Modern Elliptic Curve Cryptography 1

Benjamin Smith

SAC 2020 Summer School // 19/10/2020

Inria + École polytechnique

A black-box group

Let $\mathcal{G} = \langle P \rangle$ be a (fixed, public) cyclic group of order N.

Group operation: $(P, Q) \mapsto P \oplus Q$.

Scalar multiplication:

$$(m, P) \longmapsto [m]P := \underbrace{P \oplus \cdots \oplus P}_{m \text{ copies of } P} .$$

For the moment, we treat $\mathcal{G}$ as a *black-box group*:

- Elements identified with labels / strings of $\log_2 N$ bits
- Group operations: a black box / oracle:
  - *Input* labels corresponding to elements P_1 and P_2
  - *Output* the label corresponding to $P_1 \oplus P_2$
  - *Computation* carried out in polynomial time (in $\log_2 N$)

Scalar multiplication is easy

Theorem: We can compute *any* scalar multiple in $O(\log N)$ $\mathcal{G}$-ops.

Algorithm 1: Classic double-and-add scalar multiplication.

Input: Scalar $m = \sum_{i=0}^{\beta-1} m_i 2^i$ in $[0..N-1]$, element P in $\mathcal{G}$

Output: $[m]P$

$R \leftarrow 0_{\mathcal{G}}$

for $i := \beta - 1$ *down to* 0 **do** `// Loop invariant:` $R = [\lfloor m/2^i \rfloor]P$

  $R \leftarrow [2]R$

  if $m_i = 1$ **then** `// Danger! Branch leaks secret` m_i `to SCA`

    $R \leftarrow R \oplus P$

return R `//` $R = [m]P$

Inverting scalar multiplication is the **Discrete Logarithm Problem** in $\mathcal{G}$:

$$\text{Given } (P, [x]P)\,, \text{compute } x\,.$$

Fact: in any $\mathcal{G}$, we can *always* solve the DLP in time $O(\sqrt{N})$.

- *Shanks' Baby-step giant-step (+ low-memory variants),*
- *Pollard's ρ and Kangaroo (λ)...*

Generic DLP: Shanks' BSGS in $\mathcal{G}$

Algorithm 2: Baby-step giant-step in $\mathcal{G}$

Input: P and Q in $\mathcal{G}$

Output: x such that $Q = [x]P$

```
β ← ⌈√(#G)⌉
(S_i) ← ([i]P : 1 ≤ i ≤ β)
Sort/hash ((S_i, i))_{i=1}^{β}
T ← Q
for j in (1, ..., β) do
    if T = S_i for some i then
        return i − jβ
    T ← T + [β]P
return ⊥                                    // Only if Q ∉ ⟨P⟩
```

The Pohlig–Hellman reduction

The largest prime-order subgroup of $\mathcal{G}$ is all that matters.

Theorem (Pohlig and Hellman)
Suppose we know primes p_i and exponents e_i such that

$$\mathcal{G} \cong \prod_{i=1}^{n} (\mathbb{Z}/p_i^{e_i}\mathbb{Z})$$

(and so $N = \#\mathcal{G} = \prod_{i=1}^{n} p_i^{e_i}$).
Then we can solve the DLP in $\mathcal{G}$ in

$$O(\sum_{i=1}^{n} e_i(\log N + \sqrt{p_i})) \quad \mathcal{G}\text{-operations.}$$

Shoup's theorem

Idea: we want to talk about algorithms that run *independently of the presentation* of a group $\mathcal{G} \cong \mathbb{Z}/N\mathbb{Z}$, treating $\mathcal{G}$ as a **black box group**. To formalize this: consider the set Σ of all *encoding* functions $\sigma : \mathbb{Z}/N\mathbb{Z} \hookrightarrow S$ for some (fixed) $S \subset \{0,1\}^*$.

Encoded group laws: oracles L which, on input $(\sigma(a), \sigma(b), \pm 1)$, output $\sigma(a \pm b)$.

A **generic algorithm** is a randomized algorithm which takes $\sigma \in \Sigma$ and $(\sigma(x_1), \ldots, \sigma(x_r)) \in S^r$ and returns some y in $\mathbb{Z}$.

Theorem (Shoup): Let p be the largest prime divisor of N, and let $\mathcal{A}$ be a generic algorithm making at most t queries to L. If $x \in \mathbb{Z}/N\mathbb{Z}$ and σ are chosen at random, then the probability that $\mathcal{A}(\sigma; (\sigma(1), \sigma(x)))$ returns x is $O(t^2/p)$.

Corollary: For $\mathcal{A}$ to solve the DLP in a group $\mathcal{G} \cong \mathbb{Z}/N\mathbb{Z}$ with probability bounded away from 0 by a constant, it must use $\Omega(p^{1/2})$ group operations.

In an ideal world...

In practice we compute with concrete groups, not abstract black-box groups.

To maximise cryptographic efficiency (ratio: security level / key length),
we need concrete groups that act like black box groups:

- **Order** Prime (or almost-prime) order N
- **Elem. Size** Elements stored in $\sim \log_2 N$ bits each
- **Elem. Ops** Operations computed in $\widetilde{O}(\log_2^c N)$ bit-ops, c small
- **DLP** Best known DLP solutions in $O(\sqrt{N})$ $\mathcal{G}$-ops

Additive groups of finite fields

First attempt at a cryptographic $\mathcal{G}$: prime-order subgroups of $\mathbb{G}_a(\mathbb{F}_q)$ *(the additive group)*.

How do subgroups of $\mathbb{G}_a(\mathbb{F}_q)$ measure up against a black-box group?

Order Automatic: $(\mathbb{F}_p, +)$ is the only prime-order subgroup.

Elem. Size $\log_2 p$ bits (ideal!)

Elem. Ops $\sim \log_2 p$ bit-ops: *very efficient*.

DLP? Solve with the **Euclidean algorithm** *(essentially linear time)*.

Multiplicative groups of finite fields

Second attempt at a cryptographic $\mathcal{G}$: prime-order subgroups of $\mathbb{G}_m(\mathbb{F}_q)$.

How do subgroups of $\mathbb{G}_m(\mathbb{F}_q)$ measure up against a black-box group?

Order need to choose q carefully

Elem. Size $\geq \log_2 N + 1$ bits (best case $q = 2N + 1$, N prime)

Elem. Ops $\sim \log_2^c N$ bit-ops $(1 < c \leq 2)$

DLP? *Good news for people who like bad news...*

Subexponential notation

Recall notation for **subexponential** complexities:

$$L_x[\alpha, c] = \exp\left((c + o(1))(\log x)^{\alpha}(\log\log x)^{(1-\alpha)}\right)$$

Idea *(with $\widetilde{O}(f)$ meaning $O(f)$ ignoring log factors)*:

- $L_x[0, c] = \widetilde{O}((\log x))^c)$: polynomial behaviour in $\log x$
- $L_x[1, c] = \widetilde{O}(x^c)$: exponential behaviour in $\log x$

Also: $L_x(\alpha) := L_x[\alpha, c]$ for any c

Discrete Logarithms in finite fields

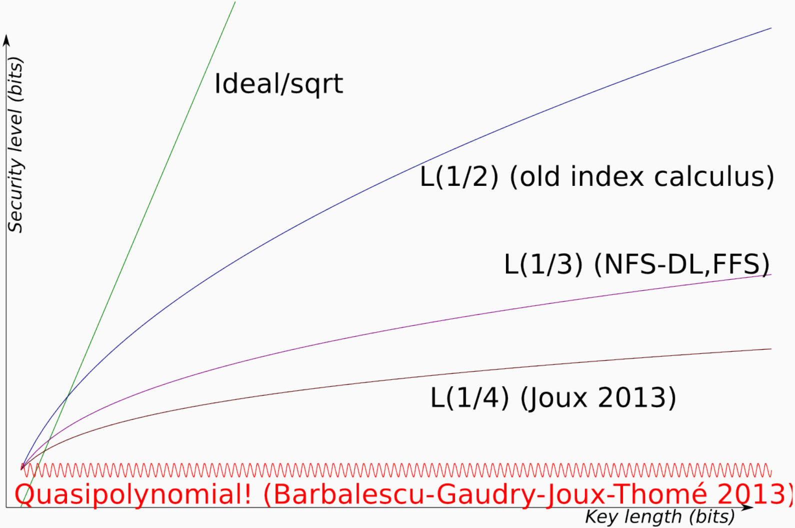

Discrete Logarithms in finite fields

This improvement isn't just asymptotic/theoretical:
records have been repeatedly (and spectacularly) broken since 2013.

The large characteristic case is still in $L(1/3)$, but small-characteristic finite fields are officially useless for discrete-log-based cryptography.

Elliptic Curves

A very short introduction

The base field

We will mostly work over $\mathbb{F}_q$, where q is a power of p,
though sometimes we will work/think over $\mathbb{Q}$, since equations over $\mathbb{Q}$ hold modulo all but finitely many p (i.e., those appearing as factors of denominators).

- Normally, $p \neq 2, 3$.
- *But* in some hardware implementations, $q = 2^n$ with n prime.
- In practice: $q = p$ or p^2.
- *But* in pairing-based crypto, we often need $q = p^n$ with $n \leq 12$.

The main unit of measure for complexity is $\log q$.

Short **Weierstrass models**: nonsingular plane cubic curves

$$\mathcal{E} : y^2 = x^3 + ax + b\,,$$

where the parameters a and b in $\mathbb{F}_q$ satisfy $4a^3 + 27b^2 \neq 0$ (**nonsingularity condition**).

There is a natural **involution** $\ominus : (x, y) \mapsto (x, -y)$ (negation).

Points on $\mathcal{E}$: $(\alpha, \beta) \in \mathbb{F}_q^2$ s.t. $\beta^2 = \alpha^3 + a\alpha + b$
Plus a unique **point at infinity**, $\mathcal{O}_{\mathcal{E}}$ (zero element)

Projective space

Consider the projective plane $\mathbb{P}^2$. Two-dimensional, with three coordinates:

$$\mathbb{P}^2(\mathbb{F}_q) = \big\{(\alpha : \beta : \gamma) \in \mathbb{F}_q^3 \setminus \{(0,0,0)\}\big\} \,/\, \sim$$

where $\sim$ is the equivalence relation defined by

$$(\alpha : \beta : \gamma) \sim (\lambda\alpha : \lambda\beta : \lambda\gamma) \text{ for all } \lambda \neq 0 \in \mathbb{F}_q \,.$$

Warning on projective coordinates

The coordinates X, Y, Z can be 0 or $\neq 0$ at a point P but they do not have any other well-defined values: *$Z(P) = 0$ is meaningful, but $X(P) = 1$ is not*.

More generally: homogeneous polynomials in X, Y, Z (*eg.* $X^2 - YZ, X + Y - Z$) can be either 0 or $\neq 0$ at points.

Functions on $\mathbb{P}^2$ are quotients of homogeneous polynomials *of the same degree*. *Functions can have proper, nontrivial values*.

Example:

$$(X/Z)(\lambda\alpha : \lambda\beta : \lambda\gamma) = \alpha/\gamma = (X/Z)(\alpha : \beta : \gamma) \quad \forall\lambda \neq 0\,.$$

Affine and projective space

The (x, y)-plane $\mathbb{A}^2$ is an open subvariety filling in almost all of $\mathbb{P}^2$: we have an inclusion $\mathbb{A}^2 \hookrightarrow \mathbb{P}^2$ defined by

$$(x, y) \longmapsto (X : Y : Z) = (x : y : 1)$$

with an inverse mapping

$$(X : Y : Z) \longmapsto (x, y) = (X/Z, Y/Z)$$

which is only defined where $Z \neq 0$.

The "missing part" where $Z = 0$ is the "line at infinity".

Exercise: Describe the points in $\mathbb{P}^2(\mathbb{F}_q) \setminus \mathbb{A}^2(\mathbb{F}_q)$.

Projective elliptic curves

Putting $(x, y) = (X/Z, Y/Z)$ gives a projective model

$$\mathcal{E} : Y^2Z = X^3 + aXZ^2 + bZ^3 \subseteq \mathbb{P}^2 .$$

Affine points (α, β) become projective points $(\alpha : \beta : 1)$

The point at infinity $\mathcal{O}_{\mathcal{E}}$ is $(0 : 1 : 0)$; it is the unique point on $\mathcal{E}$ with $Z = 0$.

This is not the only projective closure/model of $\mathcal{E}$...

Rational points

For any commutative $\mathbb{F}_q$-algebra K (*ie, a ring with a homomorphism* $\mathbb{F}_q \to K$), the set of *K-rational points* of $\mathcal{E}$ is

$$\mathcal{E}(K) := \left\{(\alpha, \beta) \in K^2 : \beta^2 = \alpha^3 + a\alpha + b\right\} \cup \{\mathcal{O}_{\mathcal{E}}\} \ .$$

In projective coordinates,

$$\mathcal{E}(K) = \left\{(\alpha : \beta : 1) : \alpha, \beta \in K, \beta^2 = \alpha^3 + a\alpha + b\right\} \cup \{(0 : 1 : 0)\} \ .$$

The point $\mathcal{O}_{\mathcal{E}} = (0 : 1 : 0)$ is the unique **point at infinity** of $\mathcal{E}$.

The group law

Projectively: all lines intersect $\mathcal{E}$ in exactly three points (with multiplicity).

If two are in $\mathcal{E}(\mathbb{F}_q)$, then so is the third.

The group law:

$$P, Q, R \text{ collinear} \iff P \oplus Q \oplus R = 0$$

Identity element: $0 = \mathcal{O}_{\mathcal{E}} = (0 : 1 : 0)$

Each "vertical" line $x = \alpha$ intersects $\mathcal{E}$ in three points $\{(\alpha : \beta : 1), (\alpha : -\beta : 1), \mathcal{O}_{\mathcal{E}}\}$ where $\beta^2 = \alpha^3 + a\alpha + b$.

Hence:

$$\ominus : (x : y : 1) \longmapsto (x : -y : 1) \text{is the negation map on } \mathcal{E}\,.$$

The group law: Adding ($R = P \oplus Q$) and Doubling ($R = [2]P$)

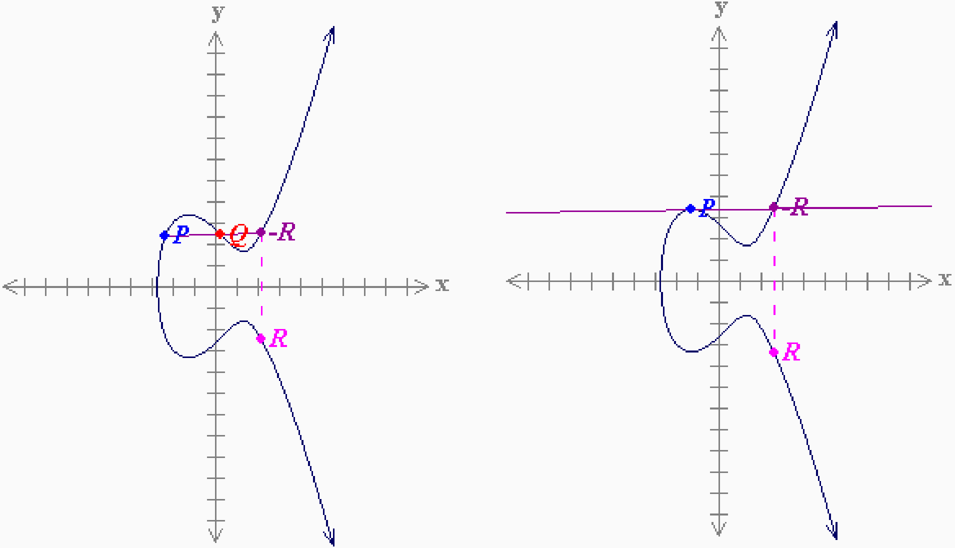

Computing $P \oplus Q$ on $\mathcal{E} : y^2 = x^3 + ax + b$

- $P = \mathcal{O}_\mathcal{E}$ or $Q = \mathcal{O}_\mathcal{E}$? Nothing to be done.
- If $P = \ominus Q$, then $P \oplus Q = \mathcal{O}_\mathcal{E}$

Otherwise: *compute $P \oplus Q$ using low-degree polynomial expressions*

$$\begin{aligned} x(P \oplus Q) &= \lambda^2 - x(P) - x(Q), \\ y(P \oplus Q) &= -\lambda x(P \oplus Q) - \nu, \end{aligned}$$

where

$$\lambda := \begin{cases} (y(P) - y(Q))/(x(P) - x(Q)) & \text{if } x(P) \neq x(Q), \\ (3x(P)^2 + a)/(2y(P)) & \text{if } P = Q \end{cases}$$

$$\nu := \begin{cases} (x(P)y(Q) - x(Q)y(P))/(x(P) - x(Q)) & \text{if } x(P) \neq x(Q), \\ -y(P)/2 + (2ax(P) + 3b)/(2y(P)) & \text{if } P = Q. \end{cases}$$

The group law, as seen by the computer

Algorithmic benefit of projective coords: avoiding costly inversions.

We need addition and general scalar multiplication

$$P \mapsto [m]P := \underbrace{P \oplus \cdots \oplus P}_{m \text{ times}};$$

implement using addition chains (naïve: double-and-add loops).

Main subroutines:

Addition $(X_1 : Y_1 : Z_1) \oplus (X_2 : Y_2 : Z_2)$

Doubling $[2](X_1 : Y_1 : Z_1)$

Mixed addition $(X_1 : Y_1 : Z_1) \oplus (x_2 : y_2 : 1)$ *(second operand fixed)*

See `http://hyperelliptic.org/EFD/g1p/auto-shortw-projective.html`

Algorithmic group law: addition

Algorithm 3: Projective adding: computes $(X_1 : Y_1 : Z_1) \oplus (X_2 : Y_2 : Z_2)$

Cost: $12M + 2S + 6\text{add} + 1\times 2$

$(X_1Z_2, Y_1Z_2, Z_1Z_2) \leftarrow (X_1 * Z_2, Y_1 * Z_2, Z_1 * Z_2)$ `// omit in "mixed" case` $Z_2 = 1$

$u \leftarrow Y_2 * Z_1 - Y_1Z_2$

$uu \leftarrow u^2$

$v \leftarrow X_2 * Z_1 - X_1Z_2$

$vv \leftarrow v^2$

$vvv \leftarrow v * vv$

$R \leftarrow vv * X_1Z_2$

$A \leftarrow uu * Z_1Z_2 - vvv - 2 * R$

$(X_3, Y_3, Z_3) \leftarrow (v * A, u * (R - A) - vvv * Y_1Z_2, vvv * Z_1Z_2)$

return $(X_3 : Y_3 : Z_3)$

Algorithmic group law: doubling

Algorithm 4: Projective doubling: computes $[2](X_1 : Y_1 : Z_1)$.

Cost: $5M + 6S + 1 \times a + 7\text{add} + 3 \times 2 + 1 \times 3$

$$(XX, ZZ) \leftarrow (X_1{}^2, Z_1{}^2)$$

$$w \leftarrow a * ZZ + 3 * XX$$

$$s \leftarrow 2 * Y_1 * Z_1$$

$$ss \leftarrow s^2$$

$$sss \leftarrow s * ss$$

$$R \leftarrow Y_1 * s$$

$$RR \leftarrow R^2$$

$$B \leftarrow (X_1 + R)^2 - XX - RR$$

$$h \leftarrow w^2 - 2 * B$$

$$(X_3, Y_3, Z_3) \leftarrow (h * s, w * (B - h) - 2 * RR, sss)$$

return $(X_3 : Y_3 : Z_3)$

Rough operation counts

In terms of $\mathbb{F}_q$-operations:

- Doubling costs $\sim 5M + 6S + 1 \times a$
- Addition costs $\sim 12M + 2S$
- Adding a fixed/normalized point costs $\sim 9M + 2S$

Exponentiation in $\mathcal{E}(\mathbb{F}_q)$ is an order of magnitude slower than in $\mathbb{G}_m(\mathbb{F}_q)$ *for the same value of q.*

Advantage: since their DLPs seem harder, we can use elliptic curves over much smaller fields to get the same level of security.

At modern security levels, exponentiation in $\mathcal{E}(\mathbb{F}_q)$ is *faster* than in $\mathbb{G}_m(\mathbb{F}_q)$.

Elliptic Curve vs $\mathbb{F}_p$/RSA parameters

Security level (bits)	Elliptic $\mathcal{E}(\mathbb{F}_p)$ ($\log_2 p$)	$\mathbb{G}_m(\mathbb{F}_p)$/RSA ($\log_2 p$)	keylength ratio
56	112	512	4.57
64	128	704	5.5
80	160	1024	6.4
96	192	1536	8.0
112	224	2048	9.14
128	256	3072	12.0
192	384	7680	20.0
256	512	15360	30.0

Bonus Track 1: Group structure and Torsion

Group structure

Cryptographers generally see elliptic curves as a replacement for $\mathbb{G}_m(\mathbb{F}_q) = \mathbb{F}_q^\times$, with **more flexibility** and a **harder DLP**.

We know that $\mathbb{G}_m(\mathbb{F}_q)$ is cyclic, of order $q - 1$. Given the factorization of $q - 1$, we know everything about the subgroups of $\mathbb{G}_m(\mathbb{F}_q)$.

Over the algebraic closure: if we write $\mathbb{G}_m(\overline{\mathbb{F}}_q)[m]$ for the m-torsion subgroup of $\mathbb{G}_m(\overline{\mathbb{F}}_q)$ (the kernel of m-powering) then we have

- $\mathbb{G}_m(\mathbb{F}_q)[\ell] \subseteq \mathbb{G}_m(\overline{\mathbb{F}}_q)[\ell] \cong \mathbb{Z}/\ell\mathbb{Z}$ for prime $\ell \neq p$
- $\mathbb{G}_m(\mathbb{F}_q)[\ell] \subseteq \mathbb{G}_m(\overline{\mathbb{F}}_q)[\ell^k] \cong \mathbb{Z}/\ell^k\mathbb{Z}$ for prime $\ell \neq p$
- $\mathbb{G}_m(\mathbb{F}_q) = \mathbb{G}_m(\overline{\mathbb{F}}_q)[p] = 0$

Analogous questions for elliptic curves $\mathcal{E}/\mathbb{F}_q$: what is the group structure of $\mathcal{E}(\mathbb{F}_q)$?

The size of the group

First question: given an elliptic curve $\mathcal{E}/\mathbb{F}_q$, what is $\#\mathcal{E}(\mathbb{F}_q)$?

First approximation: $\mathcal{E}$ is a curve: a **one-dimensional** object over $\mathbb{F}_q$, so we might guess that $\#\mathcal{E}(\mathbb{F}_q)$ has the same order of magnitude as a line over $\mathbb{F}_q$.

That is, we naïvely expect $O(q)$ points in $\mathcal{E}(\mathbb{F}_q)$.

The size of the group

Second approximation: consider $\mathcal{E} : y^2 = x^3 + ax + b$ over $\mathbb{F}_q$. We have

- Exactly one point at infinity, and
- q potential values for x, each of which corresponds to
  - 0 points in $\mathcal{E}(\mathbb{F}_q)$ if $x^3 + ax + b$ is not a square in $\mathbb{F}_q$
  - 1 point in $\mathcal{E}(\mathbb{F}_q)$ if $x^3 + ax + b = 0$
  - 2 points in $\mathcal{E}(\mathbb{F}_q)$ if $x^3 + ax + b$ is a nonzero square in $\mathbb{F}_q$

So *a priori*, there is **at least 1** and **at most** $2q + 1$ points in $\mathcal{E}(\mathbb{F}_q)$.

The size of the group

On $\mathcal{E} : y^2 = x^3 + ax + b$ over $\mathbb{F}_q$ we have

- Exactly one point at infinity, and
- q potential values for x, each of which corresponds to
  - 0 points in $\mathcal{E}(\mathbb{F}_q)$ if $x^3 + ax + b$ is not a square in $\mathbb{F}_q$
  - 1 point in $\mathcal{E}(\mathbb{F}_q)$ if $x^3 + ax + b = 0$
  - 2 points in $\mathcal{E}(\mathbb{F}_q)$ if $x^3 + ax + b$ is a nonzero square in $\mathbb{F}_q$

Take q odd: there are exactly $\frac{q-1}{2}$ nonzero squares and $\frac{q-1}{2}$ nonsquares in $\mathbb{F}_q$.

If we model $x \mapsto x^3 + ax + b$ as a random function, then we would expect

$$\#\mathcal{E}(\mathbb{F}_q) = q + 1 + O(\sqrt{q}) .$$

Problem: $x \mapsto x^3 + ax + b$ is **not** random...

Hasse's theorem

Efficiently computing $\#\mathcal{E}(\mathbb{F}_q)$ *in general* is a fascinating algorithmic problem *(for more, see the Schoof and SEA algorithms).*

Hasse's theorem:

$$\mathcal{E}(\mathbb{F}_q) = q + 1 - t_{\mathcal{E}} \qquad \text{with} \qquad |t_{\mathcal{E}}| \leq 2\sqrt{q}$$

Deuring's theorem: Let p be prime. Then for every t in the interval $[-2\sqrt{p}, 2\sqrt{p}]$, there exists an elliptic curve $\mathcal{E}/\mathbb{F}_p$ with $\#\mathcal{E}(\mathbb{F}_p) = p + 1 - t$.

Deuring's theorem becomes more complicated when we replace p with a general prime power q, but the result is the same except when $p \mid t$.

Torsion points

Let $P = (x : y : 1) \neq \mathcal{O}_{\mathcal{E}}$ be a generic point of $\mathcal{E}$.

Formally iterating $\oplus$ on P yields polynomial expressions for the coordinates of $[m]P$ in terms of x and y for every integer m: that is,

$$[m](x : y : 1) = \left(\Phi_m(x)\Psi_m(x) : \Omega_m(x, y) : \Psi_m^3(x)\right)$$

where Φ_m, Ω_m, Ψ_m depend only on m (and $\mathcal{E}$) *(in fact, they are in $\mathbb{Z}[a, b][x, y]$.)*

We can compute Φ_m, Ω_m, and Ψ_m using recurrences derived from the group law.

Ψ_m is the most fundamental: it is called the m-th **division polynomial**.

Division polynomials

The **division polynomials** for $\mathcal{E} : y^2 = x^3 + ax + b$ are defined by

$$\begin{aligned}
\Psi_{-1} &:= -1 \\
\Psi_0 &:= 0 \\
\Psi_1 &:= 1 \\
\Psi_2 &:= 2y \\
\Psi_3 &:= 3x^4 + 6ax^2 + 12bx - a^2 \\
\Psi_4 &:= 4y(x^6 + 5ax^4 + 20bx^3 - 5a^2x^2 - 4abx - 8b^2 - a^3) \\
\Psi_{2k} &:= \Psi_k(\Psi_{k+2}\Psi_{k-1}^2 - \Psi_{k-2}\Psi_{k+1}^2)/2y \text{ for all } k > 2 \\
\Psi_{2k+1} &:= \Psi_{k+2}\Psi_k^3 - \Psi_{k+1}^3\Psi_{k-1} \text{ for all } k \geq 2
\end{aligned}$$

The division polynomials have analogous (but more complicated) definitions for elliptic curves with more general defining equations. In particular, there exist division polynomials for curves defined over fields of characteristic 2 and 3.

Division polynomials

The Ω_m and Φ_m can be expressed in terms of x, y, and the Ψ_m:

$$\Phi_m(x,y) = x\Psi_m(x,y)^2 - \Psi_{m-1}(x,y)\Psi_{m+1}(x,y)$$

and

$$\Omega_m(x,y) = \frac{\Psi_{m+2}\Psi_{m-1}^2 - \Psi_{m-2}\Psi_{m+1}^2}{4y} .$$

We can rewrite Φ_m, Ψ_m^2, Ψ_{2m+1}, and Ψ_{2m}/y as polynomials in x only (using $y^2 = x^3 + ax + b$):

$$\begin{aligned} \Psi_m(x) &= mx^{(m^2-1)/2} + \cdots && \text{if } m \text{ is odd;} \\ \Psi_m(x) &= y(mx^{(m^2-4)/2} + \cdots) && \text{if } m \text{ is even;} \\ \Psi_m^2(x) &= m^2x^{m^2-1} + \cdots && \text{for all } m; \\ \Phi_m(x) &= 1x^{m^2} + \cdots && \text{for all } m . \end{aligned}$$

What do division polynomials tell us about torsion?

We have

$$[m](x,y) = \mathcal{O}_{\mathcal{E}} \iff \Psi_m(x,y) = 0 \ .$$

Use $\deg \Psi_m$ to bound torsion rank, hence group structure.

Let ℓ^k be any prime power. If $\mathbb{F}_q \supset \mathbb{Q}$, then

$$\mathcal{E}[\ell^k](\overline{\mathbb{F}}_q) \cong (\mathbb{Z}/\ell^k\mathbb{Z})^2 \ .$$

If $\mathcal{E}$ is defined over a finite field then

$$\mathcal{E}[\ell^k](\overline{\mathbb{F}}_p) \cong \begin{cases} (\mathbb{Z}/\ell^k\mathbb{Z})^2 & \text{if } \ell \neq p \\ (\mathbb{Z}/p^k\mathbb{Z}) & \text{if } \ell = p \text{ and } \mathcal{E} \text{ is ``ordinary''} \\ 0 & \text{if } \ell = p \text{ and } \mathcal{E} \text{ is ``supersingular''} \end{cases}$$

Possible group structures

The possible group structures for elliptic curves over $\mathbb{F}_q$ are extremely limited.

Theorem: If $\mathcal{E}$ is defined over $\mathbb{F}_q$, then

$$\mathcal{E}(\mathbb{F}_q) \cong \mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z}$$

where

$$d_2 \mid d_1 \quad \text{and} \quad d_2 \mid (q-1) .$$

Why does d_2 divide $q-1$? Because of the non-degeneracy of the Weil pairing.

Exercise: Prove that $\mathcal{E}(\mathbb{F}_q) \cong \mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z}$ with $d_2 \mid d_1$.
Hint: Use the fact that $\mathcal{E}(\mathbb{F}_q)[\ell^k](\overline{\mathbb{F}}_q) \cong (\mathbb{Z}/\ell\mathbb{Z})^2$ for $\ell \neq p$, etc.

Bonus Track 2:

The Maurer Reduction

Relating DLP and CDHP hardness

Relating the DLP and the CDHP

Why do we believe the Computational Diffie–Hellman Problem is hard?

Clearly, if we can solve DLP instances $(P, Q = [x]P) \mapsto x$ in an abstract group $\mathcal{G}$, then we can also solve CDHP instances $(P, A = [a]P, B = [b]P) \mapsto [ab]P$.

Converse *(Den Boer, Maurer, Wolf, ...; under "reasonable" conditions):*
If we can solve CDHPs in $\mathcal{G}$, then we can solve DLPs in $\mathcal{G}$.

1. Reduce to the case of prime N;
2. View $\mathcal{G}$ as a representation of $\mathbb{F}_N$, via $\mathcal{G} \ni [a]P \leftrightarrow a \in \mathbb{F}_N$, with the group operation as $+$ and a $\mathcal{G}$-DH oracle for $\times$.
3. This allows Boneh–Lipton-style **black-box field** arguments, which give subexponential (or better) reductions.

The Maurer reduction: how does it work?

We want to **solve a DLP** instance $Q = [x]P$ in $\mathcal{G}$ of prime order N,
given a DH oracle for $\mathcal{G}$ (so we can compute $[F(x)]P$ $\forall$ poly F) in $\mathbb{F}_N[X]$:

1. Find an $\mathcal{E}/\mathbb{F}_N : Y^2 = X^3 + aX + b$ s.t. $\mathcal{E}(\mathbb{F}_N)$ is cyclic with **polynomially smooth order** (*this is the hard part!*), and let (x_0, y_0) be a generator for $\mathcal{E}(\mathbb{F}_N)$.
2. Compute $[x^3 + ax + b]P$ using the DH oracle
3. Use Tonelli–Shanks to compute a $Y = [y]P$ s.t. $[y^2]P = [x^3 + ax + b]P$.
   If this fails: replace $Q = [x]P$ with $Q' = Q + [\delta]P = [x + \delta]P$ and try again...
   Now (Q, Y) is a point in $\mathcal{E}(\mathcal{G})$; we still don't know x or y.
4. Solve the DLP instance $(Q, Y) = [e]([x_0]P, [y_0]P)$ in $\mathcal{E}(\mathcal{G})$ for e.
   Pohlig–Hellman: solve DLPs in $\mathcal{E}(\mathcal{G})$ in polynomial time.
5. Compute $(x, y) = [e](x_0, y_0)$ in $\mathcal{E}(\mathbb{F}_N)$ and return x.

Why is it conditional?

ECC depends on the fact that finding almost-prime-order curves is easy.

Weird (reassuring) converse: finding smooth-order curves is extremely hard (unless we get to choose the field size).

Theory: nothing guarantees that there are polynomially smooth orders of constructible curves in the Hasse interval.

Practice: we seem to be able to find sufficiently smooth auxiliary curves for cryptographically useful N.

Theory again: relax to **subexponential** smoothness.

See Muzereau–Smart–Vercauteren (2004) and Bentahar (2005) for sharper plausible/unconditional subexponential reductions.

Bonus Track 3:

Degenerate Elliptic Curves

The group law on singular curves

Singular plane cubics

Recall that when we defined elliptic curves in short Weierstrass form

$$\mathcal{E} : y^2 = x^3 + ax + b$$

we imposed the **nonsingularity condition** $4a^3 + 27b^2 \neq 0$.

Question*: What is a singularity?*

What happens to the geometric group law *(any three collinear points sum to zero)* for **singular curves**, where $4a^3 + 27b^2 = 0$?

Nodal and cuspidal cubics

Consider projective Weierstrass models $\mathcal{E} : Y^2Z = X^3 + b_2X^2Z + b_4XZ^2 + b_6Z^3$.

Up to isomorphism, there are **two kinds** of singular cubics:

Nodal $\mathcal{E} : Y^2Z = X^2(X - cZ)$ with $c \in \mathbb{F}_q \neq 0$:
A single "node" (like a self-intersection) at $(0, 0)$. Notice that there are *two* tangent lines at $(0, 0)$: $X = \sqrt{c}Y$ and $X = -\sqrt{c}Y$.

Cuspidal $\mathcal{E} : Y^2Z = X^3$.
A single "cusp" (like a sharp point) at $(0, 0)$.
The tangent cone at $(0, 0)$ is the entire plane!

Cuspidal cubics and the additive group

Consider the **cuspidal cubic** $\mathcal{E} : Y^2Z = X^3$.

The singular point is $S = (0 : 0 : 1)$.

We still have a unique point at infinity, $\mathcal{O}_{\mathcal{E}} = (0 : 1 : 0)$.

We want to define the "usual" group law on $\mathcal{E} \setminus \{S\}$:

zero is $\mathcal{O}_{\mathcal{E}}$;

negation is reflection in the x-axis, $(X : Y : Z) \mapsto (X : -Y : Z)$;

addition is defined for P and Q in $\mathcal{E} \setminus \{S\}$ by

1. taking the line through P and Q,
2. finding the third point R of intersection, then
3. negating R to get $P \oplus Q$ (so $P \oplus Q \oplus R = 0$).

Cuspidal cubics and the additive group

The points in $\mathcal{E}(\mathbb{F}_p) \setminus \{S\}$ are

$$P_\alpha = (\alpha : 1 : \alpha^3) \quad \text{for each} \quad \alpha \in \mathbb{F}_p$$

(notice that $P_0 = (0 : 1 : 0) = \mathcal{O}_\mathcal{E}$, the point at infinity).

Negation: $\ominus : (X : Y : Z) \mapsto (X : -Y : Z)$ sends P_α to $\ominus P_\alpha = P_{-\alpha}$

Addition: $P_\alpha \oplus P_\beta = P_{\alpha+\beta}$. The line through P_α and P_β is

$$L_{\alpha,\beta} : \alpha\beta(\alpha+\beta)Y = (\alpha^2 + \alpha\beta + \beta^2)X - Z\,,$$

and the three points of intersection are

$$L_{\alpha,\beta} \cap \mathcal{E} = \left\{P_\alpha\,, P_\beta\,, (-(\alpha+\beta) : 1 : -(\alpha+\beta)^3) = P_{-(\alpha+\beta)}\right\}\,,$$

so $\mathcal{E}(\mathbb{F}_q) \setminus \{(0 : 0 : 1)\} \cong (\mathbb{F}_q, +)$, the additive group.

Nodal cubics and the multiplicative group

Now consider the nodal cubic $\mathcal{E} : Y^2Z = X^2(X - cZ)$.

The “group law” is more complicated here.

Singular point: $S = (0 : 0 : 1)$. There are *two* lines tangent to $\mathcal{E}$ there, $Y = \sqrt{c}X$ and $Y = -\sqrt{c}X$. To simplify, **suppose** c is square.

Change coordinates to a system defined by the tangent lines:
let $U = Y + \sqrt{c}X$ and $V = Y - \sqrt{c}X$. Now $\mathcal{E}$ is defined by

$$\mathcal{E} : \sqrt{c}^3 UVZ = (U - V)^3 ;$$

the singularity S is still at $(0 : 0 : 1)$.

Questions:

1. Where does $\mathcal{O}_\mathcal{E}$ map to in this coordinate system?
2. What is the “negation” operation in these coordinates?

The nodal cubic and the multiplicative group

For **addition**: first observe that any line in $\mathbb{P}^2$ that does not pass through $S = (0 : 0 : 1)$ has the form $Z = lU + mV$, and it meets $\mathcal{E}$ where $(U - V)^3 = 8\sqrt{c}^3 UV(lU + mV)$.

Exercises:

1. Check that if $(U_i : V_i : Z_i)$ for $i \in \{1, 2, 3\}$ are the three points of intersection of $\mathcal{E}$ with a line, then
   $(U_1/V_1) \cdot (U_2/V_2) \cdot (U_3/V_3) = 1$.
2. Conclude that $\mathcal{E}(\mathbb{F}_q) \setminus \{(0 : 0 : 1)\}$, with this chord-and-tangent "group law", is a model for $\mathbb{F}_q^\times$.
3. Can you find a nice parametrization $\alpha \in \mathbb{F}_p^\times \longmapsto P_\alpha = (X_\alpha : Y_\alpha : Z_\alpha) \in \mathcal{E}(\mathbb{F}_p)$?
4. What happens when c is not a square?

Conclusion

Mathematical perspective:

- Elliptic curves are not just a formal replacement for the multiplicative group $\mathbb{G}_m$: they are a sort of **deformation** of $\mathbb{G}_m$ (and also of the additive group $\mathbb{G}_a$).
- We can see both the multiplicative and the additive group as **degenerate elliptic curves**.

Algorithmic consequences:

- Any elliptic-curve algorithm has an immediate analogue for $\mathbb{F}_q^\times$ (and $(\mathbb{F}_q, +)$).
- Any algorithm for $\mathbb{F}_q$ that requires *only multiplications and divisions* has an immediate elliptic-curve analogue.
- Similarly, any algorithm for $\mathbb{F}_q$ that requires *only additions and subtractions* has an elliptic-curve analogue.