Index Calculus in Class Groups of Plane Curves of Small Degree

Claus Diem

University of Leipzig

Motivation

Additionally to the DLP in elliptic curves, the DLP in class groups of hyperelliptic curves has been suggested as a cryptographic primitive.

Motivation

Additionally to the DLP in elliptic curves, the DLP in class groups of hyperelliptic curves has been suggested as a cryptographic primitive.

However, it is well known that one can attack the DLP in class groups (Jacobian groups / Picard groups) of hyperelliptic curves via index calculus.

Motivation

Additionally to the DLP in elliptic curves, the DLP in class groups of hyperelliptic curves has been suggested as a cryptographic primitive.

However, it is well known that one can attack the DLP in class groups (Jacobian groups / Picard groups) of hyperelliptic curves via index calculus.

State of the art: Such attacks are more efficient than generic attacks if the genus is ≥ 3.

Motivation

Additionally to the DLP in elliptic curves, the DLP in class groups of hyperelliptic curves has been suggested as a cryptographic primitive.

However, it is well known that one can attack the DLP in class groups (Jacobian groups / Picard groups) of hyperelliptic curves via index calculus.

State of the art: Such attacks are more efficient than generic attacks if the genus is ≥ 3.

What about non-hyperelliptic curves?

Why is this important?

1. The DLP in class groups of non-hyperelliptic genus 3 curves has been suggested as a cryptographic primitive.

- Basiri, Enge, Faugère, Gürel. The arithmetic of Jacobian groups of superelliptic cubics (Math. Comp, 2005)

- Basiri, Enge, Faugère, Gürel. Implementing the Arithmetic of $C_{3,4}$-curves. ANTS VI (2004)

- Flon, Oyono. Fast arithmetic on Jacobians of Picard curves. PKC (2004)

- Koike, Weng. Construction of CM-Picard curves. (Math. Comp. 2004)

- Bauer, Teske, Weng. Point Counting on Picard Curves in Large Characteristic. (Math. Comp.)

Why is this important?

2. It is sometimes possible to transfer the DLP in elliptic curves or in class groups of hyperelliptic curves over $\mathbb{F}_{q^n}$ to the DLP in class groups of curves of higher genus over $\mathbb{F}_q$.

The idea is: If the genus of the resulting curve is not too large, it should be more efficient to attack the resulting DLP via index calculus than the original one with generic methods.

Examples of such attacks:

- GHS
- work by D. & Scholten (talk at last ECC).

Very often the resulting curve is not hyperelliptic anymore.

General result (informally)

An observation.

There is no principal problem in adapting the well known index calculus algorithms from hyperelliptic curves to arbitrary curves.

Heuristically, the running times stay the same except maybe for logarithmic factors.

General result (informally)

An observation.

There is no principal problem in adapting the well known index calculus algorithms from hyperelliptic curves to arbitrary curves.

Heuristically, the running times stay the same except maybe for logarithmic factors.

General result.

The DLP in class groups of non-hyperelliptic curves of small genus can often be solved *faster* than the DLP in class groups of hyperelliptic curves of the same genus over the same field.

The case of genus 3 curves

Let us consider genus 3 curves over $\mathbb{F}_q$.

- The rho method: $\tilde{O}(q^{3/2})$.

The case of genus 3 curves

Let us consider genus 3 curves over $\mathbb{F}_q$.

- The rho method: $\tilde{O}(q^{3/2})$.

- Gaudry's algorithm with optimal factor base (the algorithm by Gaudry and Harley): $\tilde{O}(q^{3/2})$.

The case of genus 3 curves

Let us consider genus 3 curves over $\mathbb{F}_q$.

- The rho method: $\tilde{O}(q^{3/2})$.

- Gaudry's algorithm with optimal factor base (the algorithm by Gaudry and Harley): $\tilde{O}(q^{3/2})$.

- ... with double large prime variation: $\tilde{O}(q^{4/3})$ (Gaudry, Thériault, Thomé).

The case of genus 3 curves

Let us consider genus 3 curves over $\mathbb{F}_q$.

- The rho method: $\tilde{O}(q^{3/2})$.

- Gaudry's algorithm with optimal factor base (the algorithm by Gaudry and Harley): $\tilde{O}(q^{3/2})$.

- ... with double large prime variation: $\tilde{O}(q^{4/3})$ (Gaudry, Thériault, Thomé).

For non-hyperelliptic genus 3 curves one can obtain $\tilde{O}(q)$ (heuristically).

Why?

Why can one obtain better results for non-hyperelliptic curves?

Why?

Why can one obtain better results for non-hyperelliptic curves?

One can exploit the fact that non-hyperelliptic curves can often be defined by equations of a smaller degree.

For example: Non-hyperelliptic genus 3 curves can be defined by equations of degree 4. But hyperelliptic genus 3 curves can only be defined by equations of degree 5 or higher.

More generally ...

Let us consider curves of a fixed genus g defined by equations of a fixed degree $d \geq 4$. Then we have heuristically:

Gaudry's algorithm with optimal factor base + double large prime variation:

$$\tilde{O}(q^{2-\frac{2}{g}}) \ .$$

New algorithm (also with double large prime variation):

$$\tilde{O}(q^{2-\frac{2}{d-2}}) \ .$$

More generally ...

Let us consider curves of a fixed genus g defined by equations of a fixed degree $d \geq 4$. Then we have heuristically:

Gaudry's algorithm with optimal factor base + double large prime variation:

$$\tilde{O}(q^{2-\frac{2}{g}}) \ .$$

New algorithm (also with double large prime variation):

$$\tilde{O}(q^{2-\frac{2}{d-2}}) \ .$$

Moreover, every "sufficiently general" curve of genus g can be defined by an equation of degree $d+1$. This gives a running time of

$$\tilde{O}(q^{2-\frac{2}{g-1}}) \ .$$

Arithmetic in class groups of curves

1. We consider curves $\mathcal{C}/\mathbb{F}_q$ represented by possibly singular plane models $\mathcal{C}_{pm}$ of a fixed degree d. The defining equation is $F(X,Y,Z)=0$.
2. Points on $\mathcal{C}$ are given by their coordinates (x,y,z).
3. Divisors on $\mathcal{C}$ are given as formal sums of points (over extension fields) (“free representation”).
4. Let us fix a point $P_0 \in \mathcal{C}(\mathbb{F}_q)$. By the Theorem of Riemann-Roch every element in $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ is given by $D - gP_0$ for some divisor D of degree $g = g(\mathcal{C})$. “Usually”, the divisor D is unique.

Arithmetic in class groups

Proposition Let us consider curves represented by plane models of a fixed degree over finite fields $\mathbb{F}_q$ with a fixed point $P_0 \in \mathcal{C}(\mathbb{F}_q)$. Then the arithmetic in $\mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ can be performed in randomized polynomial time in $\log(q)$.

This means: Given two divisors D_1, D_2 of degree g in free representation, one can calculate a third divisor D_3 of degree g in free representation with

$$([D_1] - g[P_0]) + ([D_2] - g[P_0]) = [D_3] - g[P_0] \ ,$$

i.e. with

$$D_1 + D_2 \sim D_3 + gP_0$$

in randomized polynomial time in $\log(q)$.

The algorithm by Gaudry and Harley

Let $\mathcal{C}/\mathbb{F}_q$ and $P_0 \in \mathcal{C}(\mathbb{F}_q)$ be as above. Let $a, b \in \mathrm{Cl}^0(\mathcal{C}/\mathbb{F}_q)$ with $b \in \langle a \rangle$.

The goal is to find an $x \in \mathbb{N}$ with $x \cdot a = b$.

We assume that $\ell := \#\langle a \rangle$ is known.

Let $r \in (0, 1]$.

1. Fix a *factor base* $\mathcal{F} = \{F_1, F_2, \ldots\} \subset \mathcal{C}(\mathbb{F}_q)$ of size $\approx q^r$.

The algorithm by Gaudry and Harley

2. Construct a sparse matrix R over $\mathbb{Z}/\ell\mathbb{Z}$ as follows:

For $i = 1, \ldots, \#\mathcal{F} + 1$ do:

{ Select α_i, $\beta_i \in \mathbb{Z}/\ell\mathbb{Z}$ independently and uniformly randomly and calculate D_i with

$$[D_i] - g[P_0] = \alpha_i a + \beta_i b \ .$$

Repeat this until

$$D_i = \sum_j r_{i,j} F_j .$$

(Now $\sum r_{i,j}[F_j] - g[P_0] = \alpha_i a + \beta_i b$.)

Store $(r_{i,j})_j$ as the i-th row of R. }

3. Calculate a random element $\gamma \in \ker(R^t)$, i.e. $\gamma R = 0$.

The algorithm by Gaudry and Harley

We now have

$$(\textstyle\sum_i \gamma_i\alpha_i) \cdot a + (\textstyle\sum_i \gamma_i\beta_i) \cdot b =$$

$$\textstyle\sum_i \gamma_i \cdot (\alpha_i \cdot a + \beta_i b) =$$

$$\textstyle\sum_i \gamma_i \cdot \sum_j r_{i,j}([F_j] - [P_0]) = \sum_{i,j} \gamma_i \, r_{i,j} \cdot ([F_j] - [P_0]) = 0.$$

Assume that $(\sum_i \gamma_i\beta_i)^{-1} \in (\mathbb{Z}/\ell\mathbb{Z})^*$. Then it follows:

$$b = -(\sum_i \gamma_i\beta_i)^{-1} \cdot (\sum_i \gamma_i\alpha_i) \cdot a$$

The algorithm by Gaudry and Harley

We now have

$$(\sum_i \gamma_i \alpha_i) \cdot a + (\sum_i \gamma_i \beta_i) \cdot b =$$

$$\sum_i \gamma_i \cdot (\alpha_i \cdot a + \beta_i b) =$$

$$\sum_i \gamma_i \cdot \sum_j r_{i,j}([F_j] - [P_0]) = \sum_{i,j} \gamma_i \, r_{i,j} \cdot ([F_j] - [P_0]) = 0.$$

Assume that $(\sum_i \gamma_i \beta_i)^{-1} \in (\mathbb{Z}/\ell\mathbb{Z})^*$. Then it follows:

$$b = \underbrace{-(\sum_i \gamma_i \beta_i)^{-1} \cdot (\sum_i \gamma_i \alpha_i)}_{x} \cdot a.$$

The algorithm by Gaudry and Harley

The probability that a divisor of degree g splits completely into rational points is asymptotically (for fixed g)

$$\frac{1}{g!}\ .$$

The probability that a completely split divisor is “smooth” (i.e. it is the sum of elements of the factor base) is (roughly)

$$\left(\frac{\#\mathcal{F}}{\#\mathcal{C}(\mathbb{F}_q)}\right)^g$$

which is asymptotically

$$q^{g\cdot(r-1)}\ .$$

The algorithm by Gaudry and Harley

This means that we can expect:

We have to generate $g! \cdot q^{g\cdot(1-r)} \cdot q^r$ relations until R has more rows than columns. For fixed g the total running time is

$$\tilde{O}(q^{g\cdot(1-r)} \cdot q^r + q^{2r}) \ .$$

For $r = g/(g+1)$ we obtain

$$\tilde{O}(q^{2\cdot\frac{g}{g+1}}) = \tilde{O}(q^{2-\frac{2}{g+1}}) \ .$$

(This running time can be proven for cyclic class groups.)

Our variant

We have a running time of

$$\tilde{O}(q^{2-\frac{2}{g+1}}) \ .$$

We now show that with a variant of this algorithm one can obtain heuristically a running time of

$$\tilde{O}(q^{2-\frac{2}{d-1}}) \ .$$

(This is without double large prime variation.)

Our variant

Let $\mathcal{C}/\mathbb{F}_q$ be as above, and let $\mathcal{F} = \{F_1, F_2, \ldots\} \subset \mathcal{C}(\mathbb{F}_q)$ be a "factor base".

Note: If $h \in \mathbb{F}_q(\mathcal{C})$ with

$$\mathrm{div}(h) = \sum_j r_j F_j \ ,$$

then we have the relation

$$\sum_j r_j [F_j] = 0 \in \mathrm{Cl}^0(C/\mathbb{F}_q) \ .$$

Our variant

Let $\mathcal{C}_{pm}$ be the fixed plane model of $\mathcal{C}$. Let us (for simplicity) assume that $\mathcal{C}_{pm}$ is non-singular "at infinity".

Let D_∞ be the intersection of $\mathcal{C}_{pm}$ with the line $Z = 0$ $(\deg(D) = d)$.

Let $L(X, Y, Z) = \lambda X + \mu Y + \nu Z$ with $\lambda, \mu, \nu \in \mathbb{F}_q$. Let D be the intersection of $L(X, Y, Z) = 0$ with $\mathcal{C}_{pm}$. Let us again assume that the intersection does not contain singular points.

Now D and D_∞ are both divisors on $\mathcal{C}$, and we have

$$\mathrm{div}(\lambda\frac{X}{Z} + \mu\frac{Y}{Z} + \nu) = \mathrm{div}(\frac{L(X,Y,Z)}{Z}) = D - D_\infty \ .$$

Our variant

We want D to be smooth, i.e. $L(X, Y, Z)$ to pass only through points of the factor base.

Note: This implies that $L(X, Y, Z)$ passes through two points of the factor base, thus we can restrict our attention to lines passing through two points of the factor base.

The algorithm is as follows: Fix some number $r \in (0, 1]$.

1. Find relations of the form

$$\sum_i [P_i] - g[P_0] = \alpha a, \ \sum_i [P_i'] - g[P_0] = \beta b.$$

2. Fix the factor base $\mathcal{F} \subset \mathcal{C}(\mathbb{F}_q)$ without singular points with $\#\mathcal{F} \approx q^r$, and include thereby P_0, P_i, P_i' into $\mathcal{F}$.

Our variant

3. Construct a matrix R over $\mathbb{Z}/\ell\mathbb{Z}$ as follows:

Consider a line $L(X, Y, Z) = 0$ through two points of $\mathcal{F}$. Let D be the intersection ($\deg(D) = d$).

If D splits over the factor base, store a corresponding row for R.

Repeat this until all lines are exhausted or the matrix R has more different rows than non-zero columns.

4. Calculate a random element $\gamma \in \ker(R^t)$.

Now $\gamma_1 \alpha a + \gamma_2 \beta b = 0$. If $\gamma_2 \neq 0$, we have a solution to the DLP.

Complexity of our variant

Complexity estimates for a fixed d.

Finding relations for αa and βb can be done in polynomial time in $\log(q)$.

Complexity of our variant

Complexity estimates for a fixed d.

Finding relations for αa and βb can be done in polynomial time in $\log(q)$.

Heuristically, the probability that the divisor of a line through two points of the factor base splits fully over the factor base is equal to the probability that a divisor of degree $d-2$ on $\mathcal{C}/\mathbb{F}_q$ splits completely over the factor base.

This probability is

$$\frac{1}{(d-2)!}q^{(d-2)\cdot(r-1)}$$

asymptotically for fixed d.

Complexity of our variant

We want to generate more different completely split lines than elements of the factor base.

On the basis of the heuristic probability above, one can prove:

This is the case if

$$\#\mathcal{F} \approx q^r \leq d!^{\frac{1}{d-1}} \cdot q^{1-\frac{1}{d-1}} \ .$$

If we have equality, we obtain a running time of

$$\tilde{O}(q^{2-\frac{2}{d-1}})$$

for both the relation generation and the linear algebra part.

Our variant

Warning. The algorithm only works if $d! > q$.

Asymptotically for fixed d, this is no problem.

Double large prime variation

One can use a double large prime variation for both the algorithm by Gaudry-Harley and our algorithm.

Idea of double large prime variation:

Consider relations of the form

$$\sum_j r_j[F_j] + [P] + [Q] - g[P_0] = \alpha a + \beta b$$

or

$$\sum_j r_j[F_j] + [P] + [Q] - [D_\infty] = 0$$

with $P, Q \in \mathcal{C}(\mathbb{F}_q)$. (The set $\mathcal{C}(\mathbb{F}_q) - \mathcal{F}$ is called the set of *large primes*.)

Double large prime variation

We construct a *graph of large prime variation* on $\mathcal{L} \,\dot{\cup}\, \{*\}$:

1. If we have a relation with two large primes P and Q, we insert the points P and Q as vertices into the *graph of large prime variation* as well as an edge from P to Q (with the data for $(r_j)_j, \alpha, \beta$).
2. If we have a relation with one large prime P, we insert an edge from $*$ to P.

... provided we do not obtain a cycle.

If we *would* obtain a cycle containing $*$, we cancel all large primes and thus have a relation over the factor base.

The result

We have the following heuristic result:

Let us consider the DLP in class groups of curves represented by plane models of a fixed degree $d \geq 4$. Then “essentially all” instances of the DLP in such groups can be solved in an expected running time of

$$\tilde{O}(q^{2-\frac{2}{d-2}}) \ .$$

Experimental results by E. Thomé

Emmanuel Thomé has implemented both

- the algorithm by Gaudry-Harley with double large prime variation for hyperelliptic genus 3 curves

- the new algorithm with double large prime variation for non-hyperelliptic genus 3 curves.

For hyperelliptic curves the largest experiment was for $q = 2^{27}$.

For non-hyperelliptic curves the largest experiment was for $q = 2^{31}$.

Experimental results by E. Thomé

Some data:

For the hyperelliptic genus 3 curve, $q = 2^{27}$:

Factor base: $\approx 130\,000$ elements

CPU time for relation search: 9 days.

Experimental results by E. Thomé

Some data:

For the hyperelliptic genus 3 curve, $q = 2^{27}$:

Factor base: $\approx 130\,000$ elements

CPU time for relation search: 9 days.

For the non-hyperelliptic genus 3 curve, $q = 2^{31}$:

Factor base: $\approx 88\,000$ elements

CPU time for relation search: 1 day.

Plane models of small degree

Let $\mathcal{C}/\mathbb{F}_q$ be some curve. We want to find plane models of small degree.

First approach:

Let D_∞ be a random divisor of degree $g+2$. Then the complete linear system $|D_\infty|$ “should” define a map $\mathcal{C} \longrightarrow \mathbb{P}^2$ which is birational to its image, and the image then has degree $g+2$. Experimentally, this always works.

Plane models of small degree

Let $\mathcal{C}/\mathbb{F}_q$ be some curve. We want to find plane models of small degree.

First approach:

Let D_∞ be a random divisor of degree $g+2$. Then the complete linear system $|D_\infty|$ “should” define a map $\mathcal{C} \longrightarrow \mathbb{P}^2$ which is birational to its image, and the image then has degree $g+2$. Experimentally, this always works.

We obtain a heuristic running time of

$$O(q^{2-\frac{2}{g}}) \ ,$$

as in the algorithm by Gaudry-Harley with double large prime variation.

Plane models of small degree

Second approach:

Let $\mathcal{C}$ be non-hyperelliptic. Let K be a canonical divisor. Now let D_0 be a random divisor of degree $g-3$. Then $|K-D_0|$ has degree $2g-2-g+3=g+1$ and is special. If $\mathcal{C}$ and D_0 are “sufficiently general”, $|K-D_0|$ defines a map $\mathcal{C} \longrightarrow \mathbb{P}^2$ which is birational to its image, and the image then has degree $g+1$.

Plane models of small degree

Second approach:

Let $\mathcal{C}$ be non-hyperelliptic. Let K be a canonical divisor. Now let D_0 be a random divisor of degree $g-3$. Then $|K-D_0|$ has degree $2g-2-g+3=g+1$ and is special. If $\mathcal{C}$ and D_0 are "sufficiently general", $|K-D_0|$ defines a map $\mathcal{C} \longrightarrow \mathbb{P}^2$ which is birational to its image, and the image then has degree $g+1$.

Note: If $g(\mathcal{C})=3$, the canonical linear system $|K|$ itself defines an embedding of $\mathcal{C}$ into $\mathbb{P}^2$ of degree 4. The image is called a canonical curve.

Plane models of small degree

Second approach:

Let $\mathcal{C}$ be non-hyperelliptic. Let K be a canonical divisor. Now let D_0 be a random divisor of degree $g-3$. Then $|K-D_0|$ has degree $2g-2-g+3=g+1$ and is special. If $\mathcal{C}$ and D_0 are "sufficiently general", $|K-D_0|$ defines a map $\mathcal{C} \longrightarrow \mathbb{P}^2$ which is birational to its image, and the image then has degree $g+1$.

Note: If $g(\mathcal{C})=3$, the canonical linear system $|K|$ itself defines an embedding of $\mathcal{C}$ into $\mathbb{P}^2$ of degree 4. The image is called a canonical curve.

For fixed g, we obtain a running time of

$$\tilde{O}(q^{2-\frac{2}{g-1}}).$$

Historical remark

The idea to use *principal divisors* to generate relations was used by before by Adleman, DeMarrais, Huang in the “large genus” case.

They obtained an algorithm with a heuristic subexponential running time of $L[\frac{1}{2}]$ for the DLP in class groups of hyperelliptic curves with $g \geq c \cdot \log(q)$.

Conclusions

- The DLP in class groups of non-hyperelliptic curves of small genus can often be solved faster than the DLP in class groups of hyperelliptic curves of the same genus over the same field.

Conclusions

- The DLP in class groups of non-hyperelliptic curves of small genus can often be solved faster than the DLP in class groups of hyperelliptic curves of the same genus over the same field.
- In the "small genus" case, the degree of an equation of a plane model of a curve is at least as important, if not more important, than the genus.

Conclusions

- The DLP in class groups of non-hyperelliptic curves of small genus can often be solved faster than the DLP in class groups of hyperelliptic curves of the same genus over the same field.
- In the “small genus” case, the degree of an equation of a plane model of a curve is at least as important, if not more important, than the genus.
- There is an argument against the usage of non-hyperelliptic genus 3 curves in cryptographic applications. But currently there is no argument for it ...

Acknowledgments

Many thanks to E. Thomé for conducting experiments.

I thank G. Frey, P. Gaudry, F. Hess and E. Viehweg for discussions and comments.