MPRI – Cours 2-12-2

F. Morain

Lecture IIa: Generic groups

2009/11/30

I. The discrete logarithm in a group.

II. A typical generic group: an elliptic curve over a finite field.

I. The discrete logarithm in a group

Def. (DLP) Given $G = \langle g \rangle$ of order n and $a \in G$, find $x \in [0..n[$ s.t. $a = g^x$.

Goal: find a resistant group.

Rem. DL is easy in $(\mathbb{Z}/N\mathbb{Z}, +)$, since $a = xg \bmod N$ is solvable in polynomial time.

Relatively easy groups: finite fields, curves of very large genus, class groups of number fields.

Probably difficult groups: elliptic curves.

Generic groups

Rem. generic means we cannot use specific properties of G, just group operations.

Known *generic* solutions:

- enumeration: $O(n)$;
- Shanks: deterministic time and space $O(\sqrt{n})$;
- Pollard: probabilistic time $O(\sqrt{n})$, space $O(1)$ elements of G.

Rem. All these algorithms can be more or less distributed.

A) The Pohlig-Hellman reduction

Idea: reduce the problem to the case n prime.

$$n = \prod_i p_i^{\alpha_i}$$

Solving $g^x = a$ is equivalent to knowing $x \bmod n$, i.e. $x \bmod p_i^{\alpha_i}$ for all i (chinese remainder theorem).

Idea: let $p^\alpha \mid\mid n$ and $m = n/p^\alpha$. Then $b = a^m$ is in the cyclic group of ordre p^α generated by g^m. We can find the log of b in this group, which yields $x \bmod p^\alpha$.

Cost: $O(\max(DL(p)))$.

Consequence: in DH, n must have at least one large prime factor.

B) Shanks

$$x = cu + d, 0 \le d < u, \quad 0 \le c < n/u$$
$$g^x = a \Leftrightarrow a(g^{-u})^c = g^d.$$

- Step 1 **(baby steps)**: compute $\mathcal{B} = \{g^d, 0 \le d < u\}$;
- Step 2 **(giant steps)**:
  - compute $f = g^{-u} = 1/g^u$;
  - for $c = 0..n/u$, if $af^c \in \mathcal{B}$, then stop.
- End: $af^c = g^d$ hence $x = cu + d$.

Analysis:

- $C_o = u + n/u$ group operations;
- $C_m = n/u$ membership tests.

If membership test $= O(1)$, then dominant term is C_o, minimal for $u = \sqrt{n} \Rightarrow$ (deterministic) time and space $O(\sqrt{n})$.
Implementation:

- use hashing to test membership in $\mathcal{B}$;
- all kinds of trade-offs possible if low memory available.

C) Pollard's ρ

Prop. Let $f : E \to E$, $\#E = m$; $X_{n+1} = f(X_n)$ with $X_0 \in E$. The functional digraph of X is:

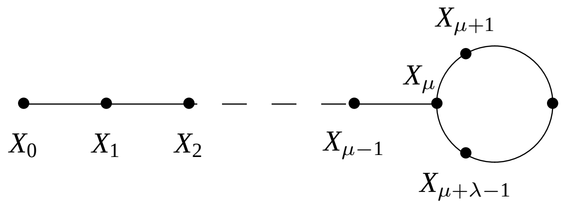

Ex1. If $E_m = G$ is a finite group with m elements, and $a \in G$ of ordre N, $f(x) = ax$ and $x_0 = a$, (x_n) is purely periodic, i.e., $\mu = 0$, and $\lambda = N$.

Ex2. Soit $E_m = \mathbb{Z}/11\mathbb{Z}$, $f : x \mapsto x^2 + 1 \bmod 11$:

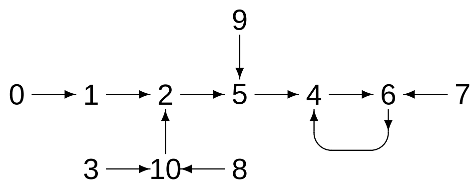

Epact

Thm. (Flajolet, Odlyzko, 1990) When $m \to \infty$

$$\overline{\lambda} \sim \overline{\mu} \sim \sqrt{\frac{\pi m}{8}} \approx 0.627\sqrt{m}.$$

Prop. There exists a unique $e > 0$ (epact) s.t. $\mu \le e < \lambda + \mu$ and $X_{2e} = X_e$. It is the smallest non-zero multiple of λ that is $\ge \mu$: if $\mu = 0$, $e = \lambda$ and if $\mu > 0$, $e = \lceil \frac{\mu}{\lambda} \rceil \lambda$.

Thm. $\overline{e} \sim \sqrt{\frac{\pi^5 m}{288}} \approx 1.03\sqrt{m}$.

Floyd's algorithm:

```
X <- X0; Y <- X0; e <- 0;
repeat
    X <- f(X); Y <- f(f(Y)); e <- e+1;
until X = Y;
```


Application to the discrete log (à la Teske)

Compute the DL of $h = g^x$:

- Choose $y_0 = g^{\alpha_0} h^{\beta_0}$ for $\alpha_0, \beta_0 \in_R [0..n[$;
- Use a function F s.t. given $y = g^\alpha h^\beta$, one can compute efficiently $F(y) = g^{\alpha'} h^{\beta'}$;
- Compute the sequence $y_{k+1} = F(y_k)$ and the exponents $y_k = g^{\alpha_k} h^{\beta_k}$ until $y_i = y_j$.

When $y_i = y_j$, one gets

$$\alpha_i + \beta_i x \equiv \alpha_j + \beta_j x \bmod n$$

or

$$x \equiv (\alpha_j - \alpha_i)(\beta_i - \beta_j)^{-1} \bmod n$$

(with very high probability $\gcd(\beta_i - \beta_j, n) = 1$).

Two versions

Storing a few points:

- Compute r random points $M_k = g^{\gamma_k} h^{\delta_k}$ for $1 \leq k \leq r$;
- use $\mathcal{H} : G \to \{1, \ldots, r\}$;
- define $F(Y) = Y \cdot M_{\mathcal{H}(Y)}$.

Experimentally, $r = 20$ is enough to have a large mixing of points.
Under a plausible model, this leads to a $O(\sqrt{n})$ method (see Teske).

Storing a lot of points:
(van Oorschot and Wiener)
Say a distinguished has some special form; we can store all of them to speed up the process.

D) Shoup's theorem (à la Stinson)

Encoding function: injective map $\sigma : \mathbb{Z}/n\mathbb{Z} \to S$ where S is a set of binary strings s.t. $\#S \geq n$.

Ex. $G = (\mathbb{Z}/q\mathbb{Z})^* = \langle g \rangle$, $n = q - 1$, $\sigma : a \mapsto g^a \bmod q$, S can be $\{0, 1\}^\ell$ where $q < 2^\ell$.

Rem. A generic algorithm should work for any σ.

Oracle $\mathcal{O}$: given $\sigma(i)$ and $\sigma(j)$, computes $\sigma(ci \pm dj \bmod n)$ for any given known integers c and d. This is the only operation permitted.

Game: given $\sigma_1 = \sigma(1)$ and $\sigma_2 = \sigma(a)$ for random a, GenLog succeeds if it outputs a.

Ex. Pollard's algorithm belongs to this class.

Reference: *Cryptography, Theory and Practice*, 2nd edition.

Stinson (2/4)

GenLog produces $(\sigma_1, \sigma_2, \ldots, \sigma_m)$ using $\mathcal{O}$ where

$$\sigma_i = \sigma(c_i + ad_i \bmod n),$$

with $(c_1, d_1) = (1, 0)$ and $(c_2, d_2) = (0, 1)$, $(c_i, d_i) \in \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$.

Key remark: since σ is injective, $\sigma_i = \sigma_j$ iff $c_i + ad_i \equiv c_j + ad_j$, hence a.

Two cases: non-adaptive (choose c_i, d_i before starting) or adaptive.

Thm. Let $\beta = \mathrm{Proba}$(GenLog succeeds). For $\beta > \delta > 0$, one must have $m = \Omega(n^{1/2})$.

Stinson (3/4)

The non-adaptive case: GenLog chooses

$$\mathcal{C} = \{(c_i, d_i), 1 \leq i \leq m\} \subset \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$$

and then computes all σ_i's.
Put

$$\text{Good}(\mathcal{C}) = \{(c_i - c_j)/(d_i - d_j)\}, \#\text{Good}(\mathcal{C}) = \mathcal{G} \leq m(m-1)/2.$$

If $a \in \text{Good}(\mathcal{C})$, GenLog returns a, otherwise some a at random.
α is the event $a \in \text{Good}(\mathcal{C})$,

$$\begin{aligned}
\mathrm{Proba}(\beta) &= \mathrm{Proba}(\beta \| \alpha)\mathrm{Proba}(\alpha) + \mathrm{Proba}(\beta \| \overline{\alpha})\mathrm{Proba}(\overline{\alpha}) \\
&= 1 \times \frac{\mathcal{G}}{n} + \frac{1}{n - \mathcal{G}} \times \frac{n - \mathcal{G}}{n} \\
&= \frac{\mathcal{G} + 1}{n} \leq \frac{m(m-1)/2 + 1}{n}.
\end{aligned}$$

$\Rightarrow$ if proba $> \delta > 0$, then m must be $\Omega(n^{1/2})$.

Stinson (4/4)

The adaptive case: For $1 \leq i \leq m$, $\mathcal{C}_i = \{\sigma_j, 1 \leq j \leq\}$. Then a can be computed at time i if $a \in \text{Good}(\mathcal{C}_i)$. If $a \notin \text{Good}(\mathcal{C}_i)$, then $a \in \mathbb{Z}/n\mathbb{Z} - \text{Good}(\mathcal{C}_i)$ with proba $1/(n - \#\text{Good}(\mathcal{C}_i))$.

And now, what? this result tells you (only) that if you want an algorithm that is faster than Pollard's ρ or Shanks, then you have to work harder...

E) Variants of the DL problem

Decisional DH problem: given (g, g^a, g^b, g^c), do we have $c = ab \bmod n$?

Computational DH problem: given (g, g^a, g^b), compute g^{ab}.

DL problem: given (g, g^a), find a.

Prop. DL $\Rightarrow$ CDH $\Rightarrow$ DCDH.

Thm. converse true for a large class of groups (Maurer & Wolf).

More problems: ℓ-SDH (given g, g^α, ..., g^{α^ℓ}, compute $g^{\alpha^{\ell+1}}$.

Rem. Generalized problems on pairings.

II. A typical generic group: an elliptic curve over a finite field

$\mathbf{K}$ field of characteristic $\neq 2, 3$. Elements of $\mathbf{K}^3 - \{(0,0,0)\}$ are equivalent iff

$$(x_1, y_1, z_1) \sim (x_1', y_1', z_1') \Longleftrightarrow \exists\, \lambda \neq 0, x_1 = \lambda x_1', y_1 = \lambda y_1', z_1 = \lambda z_1'.$$

Projective space: $\mathbf{P}^2(\mathbf{K})$ = equivalence classes of $\sim$.

Elliptic curve defined for points in $\mathbf{P}^2(\mathbf{K})$:

$$Y^2 Z = X^3 + aXZ^2 + bZ^3 \tag{1}$$

with $4a^3 + 27b^2 \neq 0$ (discriminant of E).

Def. $E(\mathbf{K}) = \{(x : y : z)$ satisfying (1)$\}$.

Prop. $E(\mathbf{K}) = \{(0 : 1 : 0)\} \cup \{(x : y : 1)$ satisfying (1)$\}$ = point at infinity $\cup$ affine part.

The group law

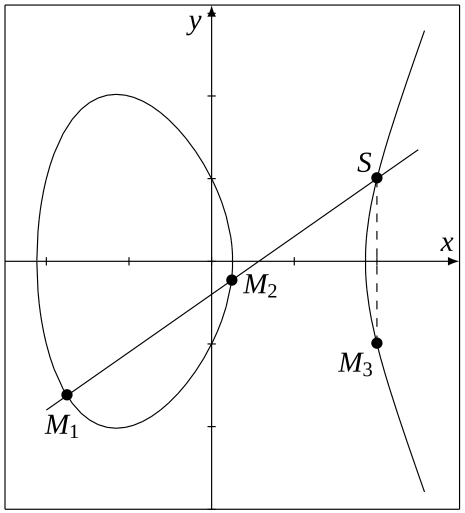

$$M_3 = M_1 \oplus M_2$$

$$\lambda = \begin{cases} (y_1 - y_2)/(x_1 - x_2) \\ (3x_1^2 + a)/(2y_1) \end{cases}$$
$$x_3 = \lambda^2 - x_1 - x_2$$
$$y_3 = \lambda(x_1 - x_3) - y_1$$

$$[k]M = \underbrace{M \oplus \cdots \oplus M}_{k \text{ times}}$$

Rem. Standard equation and group law formulas for any field. Can be improved in many ways, see later.

Cardinality (1/2)

Thm. (Hasse) $\#E(\mathbb{F}_p) = p + 1 - t$, $|t| \leq 2\sqrt{p}$.

Pb: no general formula for $\#E$ except in some special cases.

Thm. (Deuring) given $|t|$, there exists E s.t. $\#E = p + 1 - t$.

Pb: no efficient way for finding E except in some special cases (complex multiplication).

Thm. (Structure) $E(\mathbb{F}_p) \simeq E_1 \times E_2$ of respective ordres m_1 and m_2 s.t. $m_2 \mid p-1$ and $m_2 \mid m_1$.

Cardinality (2/2): do it yourself

Invent a method in time:

- $O(p)$:
- $O(p^{1/2})$:
- $O(p^{1/4})$:

Algorithms:

- $g = 1$, p large: Schoof (1985). $\tilde{O}((\log p)^5)$, completely practical after improvements by Elkies, Atkin, and implementations by M., Lercier, etc. New recent record Enge+M. for $p = 10^{2499} + 7131$ (400 days of AMD 64 Processor 3400+ (2.4GHz)).
- $p = 2$: p-adic methods (Satoh, Fouquet/Gaudry/Harley). Completely solved.

ECDLP

DLP in general resistant on an elliptic curve except

- supersingular curves ($t = 0$), due to the MOV reduction;
- anomalous curves ($t = 1$).

ECC112b: taken from `http://lacal.epfl.ch/page81774.html`, Bos/Kaihara/Kleinjung/Lenstra/Montgomery (EPFL/Alcatel-Lucent Bell Laboratories/MSR) $p = (2^{128} - 3)/(11 * 6949)$, curve secp112r1

- 3.5 months on 200 PS3; 8.5×10^{16} ec additions ($\approx$ 14 full 56-bit DES key searches); started on January 13, 2009, and finished on July 8, 2009.
- half a billion distinguished points using 0.6 Terabyte of disk space.