto \infty = (0 : 1 : 0)$.

Application: canonical models for genus 2 curves

Suppose $\mathcal{X}$ is a curve of genus 2.

- We have $\deg K = 2g - 2 = 2$,
  so $L(-nK) = 0$ for $n > 1$.
- Apply R–R to $D = 0 \implies \dim L(K) = 2$,
  so $L(K) = \langle 1, x \rangle$ for some x.
- Apply R–R to $D = nK$, $n > 1$: $\dim L(nK) = 2n - 1$ for $n > 1$.
- $L(2K) \supseteq \langle 1, x, x^2 \rangle$ but $\dim L(2K) = 3$,
  so $L(2K) = \langle 1, x, x^2 \rangle$.
- $L(3K) \supseteq \langle 1, x, x^2, x^3 \rangle$ but $\dim L(3K) = 5$,
  so $L(3K) = \langle 1, x, x^2, x^3, y \rangle$ for some new y
- ...$L(4K) = \langle 1, x, x^2, x^3, x^4, y, xy \rangle$
- ...$L(5K) = \langle 1, x, x^2, x^3, x^4, x^5, y, xy, x^2y \rangle$

...Every genus 2 curve is hyperelliptic

Now $L(6K) \supseteq \langle 1, x, x^2, x^3, x^4, x^5, x^6, y, xy, x^2y, x^3y, y^2 \rangle$,
but R–R says $\dim L(6K) = 11$, so
there is a nontrivial $\Bbbk$-linear relation between the 12 functions:

$$y^2 + \sum_{i=0}^{3}(a_i x^i y) = \sum_{i=0}^{6} b_i x^i \quad \text{with the } a_i, b_i \in \Bbbk \ .$$

$\mathrm{char}(\Bbbk) \neq 2$: replace y with $y - \frac{1}{2}\sum_{i=0}^{3} a_i x^i$ to get $y^2 = \sum_{i=0}^{6} f_i x^i$.

Now $P \mapsto (x(P), y(P))$ defines a map from $\mathcal{X}$ into the plane;
its image is the hyperelliptic curve

$$\mathcal{X} : y^2 = f(x) = \sum_{i=0}^{6} f_i x^i \ .$$

Hyperelliptic Jacobians

Benjamin Smith

Team **GRACE**

INRIA Saclay–Île-de-France

Laboratoire d'Informatique de l'École polytechnique (LIX)

`smith@lix.polytechnique.fr`

ECC Summer School 2015

Hyperelliptic Jacobians

Suppose $\mathcal{X} : y^2 = f(x)$ is hyperelliptic of genus $g > 1$.

In what follows, we suppose f has odd degree,
so $\mathcal{X}$ has a single point ∞ at infinity.

Even degree case is (only) slightly more complicated.

Our mission: to define a compact (and algebraic) representation for $\mathrm{Pic}^0(\mathcal{X})$.

Reduced representatives for classes

If $[D]$ is in $\mathrm{Pic}^0(\mathcal{X})$,
then $[D]$ has a unique *reduced representative*:

$$[D] = [P_1 + \cdots + P_r - r\infty]$$

for some $P_1, \ldots, P_r \in \mathcal{X}$ depending on $[D]$ (not D) such that

- $P_i \neq \infty$ and $P_i \neq \iota(P_j)$ for $i \neq j$ (semi-reducedness)
- $r \leq g$ (reducedness)

$$[D] \in \mathrm{Pic}^0(\mathcal{X})(\Bbbk) \iff P_1 + \cdots + P_r \in \mathrm{Div}(\mathcal{X})(\Bbbk)$$

Note: the individual P_i need not be in $\mathcal{X}(\Bbbk)$!

Why?

Because of Riemann–Roch (*quelle surprise*).

If $[D]$ is in $\mathrm{Pic}^0(\mathcal{X})$, then applying R–R to $D + g\infty$ yields a function f such that $D + g\infty + \mathrm{div}(f) = D'$ is effective; so $[D' - g\infty] = [D]$ with $\deg D' = g$.

$D' - g\infty$ is *almost* a reduced representative: it remains to remove any $P + \iota(P) - 2\infty$ from D'.

The Mumford representation

Suppose we have a class $[D]$ in $\mathrm{Pic}^0(\mathcal{X})(\Bbbk)$, with reduced representative

$$D = P_1 + \cdots + P_r - r\infty \in \mathrm{Div}^0(\mathcal{X})(\Bbbk) .$$

The **Mumford representation** of $[D]$ is the (unique) pair of polynomials $\langle a(x), b(x)\rangle$ in $\Bbbk[x]$ such that

- $a(x) = \prod_{i=1}^{r}(x - x(P_i))$, and
- $b(x(P_i)) = y(P_i)$ for $1 \leq i \leq r$;

so for each of the x-coordinates appearing as a root of a, b gives the corresponding y-coordinate.

If necessary, compute b by Lagrange interpolation.

The Mumford representation

If $\langle a(x), b(x)\rangle$ represents a class on $\mathcal{X} : y^2 = f(x)$, then

1. a is monic of degree $r \leq g$, and
2. b satisfies $\deg b < r$ and $b^2 \equiv f \pmod{a}$.

Theorem: Any pair $\langle a(x), b(x)\rangle$ in $\mathbb{k}[x]^2$ satisfying these conditions represents a divisor class in $\mathrm{Pic}^0(\mathcal{X})(\mathbb{k})$.

$\implies$ identify divisor classes with Mumford reps of their reduced representatives: we simply write $[D] = \langle a, b\rangle$.

We associate $\langle a(x), b(x)\rangle$ *with the ideal* $(a(x), y - b(x))$.

Hyperelliptic Jacobians

We can collect the Mumford representations by degree $0 \leq d \leq g$:

$$M_d := \{\langle a, b\rangle : \deg(b) < \deg(a) = d, b^2 \equiv f \pmod{a}\} \ .$$

We view the coefficients of $a(x) = x^d + a_{d-1}x^{d-1} + \cdots + a_1 x + a_0$ and $b(x) = b_{d-1}x^{d-1} + \cdots + b_0$ as coordinates on $\mathbb{A}^{2d}$.

$b^2 \pmod{a}$ and $f \pmod{a}$ are polynomials of degree $d-1$ in $\Bbbk[a_i, b_i][x]$; the vanishing of their coefficients defines d independent equations in the $2d$ coordinates, cutting out M_d as a d-dimensional subvariety in $\mathbb{A}^{2d}$.

Observe: M_0 is a point; M_1 is an affine copy of $\mathcal{X}$; and $\#M_d(\mathbb{F}_q) = O(q^d)$ for $0 \leq d \leq g$.

The Jacobian

Glueing together $M_0, \ldots, M_g$, we give $\mathrm{Pic}^0(\mathcal{X})$ the structure of a g-dimensional algebraic variety $\mathcal{J}_\mathcal{X}$, called the **Jacobian**.

Over $\mathbb{F}_q$, we have $\#\mathcal{J}_\mathcal{X} = O(q^g)$.
(more precision later)

We want an expression of the group law on $\mathcal{J}_\mathcal{X}$ in terms of its coordinates;
Cantor's algorithm does this using an explicit form of *(guess what?)* Riemann–Roch *(quelle surprise!)*.

Cantor's algorithm: addition on $\mathcal{J}_{\mathcal{X}}$

Input: Reduced divisors $D_1 = \langle a_1, b_1 \rangle$ and $D_2 = \langle a_2, b_2 \rangle$ on $\mathcal{X}$.

Output: A reduced $D_3 = \langle a_3, b_3 \rangle$ s.t. $[D_3] = [D_1 + D_2]$ in $\mathrm{Pic}^0(\mathcal{X})$.

1. $(d, u_1, u_2, u_3) := \mathrm{XGCD}(a_1, a_2, b_1 + b_2)$
   $//(\text{so } d = \gcd(a_1, a_2, b_1 + b_2) = u_1 a_1 + u_2 a_2 + u_3(b_1 + b_2))$.
2. Set $a_3 := a_1 a_2 / d^2$;
3. Set $b_3 := b_1 + (u_1 a_1 (b_2 - b_1) + u_3(f - b_1^2))/d \pmod{a_3}$;
4. If $\deg a_3 \leq g$ then go to Step 9;
5. Set $\tilde{a}_3 := a_3$ and $\tilde{b}_3 := b_3$;
6. Set $a_3 := (f - b_3^2)/a_3$;
7. Let $(Q, b_3) := \mathrm{Quotrem}(-b_3, a_3)$;
8. While $\deg a_3 > g$
   - 8a Set $t := \tilde{a}_3 + Q(b_3 - \tilde{b}_3)$;
   - 8b Set $\tilde{b}_3 := b_3$, $\tilde{a}_3 = a_3$, and $a_3 := t$;
   - 8c Let $(Q, b_3) := \mathrm{Quotrem}(-b_3, a_3)$;
9. Return $\langle a_3, b_3 \rangle$.

How does Cantor reduction work?

Suppose we want to add the Mumford/reduced representatives

$$\langle a_1, b_1 \rangle \longleftrightarrow D_1 = \sum_{i=1}^{r} P_i - r\infty$$

$$\langle a_2, b_2 \rangle \longleftrightarrow D_2 = \sum_{i=1}^{s} Q_i - s\infty$$

- Step 1: $d(x(P_i)) = 0$ iff $P_i = \iota(Q_j)$ for some j
- Steps 2, 3: sum D_1 and D_2, remove contribution of d $\longrightarrow$ pre-reduced D_3 such that $[D_3] = [D_1 + D_2]$
- Loop: reduces degree of the representative until reduced.
- Exercise: *how many steps until the result is reduced?*

Cryptographic questions

We've seen that hyperelliptic curves of genus g over $\mathbb{F}_q$ yield algebraic groups with $O(q^g)$ elements and a conveniently computable group law.

- How can we compute $\#\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q)$?
- How hard is the DLP in $\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q)$?
- How can we construct strong and fast Jacobians?
- How efficient are hyperelliptic cryptosystems, and how do they compare with elliptic cryptosystems?
- Do hyperelliptic curves have destructive applications?

Facts about Jacobians

What are the analogues of the elliptic curve group structure theorems for $\mathcal{J}_{\mathcal{X}}$?

- $\mathcal{J}_{\mathcal{X}}[\ell^n](\overline{\mathbb{k}}) \cong (\mathbb{Z}/\ell^n\mathbb{Z})^{2g}$ for $\ell \neq \mathrm{char}(\mathbb{k})$
- $\mathcal{J}_{\mathcal{X}}[p^n](\overline{\mathbb{k}}) \cong (\mathbb{Z}/p^n)^r$ for some $0 \leq r \leq g$
  (p-rank r is independent of n)
- $\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q) \cong \prod_{i=1}^{2g}(\mathbb{Z}/n_i\mathbb{Z})$ with each $n_{i+1} \mid n_i$

Facts about Jacobians

We have $\#\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q) = \chi_\pi(1)$, where

$$\begin{aligned}\chi_\pi(T) = {} & T^{2g} + a_1 T^{2g-1} + \cdots + a_g T^g \\ & + q a_{g-1} T^{g-1} + \cdots + q^{g-1} a_1 T + q^g\end{aligned}$$

is the characteristic polynomial of Frobenius.

Weil bounds: $(\sqrt{q}-1)^{2g} \le \#\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q) \le (\sqrt{q}+1)^{2g}$

Generically, $\mathrm{End}(\mathcal{J}_{\mathcal{X}})$ is an order in a CM-field of degree $2g$
(a totally imaginary extension of a totally real field of degree g)

Moduli space $\mathcal{H}_g$ (j-invariant analogue) is $(2g-1)$-dimensional
$\implies O(q^{2g-1})$ non-isomorphic $\mathcal{X}$ of genus g over $\mathbb{F}_q$

Point counting on Jacobians

How do we compute $\#\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q)$, where $q = p^n$?
Ultimate goal: polynomial time in $\log p$, n, and g.

- Generic group order methods (eg. BSGS): $\widetilde{O}(q^{g/2})$
  Sutherland's algorithms: faster but still exponential
  Easy to implement, impossible to run on big inputs
- Small p: Kedlaya's algorithm $\widetilde{O}(pg^4n^3)$
  —polynomial in g and n, but exponential in $\log p$.
  (Uses MW cohomology on p-adic differentials)
  Harvey's improvements: $\widetilde{O}(p^{1/2}g^4n^3)$

Point counting on Jacobians: large p

For large p: Pila's generalization of Schoof's algorithm.

- *In theory:* exponential in g, polynomial in $\log p$ and n
- *In practice:* never implemented for $g \geq 3$:
- *General genus 2* over $\mathbb{F}_p$: $\widetilde{O}(\log^8 p)$. Crushingly slow.
  - Gaudry–Schost 2008 record: one CPU-month per 127-bit curve
  - For comparison, equivalent elliptic curve < 10 CPU-seconds
- *Special genus 2* over $\mathbb{F}_p$: $\widetilde{O}(\log^5 p)$.
  - *(2-param. families with efficiently computable "real" endomorphisms)*
  - Gaudry–Kohel–S.: three CPU-hours per 128-bit curve
  - With early abort: practical generation of industrial-sized random cryptographic curves
  - Gaudry–Kohel–S. record: 80 CPU-days per 512-bit curve

Embeddings of Jacobians

The Mumford representation lets us compute with a hyperelliptic Jacobian by dividing it up into affine pieces:

$$\mathcal{J}_{\mathcal{X}} = M_0 \cup M_1 \cup \cdots \cup M_g \ .$$

In fact, $\mathcal{J}_{\mathcal{X}}$ is projective (it's an *abelian variety*) —so what are its projective embeddings?

This is a nontrivial question

$$\mathcal{J}_{\mathcal{X}} = M_0 \cup M_1 \cup \cdots \cup M_g \quad \text{with each } M_i \subset \mathbb{A}^{2i}$$

recalls the usual decomposition $\mathbb{P}^n = \mathbb{A}^0 \cup \mathbb{A}^1 \cup \cdots \cup \mathbb{A}^n$
—but it's not the same thing at all!

As cryptographers, we're used to thinking of projective coordinates as nothing more than convenient denominator elimination, which we carry out by homogenization.

But if you just homogenize Mumford representations, then you get something totally wrong.

The Jacobi intersection model

To create projective embeddings of curves,
we used divisors and Riemann–Roch.

For example: given a point $\mathcal{O}$ on an elliptic $\mathcal{E}$,
we embedded $\mathcal{E}$ in $\mathbb{P}^2 = \mathbb{P}(L(3\mathcal{O})) = \mathbb{P}(\langle x, y, 1\rangle)$.

Alternative embeddings: for example, use $D = 4\mathcal{O}$.

- $L(4\mathcal{O}) = \langle x, y, u, v\rangle$ (because $\dim L(4\mathcal{O}) = \deg(4\mathcal{O}) = 4$);
- $L(8\mathcal{O}) \supseteq L(4\mathcal{O})^2 = \langle x^2, xy, xu, xv, y^2, yu, yv, u^2, uv, v^2\rangle$
- but $\dim L(8\mathcal{O}) = 8 \implies$ 2 quadratic relations in x, y, u, v.
- $\implies$ the *Jacobi intersection model* of $\mathcal{E}$:

$$\mathcal{E} : F_2(x, y, u, v) = G_2(x, y, u, v) = 0 \quad \subset \mathbb{P}^3 = \mathbb{P}(L(4\mathcal{O})) .$$

Theta

So, if $\mathcal{E}$ is an elliptic curve and $\mathcal{O}$ is a point on $\mathcal{E}$, then:

- $L(3\mathcal{O})$ embeds $\mathcal{E}$ in $\mathbb{P}^2$ with one cubic equation;
- $L(4\mathcal{O})$ embeds $\mathcal{E}$ in $\mathbb{P}^3$ with two quadratic equations.

What are the hyperelliptic analogues?

We need a divisor on $\mathcal{J}_{\mathcal{X}}$ to take the place of $\mathcal{O}$ on $\mathcal{E}$:

$$\Theta := \{[P_1 + \cdots + P_{g-1} - (g-1)\infty] : P_1, \ldots, P_{g-1} \in \mathcal{X}(\overline{\mathbb{k}})\}$$

(Note: $\Theta = M_0 \cup \cdots \cup M_{g-1}$).

Projective embeddings of $\mathcal{J}_{\mathcal{X}}$

$$\Theta := \{[P_1 + \cdots + P_{g-1} - (g-1)\infty] : P_i \in \mathcal{X}(\overline{\mathbb{k}})\}$$

We have $\dim L(n\Theta) = n^g$, so

- $L(3\Theta)$ embeds $\mathcal{J}_{\mathcal{X}}$ in $\mathbb{P}^{3^g-1}$
- $L(4\Theta)$ embeds $\mathcal{J}_{\mathcal{X}}$ in $\mathbb{P}^{4^g-1}$.

The dimension of the space is exponential in g
(and so is the number of equations!)

Generally, $\mathcal{J}_{\mathcal{X}}$ does not embed in a smaller projective space than $\mathbb{P}^{3^g-1}$!

Projective embeddings of $\mathcal{J}_{\mathcal{X}}$ *for* $g = 2$

For $g = 2$: $\mathcal{J}_{\mathcal{X}}$ is a surface, Θ is a copy of $\mathcal{X}$ inside $\mathcal{J}_{\mathcal{X}}$.

- $L(3\Theta)$ gives the "Grant" embedding in $\mathbb{P}^8$ with 10 quadratic and 3 cubic equations.
- $L(4\Theta)$ gives the "Flynn" embedding in $\mathbb{P}^{15}$ with 72 quadratic equations.
- $\mathcal{J}_{\mathcal{X}}$ *never* embeds in $\mathbb{P}^3$.
- $\mathcal{J}_{\mathcal{X}}$ embeds in $\mathbb{P}^4$ if and only if $\mathrm{End}(\mathcal{J}_{\mathcal{X}})$ contains $\mathbb{Z}[(1+\sqrt{5})/2]$ (!! ...Horrocks–Mumford, etc.)

The future of Jacobian arithmetic

Mumford representations are convenient, but Cantor's algorithm does not have a uniform execution path $\Longrightarrow$ vulnerable to simple side-channel attacks.

The existing (smooth) projective embeddings are fine for one-off computations and experiments, but they are totally unsuitable for cryptographic applications.

Deriving convenient, compact models with efficient and uniform group laws is a serious open problem.

Get involved (and tell us about it at ECC next year)!

The DLP in hyperelliptic Jacobians

What about the DLP in hyperelliptic Jacobians?

We have $N = \#\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q) = O(q^g)$.

Cryptographic contexts: N is prime (or almost).
Gold standard: $\widetilde{O}(\sqrt{N}) = \widetilde{O}(q^{g/2})$ operations in $\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q)$
(Pollard/BSGS generic group methods).

If the DLP is easier than this, then we are better off using an elliptic curve over $\mathbb{F}_p$ with $p \sim q^g$.

(Oversimplified) Index Calculus

Suppose we want to solve a DLP $D_1 = [m]D_2$
in a cyclic group $\mathcal{G} \cong \mathbb{Z}/N\mathbb{Z}$.

- Choose a distinguished subset $\mathcal{F} \subset \mathcal{G}$, called a *factor base*.
- Set up a matrix M over $\mathbb{Z}/N\mathbb{Z}$ with a column for each element F_j of the factor base $\mathcal{F}$.
- Generate random combinations $[a_i]D_1 \oplus [b_i]D_2$, and test each one for *smoothness*: if $[a_i]D_1 \oplus [b_i]D_2 = \bigoplus_j [n_j]F_j$, then add a row (n_j) to M.
- Once M has more rows than columns, solve to find a kernel vector (x_i) (such that $(x_i)M = 0$).
- Then $\bigoplus_i [x_i a_i]D_1 \oplus \bigoplus_i [x_i b_i]D_2 = 0$, so $m = -(\sum_i x_i b_i)/(\sum_i x_i a_i) \pmod{N}$.

Hyperelliptic Index Calculus

For basic hyperelliptic index calculus: the factor base

$$\mathcal{F} := M_1(\mathbb{F}_q) = \{\langle x - \alpha, \beta \rangle : \beta^2 = f(\alpha)\}$$

has $O(q)$ elements.

To generate $\mathcal{F}$: iterate over α in $\mathbb{F}_q$,
keep $\langle x - \alpha, \sqrt{f(\alpha)} \rangle$ where the square root is in $\mathbb{F}_q$.

$[D] = \langle a, b \rangle \in \mathcal{J}_{\mathcal{X}}(\mathbb{F}_q)$ is smooth
if a splits completely over $\mathbb{F}_q$
(smoothness testing = polynomial factorization)

Expect: $1/g!$ divisor classes are smooth
$\implies O(g!q)$ divisors to be tested

Index Calculus Complexity (in $\mathbb{F}_q$-ops)

Group ops in $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$ (via Cantor) cost $O(g^2 \log^2 q)$
Degree-g poly factorizations $/\mathbb{F}_q$ cost $O(g^2 \log^3 q)$

Need $O(q)$ relations; each costs
$O(g!(g^2 \log^2 q + g^2 \log^3 q))$ to acquire.

Sparse linear algebra (eg. Lanczos): $O(gq^2)(g \log q)$

Total: $O((g^2 \log^3 q)g!q + (g^2 \log q)q^2)$,
$= \tilde{O}(q^2)$ for fixed g as $q \to \infty$

Small g: index calculus improvements

Harley: use only a small fraction of $\mathcal{F}$.
$\implies$ cost drops from $\widetilde{O}(q^2)$ to $\widetilde{O}(q^{2g/(g+1)})$.

Thériault: single large prime variant
$\implies$ $\tilde{O}(q^{2-4/(2g+1)})$.

Gaudry–Thomé–Thériault–Diem:
double large prime variant $\implies$ $\tilde{O}(q^{2-2/g})$

Genus $\to \infty$ and $q \to \infty$:

$\implies$ $L_{q^g}(1/2, \sqrt{2})$

(But let's be serious: $g \ll \infty$)

Bad news for genus ≥ 3

Observe: $\widetilde{O}(q^{2-2/g})$ is easier than $\widetilde{O}(q^{g/2})$ for $g > 2$.

We can do even better in genus 3 (S., Eurocrypt08):
use an explicit isogeny to move the DLP
into the Jacobian of a *non-hyperelliptic* genus 3 curve
(a smooth plane quartic),
where Diem's plane curve index calculus
solves the DLP in $\widetilde{O}(q)$ group operations.

How many bits for a given security level?

Suppose we want b bits of security
(ie, the attacker must use $\sim 2^b$ operations to solve the DLP).

Curve	$\log q$	element size	$\#\mathcal{J}_{\mathcal{X}}(\mathbb{F}_q)$
Elliptic	$\sim 2b$	$\sim 2b$	$\sim 2b$
Genus 2	$\sim b$	$\sim 2b$	$\sim 2b$
Genus 3	$\sim b$	$\sim 3b$	$\sim 3b$
Genus $g \geq 4$	$\sim \frac{g}{2g-2}b$	$\sim \frac{g^2}{2g-2}b$	$\sim \frac{g^2}{2g-2}b$

- Efficiency is already suboptimal for $g = 3$:
  genus 3 cryptosystems require 50% more space
  than elliptic or genus 2 systems at the same security level.
- Higher genus: even worse!
- $\implies$ Moral: for constructive work, stick to genus 1 and 2.

Restriction of scalars

Suppose $\mathcal{E}$ is defined over an extension field $\mathbb{F}_{q^n}$, $n > 1$.

$\mathcal{E}$ is a one-dimensional object over a degree-n field.

Weil descent is a direct tradeoff of dimension vs degree.

Think of the complex numbers:
We can see $\mathbb{C}$ as the line (one-dimensional)
over a quadratic extension $\mathbb{R}(\sqrt{-1})$,
but we can also visualise it as the real plane $\mathbb{R}^2$.

In the same way: the one-dimensional vector space $\mathbb{F}_{q^n}$
is isomorphic to the n-dimensional vector space $\mathbb{F}_q^n$.

Weil descent

The Weil restriction $\mathcal{W}$ of $\mathcal{E}$ is
an n-dimensional algebraic group over $\mathbb{F}_q$ *(not $\mathbb{F}_{q^n}$)*
whose $\mathbb{F}_q$-points correspond to $\mathbb{F}_{q^n}$-points of $\mathcal{E}$.

The Weil restriction always exists,
and doesn't weaken $\mathcal{E}$ in itself.

But if we're lucky,
we might be able to transform all (or part) of $\mathcal{W}$ into
the Jacobian of a higher-genus curve,
which we can attack using index calculus.

Weil descent of an elliptic curve

Let's try $n = 3$, with $q = 2^e$ for some e and $\mathbb{F}_{q^3} = \mathbb{F}_q[\theta]/(\theta^3 + \theta + 1)$.

$$\mathbb{F}_{q^3} = \left\langle \psi_0 = 1, \psi_1 = \theta^2, \psi_2 = \theta^4 \right\rangle_{\mathbb{F}_q}$$

Any elliptic curve over $\mathbb{F}_{q^3}$ is $\cong$ to one in the form

$$\mathcal{E}/\mathbb{F}_{q^3} : y^2 + xy = x^3 + (b_0\psi_0 + b_1\psi_1 + b_2\psi_2) \ .$$

Equations for Weil restriction $\mathcal{W}$: substitute

$$x = x_0\psi_0 + x_1\psi_1 + x_2\psi_2 \ , \quad y = y_0\psi_0 + y_1\psi_1 + y_2\psi_2 \ ,$$

get 3 equations over $\mathbb{F}_q$ by collecting coefficients of the ψ_i.

Explicit Weil restrictions

So: Weil restriction $\mathcal{W}$ of $\mathcal{E} : y^2 + xy = x^3 + (b_0\psi_0 + b_1\psi_1 + b_2\psi_2)$ is defined in $(x_0, x_1, x_2, y_0, y_1, y_2)$-space by the three equations

$$x_0^3 + x_0^2x_2 + x_0x_1^2 + x_0y_1 + x_0y_2 + x_1^3 + x_1x_2^2 + x_1y_0 + x_1y_2 + x_2^3 + x_2y_0 + x_2y_1$$

$$x_0^3 + x_0^2x_1 + x_0x_1^2 + x_0y_1 + x_0y_2 + x_1^2x_2 + x_1x_2^2 + x_1y_0 + x_1y_1 + x_2^3 + x_2y_0 + x_2y_2 + y_1^2 + y_2^2 + b_2 + b_0$$

$$x_0^2x_1 + x_0^2x_2 + x_0x_1^2 + x_0x_2^2 + x_0y_0 + x_0y_2 + x_1^3 + x_1y_1 + x_1y_2 + x_2^3 + x_2y_0 + x_2y_1 + y_0^2 + y_1^2 + b_2 + b_1$$

To get a curve in $\mathcal{W}$, intersect with (say) $x_0 = u$, $x_1 = u$, $x_2 = u$:

$$\mathcal{C} : \left(y_2^2 + uy_0 = u^3 + b_0\ ,\ y_0^2 + uy_1 = u^3 + b_1\ ,\ y_1^2 + uy_2 = u^3 + b_2\right)$$

Irreducible unless $b_0 = b_1 = b_2$ (so $\beta \in \mathbb{F}_q$). Eliminate y_1, y_2, put $v = y_0$:

$$\mathcal{C} : v^8 + u^7v + u^{12} + u^{10} + u^9 + b_0u^6 + b_2^2u^4 + b_1^4\ .$$

It may not be obvious, but $\mathcal{C}$ is hyperelliptic of genus 3.

Desingularize $\widetilde{\mathcal{C}} \to \mathcal{C} \implies$ explicit isogeny $\Phi : \mathcal{W} \to \mathrm{Jac}(\widetilde{\mathcal{C}})$.

Discrete logarithms on the Weil restriction

Start with a DLP instance in $\mathcal{E}(\mathbb{F}_{q^3})$:

$$Q = (x^Q, y^Q) = [m](x^P, y^P) = [m]P$$

Weil-restricting, we get a DLP instance in $\mathcal{W}(\mathbb{F}_q)$:

$$(x_0^Q, x_1^Q, x_2^Q, y_0^Q, y_1^Q, y_2^Q) = [m](x_0^P, x_1^P, x_2^P, y_0^P, y_1^P, y_2^P)\ ;$$

map through Φ to get a DLP instance in $\mathrm{Jac}(\mathcal{C})$:

$$\left[\sum_{i=1}^{3}(u_i^Q, v_i^Q) - D_0\right] = m\left[\sum_{i=1}^{3}(u_i^P, v_i^P) - D_0\right]$$

Solve DLP instance using index calculus in $\mathcal{J}_{\widetilde{\mathcal{C}}}$ in time $\widetilde{O}(q^{4/3})$:
beats $\widetilde{O}(q^{3/2})$ *using generic algorithms in* $\mathcal{E}(\mathbb{F}_{q^3})$.

Gaudry–Hess–Smart

In more generality:

Theorem (Gaudry–Hess–Smart, 2000)

Let $n \geq 4$ be fixed. Write $q = 2^e$. Then as $e \to \infty$, we can solve the DLP in $\mathcal{E}(\mathbb{F}_{q^n})$ for a significant proportion of all elliptic curves $\mathcal{E}/\mathbb{F}_{q^n}$ in time $O(q^{2+\epsilon})$.

For comparison: generic attacks require time $\widetilde{O}(q^{n/2})$.

Reading guide:

http://www.cs.bris.ac.uk/~nigel/weil_descent.html