Low Memory Attacks on Small Key CSIDH

Talk at Université de Picardie Jules Verne

Sabrina Kunzweiler
Inria, Institut de Mathématiques Bordeaux

March 25, 2024

(joint work with Jesús-Javier Chi-Dominguez, Andre Esser and Alexander May)

Public key exchange:
Alice and Bob want to create a secure session key.
They can only communicate over a public channel.

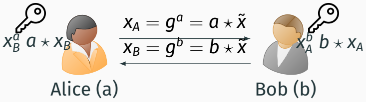

Classical Solution:
Diffie-**H**ellman key exchange based on groups
e.g. $\mathbb{Z}/p\mathbb{Z}$, elliptic curves.

! Shor's algorithm solves Discrete Logarithm in *quantum* polynomial time.

Post-quantum candidate:
Commutative **S**upersingular **I**sogeny **D**iffie-**H**ellman (CSIDH) key exchange based on group actions.

Isogeny-based Group Actions

Elliptic Curves

An **Elliptic Curve** E over $\mathbb{F}_{p^k}$ is defined by an equation

$$E : y^2 = x^3 + ax + b,$$

where $4a^3 + 27b^2 \neq 0$.

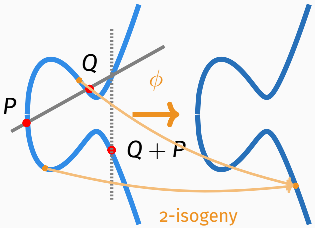

- Points of E form an additive group.

This group is used in the classical Diffie-Hellman protocol.

- An **isogeny** is a non-zero group homomorphism between elliptic curves $\phi : E \to E'$.
- For $p \nmid \ell$, an **ℓ-isogeny** is an isogeny with $\ker(\phi) \equiv \mathbb{Z}/\ell\mathbb{Z}$.

Isogenies are the basis for a post-quantum Diffie-Hellman protocol.

CSIDH [CLMPR, AsiaCrypt'18] Isogeny Graph

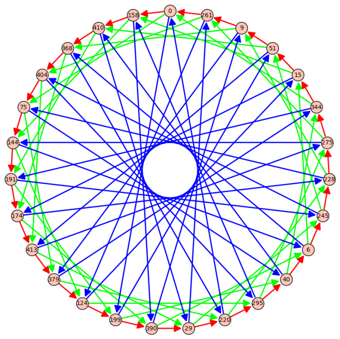

Isogeny Graph over $\mathbb{F}_{419}$ with 3-, 5-, and 7- isogenies.

Vertices: supersingular elliptic curves over $\mathbb{F}_p$

- cardinality: $O(\sqrt{p})$
- labeled by Montgomery coefficient A
  $\Rightarrow E_A : y^2 = x^3 + Ax^2 + x$

Edges: ℓ_i-isogenies for different small primes $\ell_1, \ldots, \ell_n$

- 2-regular for each ℓ_i
- directed graph
- *dual isogenies* allow to go back

Key Idea: Alice and Bob take secret walks on the isogeny graphs.
They only exchange the end vertices.

An example with $p = 59$. The starting vertex is fixed to (0).

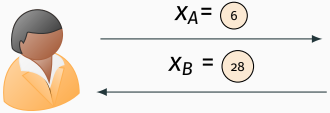

Alice: $a = (2, -1)$

$\Rightarrow x_A = $ (6)

Bob: $b = (-1, -2)$

$\Rightarrow x_B = $ (28)

$K_{ab} = $ (11)

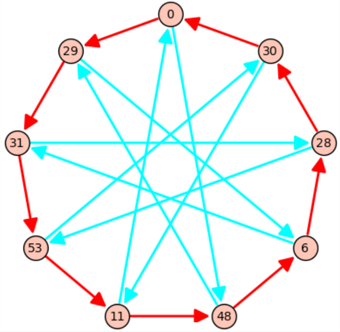

Graph with 3- and 5- isogenies.

Abstract view on CSIDH: Cryptographic group actions [ADMP, AsiaCrypt '20].

Group Action

A map $\star : \mathcal{G} \times \mathcal{X} \to \mathcal{X}$, with $\mathcal{G}$ a group and $\mathcal{X}$ a set, is a group action if:

1. $id \star x = x$ for all $x \in \mathcal{X}$ (identity),
2. $(g \circ h) \star x = g \star (h \star x)$ for all $g, h \in \mathcal{G}$, $x \in \mathcal{X}$ (compatibility).

Cryptographic assumptions

$\mathcal{G}$ is commutative and the following problems are required to be hard.

- **DLOG** Given $x, y \in \mathcal{X}$, find $g \in \mathcal{G}$ with $y = g \star x$.
- **CDH** Given $x, y, z \in \mathcal{X}$, determine $w \in \mathcal{X}$ so that $w = \mathrm{DLOG}(x, y) \star z$.

Diffie Hellman key exchange with group actions

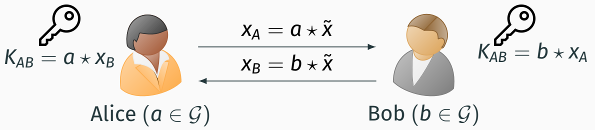

Restricted Effective Group Actions (REGA)

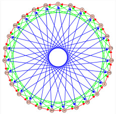

CSIDH as a cryptographic group action $(\mathcal{G}, \mathcal{X}, \star)$

	Formally	Concretely
$\mathcal{X}$	supersingular elliptic curves over $\mathbb{F}_p$	vertices in the isogeny graph
$\mathcal{G}$	the class group $\mathrm{cl}(\mathcal{O})$	exponent vectors
$\star$	isogenies of elliptic curves	paths in the graph

Random Sampling:
We fix $g_1, \ldots, g_n \in \mathcal{G}$ (the colors) and sample

$$g = \prod g_i^{e_i} \leftarrow \mathcal{G}$$

with $(e_1, \ldots, e_n) \leftarrow \{-m, \ldots, m\}^n$.

With a good choice for $g_1, \ldots, g_n$ and m, this sampling is expected to be close to uniform.

Restricted Effective Group Action (REGA)
Notation: $e \star x := \prod g_i^{e_i} \star x$
for $e = (e_1, \ldots, e_n) \in \mathbb{Z}^n$.

Security assumptions in CSIDH

Different DLOG Assumptions

GA − DLOG	Given $x, y \in \mathcal{X}$, find $g \in \mathcal{G}$ with $y = g \star x$	Given vertices in the isogeny graph, find an isogeny connecting them.[1]
REGA − DLOG	Given $x, y \in \mathcal{X}$, find a (small) exponent vector $(e_1, \ldots, e_n)$ with $y = \prod g_i^{e_i} \star x$	Given vertices in the isogeny graph, find a (short) path connecting them.

[1] Here, one can also use more general edges not present in ”our” isogeny graph.

Attacks on REGA-DLOG

Given $x, y \in \mathcal{X}$, find small $e \in \mathbb{Z}^n$, so that $y = e \star x$.

Notation: $N = \#\mathcal{G}$ and $N_m = \#\{-m, \ldots, m\}^n = (2m+1)^n$.

Classic	Quantum
Pollard-style random walk $\mathcal{O}(\sqrt{N})$	Kuperberg $2^{\mathcal{O}(\sqrt{\log N})}$
Meet-in-the-middle [2] $\mathcal{O}(\sqrt{N_m})$	Grover / Claw finding $\mathcal{O}(\sqrt[3]{N_m})$

Idea $N_m \ll N$

- Smaller secret keys
- Faster computations

[2] In practice, $\mathcal{O}\left(\frac{N_m^{3/4}}{\sqrt{W}}\right)$ with Parallel Collision Search (PCS) is more realistic. More details later.

Ternary key spaces ($m = 1$)

Instantiation proposed in the *SQALE of CSIDH* (by Chávez-Saab, Chi-Domínguez, Jaques, Rodríguez-Henríquez ’22)

NIST Level 1: $p \approx 2^{4096}$, $N_m = 3^{139} \cong 2^{220} \ll N = 2^{2048}$.
Starting vertex is fixed to x_0.

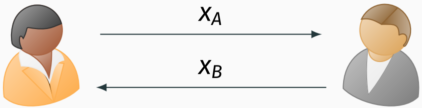

Alice: $a = (0, -1, \ldots, -1)$ Bob: $b = (1, 1, \ldots, -1)$

$\Rightarrow x_A = \prod g_i^{a_i} \star x_0$ $\Rightarrow x_B = \prod g_i^{b_i} \star x_0$

$$x_{ab} = \prod g^{a_i + b_i} \star x_0$$

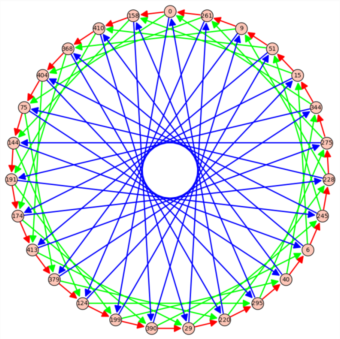

imagine a graph with 139 colors

Refined (classical) security analysis for CSIDH with ternary key spaces

Warm-up: Golden Collision Search

Given $x, y \in \mathcal{X}$, find $e \in S_n = \{-1, 0, 1\}^n$, so that $y = e \star x$.

- Write $S_{n,1} = \{-1, 0, 1\}^{n/2} \times \{0\}^{n/2}$, and $S_{n,2} = \{0\}^{n/2} \times \{-1, 0, 1\}^{n/2}$.
  $\Rightarrow$ Each $e \in S_n$ has a unique representation $e = e_1 + e_2$ with $e_i \in S_{n,i}$.
- For a hash function $H : \{0, 1\}^* \to S_{n/2}$, define $f_i : S_{n,i} \to S_{n/2}$ with

$$f_1 : e \mapsto H(e \star x), \quad f_2 : e \mapsto H(-e \star y).$$

For $y = e \star x$ and $e = e_1 + e_2$, we have $f_1(e_1) = f_2(e_2)$, the *golden collision*.

- In total: $\approx 3^{n/2} = \sqrt{N_m}$ collisions between f_1 and f_2.
  $\Rightarrow$ **Parallel Collision Search (PCS)**: Finds W collisions in time $T = \tilde{\mathcal{O}}\left(\sqrt{\sqrt{N_m} \cdot W}\right)$ with memory $M = \tilde{\mathcal{O}}(W)$.
  $\Rightarrow$ Running PCS $\mathcal{O}(\sqrt{N_m}/W)$ times, we find the golden collision.

In total: $T = \tilde{\mathcal{O}}(N_m^{3/4}/\sqrt{W}), \quad M = \tilde{\mathcal{O}}(W).$

First representation-based approach I

Given $x, y \in \mathcal{X}$, find $e \in S_n = \{-1, 0, 1\}^n$, so that $y = e \star x$.
Simplifying assumption: $\#\{i \mid e_i = a\} = n/3$ for $a \in \{-1, 0, 1\}$.

- For a parameter $\alpha \in (0, 1)$, define:

$$T_n(\alpha) = \{e \in S_n \mid \#\{i \mid e_i = a\} = \alpha \cdot n \text{ for } a = \pm 1\}.$$

Note: $e \in T_n(1/3)$.
$\Rightarrow$ Each $e \in T_n(1/3)$ has r different representations $e = e_1 + e_2$ with $e_1, e_2 \in T_n(\alpha)$, where

$$r = \binom{n/3}{n/6} \cdot \binom{n/3}{\epsilon, \epsilon, n/3 - 2\epsilon}, \quad \epsilon = (\alpha - 1/6)n.$$

- For a hash function $H : \{0, 1\}^* \to T_n(\alpha)$, define $f_i : T_n(\alpha) \to T_n(\alpha)$ with

$$f_1 : e \mapsto H(e \star x), \quad f_2 : e \mapsto H(-e \star y).$$

First representation-based approach II

For $y = e \star x$ and $e = e_1 + e_2$ one of the r representations, we have $f_1(e_1) = f_2(e_2)$, a *good collision*.

- In total: $\approx \#T_n(\alpha) = \binom{n}{\alpha n, \alpha n, (1-2\alpha)n}$ between f_1 and f_2.
  $\Rightarrow$ **PCS**: Finds W collisions in time $T = \tilde{\mathcal{O}}\left(\sqrt{\#T_n(\alpha) \cdot W}\right)$ with memory $M = \tilde{\mathcal{O}}(W)$.
  $\Rightarrow$ Running PCS $\mathcal{O}(\#T_n(\alpha)/r)$ times, we expect to find one of the good collisions.

$$\text{In total: } T = \tilde{\mathcal{O}}((\#T_n(\alpha))^{3/2}/(r\sqrt{W})), \quad M = \tilde{\mathcal{O}}(W).$$

- Given W, the optimal value for α is determined by numerical methods.

New time-memory trade-offs for ternary keys

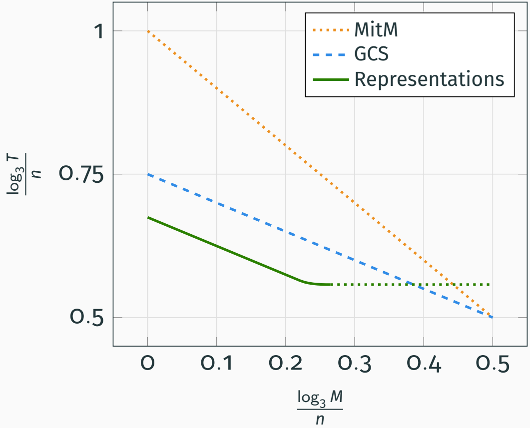

- Memoryless version:
  $T_{rep} = \tilde{O}(3^{0.675n}) < \tilde{O}(3^{0.75n}) = T_{GCS}$.
- $M \leq 3^{0.22n}$:
  $T_{rep} = \tilde{O}(3^{0.675n}/\sqrt{M})$.
- $M \geq 3^{0.265n}$:
  no more improvements.

First improvement: Partial representations

Idea: Mix of standard GCS and the first representation-based approach when M large.

- For a parameter $\delta \in (0,1)$, let:

$$e = e_1 + e_2 = (a_0, 0, c_0) + (0, a_1, c_1) = (\underbrace{a_0, a_1}_{(1-\delta)n}, \underbrace{c_0 + c_1}_{\delta n})$$

  with $a_0, a_1 \in T^{(1-\delta)n/2}(1/3),\ c_0, c_1 \in T^{\delta n}(\alpha)$. [3]

- Similar to before, we define functions

$$f_1 : T^{(1-\delta)n/2}(1/3) \times \{0\}^{(1-\delta)n/2} \times T^{\delta n}(\alpha) \to T^{(1-\delta)n/2}(1/3) \times T^{\delta n}(\alpha),$$
$$f_2 : \{0\}^{(1-\delta)n/2} \times T^{(1-\delta)n/2}(1/3) \times T^{\delta n}(\alpha) \to T^{(1-\delta)n/2}(1/3) \times T^{\delta n}(\alpha).$$

[3] This asserts proportional distribution of $1, -1$ among the three segments which can be obtained by random permutations of the indices.

Time-memory trade-offs with partial representations

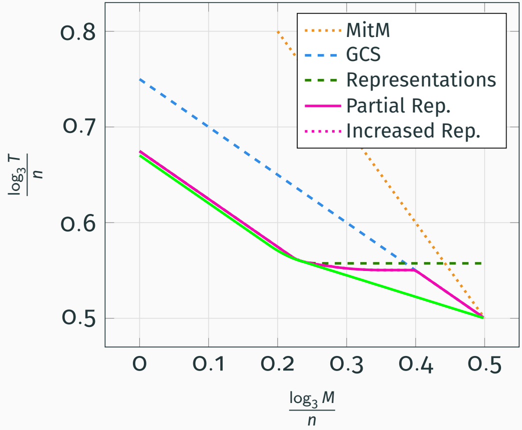

- Partial representations provide a smooth interpolation between GCS and the first representation-based approach.
- $3^{0.25n} \leq M \leq 3^{0.4n}$:
  partial representations are better than all previous methods.
- **Further improvement** by increasing the number of representations (see our paper)

Consequences for CSIDH with ternary keys

Example: NIST security level 1

$M = 2^{80} \approx 3^{50.47}$, $T = 2^{128} \approx 3^{80.76}$

Suggested parameters in the SQALE of CSIDH

$n = 139$, i.e. secret key space $\{-1, 0, 1\}^{139}$.

- $M \approx 3^{0.36n}$
- Increased representation attack:
  $T \approx 3^{0.53n} < 3^{0.57n} = T_{GCS}$

$\Rightarrow$ Security loss of around 8 bits.

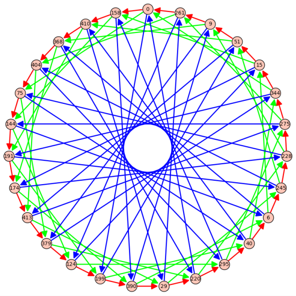

Similarly, for the parameters suggested for level 2 and level 3 security, we show a security loss of 4.57 bits and 12.75 bits, respectively.

Conclusion

Summary

- Representation-based techniques can be applied to attack CSIDH.
- This is relevant for CSIDH designs with small secret keys.

Further results in our paper

Analysis for different key spaces suggested in the CSIDH setting:

- ternary: $\{0, 1, 2\}^n$, $\{-2, 0, 2\}^n$
- non-ternary: $\{-m, \ldots, m\}^n$ for $m \in \{2, 3\}$.

Thanks for your attention!