The supersingular isogeny problem in genus ≥ 2

Benjamin Smith + Craig Costello

ECC 2019 // Bochum, DE // 2/12/2019

Inria + **École polytechnique** + Microsoft Research

(Supersingular) isogeny-based crypto

Set of supersingular elliptic curves:

$$S_1(p) := \{\mathcal{E}/\mathbb{F}_{p^2} \text{ supersingular}\} / \cong$$

Isogeny graph $\Gamma_1(\ell; p)$: vertices = $S_1(p)$, edges = ℓ-isogenies.
An $(\ell + 1)$-regular Ramanujan graph with $\#S_1(p) \approx p/12$ vertices.

Isogeny problem: given $\mathcal{E}$ and $\mathcal{E}'$ in $S_1(p)$, find a path $\mathcal{E} \to \cdots \to \mathcal{E}'$ in $\Gamma_1(\ell; p)$.

- classical algorithms: $O(\sqrt{\#S_1(p)}) = O(\sqrt{p})$
- quantum algorithms: $O(\#S_1(p)^{1/4}) = O(p^{1/4})$

Inevitable question: what happens if we do the equivalent of ECC→HECC, i.e. replace elliptic curves with g-dimensional abelian varieties?

What happens in dimension $g > 2$

Replace supersingular elliptic curves (dimension $g = 1$)
with **superspecial** g-dimensional principally polarized abelian varieties over $\mathbb{F}_{p^2}$.

$\mathcal{A}$ in $S_g(p) \implies \mathcal{A}$ is isogenous to a **product** $\mathcal{E}_1 \times \cdots \times \mathcal{E}_g$ of **supersingular** ECs.

Set $S_g(p)$ with $O(p^{g(g+1)/2})$ elements.

Graph $\Gamma_g(\ell; p)$ are connected $(\ell^{g(g+1)/2} + \cdots)$-regular graphs.

First examples of higher-dimension superspecial cryptosystems:

- Takashima hash function in $\Gamma_2(2; p)$
- Castryck–Decru–Smith hash function in $\Gamma_2(2; p)$
- Flynn–Ti SIDH analogue in $\Gamma_2(2; p)$ and $\Gamma_2(3; p)$

Expected tradeoff

Balancing graph sizes:

$$\#S_g(p) \approx \#S_1(q) \qquad \text{with } \log q \approx \frac{1}{2} g(g+1) \log p\,.$$

Implicit hypothesis in existing work:
solving isogeny problems in $\Gamma_g(\ell; p)$ *is as hard as solving them in* $\Gamma_1(\ell; q)$.

classical $O(p^{g(g+1)/4})$ with random walks,
quantum $O(p^{g(g+1)/8})$ with Grover etc.

Notice: complexities exponential in p, with exponent quadratic in g.

$\implies$ **Tradeoff**: work in dimension g and use p of much smaller bitlength.

E.g. moving from $g = 1$ to $g = 2$: use $\mathbb{F}_p$ with p one-third the size.

It doesn't work out that way

Theorem: (Costello–S. 2019): path-finding in $\Gamma_g(\ell; p)$ is only classical $O(p^{g-1})$ and quantum $O(p^{(g-1)/2})$. *Exponents linear, not quadratic, in g.*

Idea: Large subgraphs corresponding to products $\mathcal{A}_g \cong \mathcal{A}_{g-1} \times \mathcal{E}$.

1. Can walk into subgraph after $O(p^{g-1})$ short walks.
2. Recurse down into $S_1(p)^g$.
3. Solve g independent elliptic isogeny problems, take the product of the results.

Conclusion: don't do $g > 1$: tradeoff unlikely to be favourable.

Eprint: later this week.