Summary of Supersingular Hash Functions

May 24, 2019

Abstract

We give a brief overview of hash functions built on expander graphs, with a focus on supersingular isogeny graphs.

1 Construction

In this section we focus on the hash functions described in [CLG09]. The construction is based on the following theorem (see [De 17] for references).

Theorem 1. *The ℓ-isogeny graph of a supersingular elliptic curve is $\ell + 1$ regular, connected, and has the Ramanujan property.*

Let G denote the 2-isogeny graph of supersingular elliptic curves (up to isomorphism) over $\overline{\mathbb{F}}_p$ and let E_0 be some fixed curve in G. The graph G is an "expander graph". Then G is 3-connected and an expander graph. We will construct a hash function H as follows. We define H as a map from

$$H : \{\text{finite length bit strings}\} \to \left\{\text{supersingular elliptic curves over } \overline{\mathbb{F}}_p\right\}.$$

Given a length n bit string $m = m_1 \cdots m_n$, we convert it into a path in G starting at E_0. Fix an isogeny $E_{-1} \to E_0$. At the ith step in our path, there are two isogenies $E_i \to E_i^{(0)}$ and $E_i \to E_i^{(1)}$ that are not dual to $E_{i-1} \to E_i$ (this is not always possible, see Remark 3). We then set $E_{i+1} = E_i^{(m_i)}$. The output of H is E_n.

Example 2. Below is a picture of what the hash of $m = 101$ may look like.

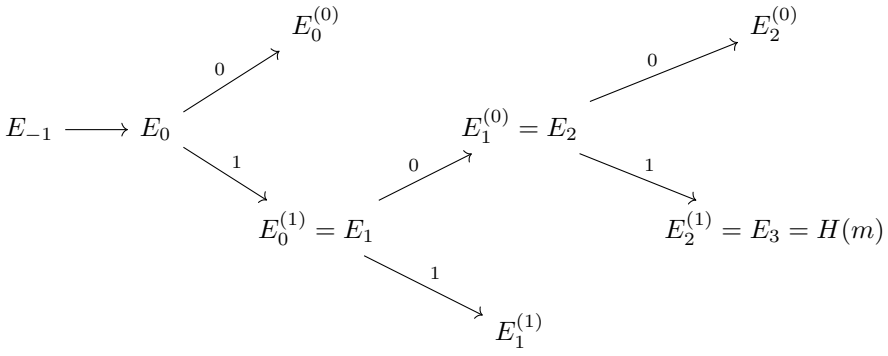

Remark 3. Technically this function is not well defined. Recall that we defined G to be elliptic curves up to isomorphism. Which means that we can only define isogenies up to isomorphism. So if $\phi : E \to E$ is an isomorphism, any isogeny $\psi : E \to E'$ must be identified with $\psi \circ \phi$. So the graph as defined may not be 3-connected. One way to fix this is to allow multi-edges, or to fix a model for every isomorphism class. But then we have the problem of enforcing some consistent choices of these edges. Another way to fix this is to restrict to primes p where no curve in G admits any extra automorphisms. This issue is discussed further in [CLG09, Sec. 4].

Remark 4. The vertices in G are usually represented by j-invariants, and the edges are defined by solutions (with multiplicity, see Remark 3) to the modular polynomial $\Phi_2(x, y)$.

2 Security

The security of a cryptographic hash function $f : A \to B$ is measured by the difficulty of the following problems:

Preimages: Given $b \in B$, find $a \in A$ such that $f(a) = b$.

Second Preimages: Given $a_1 \in A$, find some other $a_2 \in A$ such that $f(a_1) = f(a_2)$.

Collisions: Find $a_1, a_2 \in A$ such that $f(a_1) = f(a_2)$.

Indistinguishability: Distinguish between the following: $\{f(a)$ for random $a \in A\}$, and $\{$random $b \in B\}$.

Remark 5. If you can solve the preimage problem, then you can (likely) solve the second preimage problem. If you can solve the second preimage problem, then you can solve the collision problem. So we have

$$\text{preimage} \Rightarrow \text{second preimage} \Rightarrow \text{collision}.$$

2.1 Expander Graphs

In this section, we explain why we consider supersingular graphs. For references to most of the theorems here, see [De 17, Part. III].

Theorem 6. *The ℓ-isogeny graph of a supersingular elliptic curve is $\ell + 1$ regular, connected, and has the Ramanujan property (so it is an ϵ-expander graph).*

Expander graphs are good because a random walk of length $\log(\#G)$ essentially ends up at a random vertex.

Theorem 7. *Let G be k-regular ε-expander graph, and let $v \in G$. Define $B_r(v)$ the set of vertices with distance $\leq r$ from v. Then there is a constant, depending only on k and ε, such that*

$$\#B_r(v) \geq \min\left\{(1+c)^r, \#G\right\}.$$

Theorem 8. *If G is a k-regular ε-expander graph, then*

$$\frac{\varepsilon}{2}k \leq \min_{\substack{F \subseteq G \\ \#F \leq \#V/2}} \frac{\#\partial F}{\#F} \leq \sqrt{2\varepsilon} \cdot k.$$

This implies that in our graph G, the output of the hash of $\log(p)$ random bits is approximately a uniformly random supersingular curve. This is an important property, because it shows that our hash function in some sense looks like a uniformly random function. Hence, the indistinguishability property is satisfied.

Remark 9. It is also true that the endomorphism class of an ordinary elliptic curve with polynomially-small isogenies is an expander graph. However, in [CJS14], Childs, Jao, and Soukharev give a quantum subexponential time algorithm to solve the general isogeny problem between two given isogeneous ordinary curves in the same endomorphism class. Their algorithm relies on the fact that this set admits an action by an abelian group.

2.2 Reductions

The security problems for hash functions can be rephrased in terms of isogenies for our construction. See [CLG09, Thm. 1,2].

Theorem 10. *If one can find a collisions for our hash function, then one can find an elliptic curve E_1 and two distinct ℓ-power isogenies $E_0 \to E_1$.*

Theorem 11. *If one can find a preimages for our hash function, then one can find an ℓ-power isogeny $E_0 \to E_1$.*

The converses to these theorems are true if the isogenies in the conclusions are given in "factored form". That is, the isogenies are given as a sequence of ℓ-isogenies.

Remark 12. It is shown in [CLG09, Prop. 1] that an ℓ^n-isogeny on E_0 is cyclic if and only if it factors into an sequence of ℓ-isogenies that does not have any "backtracking".

3 Attacks and Defenses

3.1 Birthday Attack

There is a generic attack on the collision problem that is based on the birthday paradox.

Theorem 13 (Birthday Attack)**.** *If $f : A \to B$ is random enough, then we can find a collision in $\sqrt{\#A}$ time, and negligible memory.*

Proof. See [CLG09, Sec. 2.5.2] for references. The main idea is to use the birthday paradox. The low-memory version uses an iterative function vaguely similar to the Pollar-Rho attack on the discrete log problem. □

3.2 Frobenius

[CLG09, Sec. 7] gives a generic attack on hash functions built on expander graphs which admit an automorphism. Let G be such a graph and ψ and automorphism of G. Let $v_0 \rightsquigarrow v_n$ be a path in G. If the distance from a vertex w to $\psi(w)$ is small, then an advesary could construct two paths $v_0 \rightsquigarrow \psi(v_0) \rightsquigarrow \psi(v_n)$ and $v_0 \rightsquigarrow v_n \rightsquigarrow \psi(v_n)$, which is means they have constructed a cycle.

The supersingular isogeny graph admits a natural automorphism via Frobenius. To show that the above attack is unlikely to succeed, we have the following theorem.

Theorem 14 ([CLG09, Lem. 6])**.** *Let $n \leq \log_\ell(p/4)$. Then the number of curves $E \in G$ such that the distance from E to $E^{(p)}$ is less than n is $\tilde{O}(\ell^{n/2}\sqrt{p})$.*

3.3 Backtracking

We require the paths in G to avoid any backtracking, as this leads to collisions. This means we do not want $E_i^{(m_i)} \cong E_{i-1}$ at any step in the walk. By Lemma 15 below, this can occur if and only if $\operatorname{Aut} E_i \neq \{\pm 1\}$. The only curves with extra automorphisms have j-invariants 0 and 1728. By restricting to primes p that split in $\mathbb{Q}(i)$ and $\mathbb{Q}(\sqrt{-3})$ (i.e. $p \equiv 1 \mod 4$ and $p \equiv 1 \mod 3$), the curves corresponding to $j = 0$ and $j = 1728$ are ordinary, so they do not appear in G.

Lemma 15. *Let E be a supersingular elliptic curves over $\overline{\mathbb{F}}_p$. Let H and K be two distinct order ℓ subgroups of E. These induce separable isogenies $H \to E \to E'$ and $K \to E \to E''$. If $E \cong E''$, then E admits a non-trivial automorphism.*

Proof. We have the following diagram

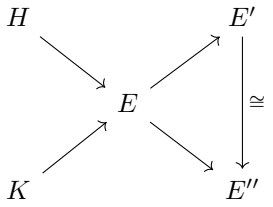

The kernel of the map $\varphi : E \to E' \to E'' \to E$, where the last map is the dual to $E \to E''$, contains H and K. Therefore it contains $E[\ell]$, and hence φ factors through multiplication by ℓ, i.e. $\varphi = \psi \circ [\ell]$. By counting degrees, ψ is an automorphism that permutes H and K. This means $\psi \notin \{\pm 1\}$. □

3.4 Small Cycles

This is covered in [CLG09, Sec. 5.3.4-5].

If G admits a small cycle, and it is possible to construct a path to the cycle from the starting curve, then collisions will be easy to find.

A cycle starting at E_0 corresponds to an endomorphism of degree ℓ^n. It is well known that End E is isomorphic to a maximal order in a quaternion algebra B ramified at p and ∞ (see Section 4).

A non-trivial cycle (one that does not allow any backtracking) corresponds to an element $\alpha \in B$ with ℓ-powered reduced norm. That means α satisfies a quadratic polynomial $x^2 - tx + \ell^n$. Let $d = t^2 - 4\ell^n$. This gives us an embedding $\mathbb{Q}(\sqrt{d}) \to B$. By Lemma 16 below, such an embedding exists if and only if p and ∞ do not split in $\mathbb{Q}(\sqrt{d})$. In particular, $d < 0$ and p satisfies some congruence condition. The fact that $d < 0$ implies that $|t| \leq 2\ell^{n/2}$. Therefore, we can ensure that no non-trivial cycles of length n occur by requiring p to not split in $\mathbb{Q}(\sqrt{t^2 - 4\ell^n})$ for all $|t| \leq 2\ell^{n/2}$. For example, if $\ell = 2$ and $p \equiv 1 \mod 840$, then there are no cycles of length 2.

Lemma 16 ([Voi18, Prop. 14.6.7])**.** *Let K be a quadratic number field and B a quaternion algebra over $\mathbb{Q}$. There is an embedding $K \to B$ if and only if every place v of $\mathbb{Q}$ that ramifies in B does not split in K (i.e. K_v is a field).*

3.5 Self Loops

Similarly to backtracking, self loops can also lead to collisions. Essentially, self loops are cycles of length 1. A self loop corresponds to an endomorphism α of degree ℓ. By the same argument as in Section 3.4, such an elemnent has a minimal polynomial of the form $x^2 - tx + \ell$ for some $|t| \leq 2\sqrt{\ell}$. If $\ell = 2$, then this means $\mathbb{Q}(i)$, $\mathbb{Q}(\sqrt{-2})$, or $\mathbb{Q}(\sqrt{-7})$ embed into the quaternion algebra B. By restricting the residue class of p, we can ensure p splits is these fields so that no such element α exists.

4 Quaternion Algebras

For references, see [Voi18].

Let $B = \mathbb{Q}_{p,\infty}$ be the unique quaternion algebra ramified at p and ∞.

Example 17. If $p \equiv 3 \mod 4$, then $B = \left(\frac{-1,-p}{\mathbb{Q}}\right)$.

It is well known that the endomorphism ring of a supersingular elliptic curve $E/\overline{\mathbb{F}}_p$ is isomorphic to a maximal order in B.

Fix such a curve E_0 and let $\mathcal{O}_0$ be a maximal order in B isomorphic to End E_0.

Theorem 18. *The isomorphism $E_0 \cong \mathcal{O}_0$ induces an equivalence of categories*

$$\{\textit{supersingular curves over } \overline{\mathbb{F}}_p\textit{, with isogenies}\}$$
$$\textit{and}$$
$$\{\textit{left } \mathcal{O}_0\textit{-ideals, with left } \mathcal{O}_0\textit{-module homomorphisms}\}.$$

The equivalence is essentially given by

$$E \mapsto \mathrm{Hom}(E, E_0).$$

An ideal $I \subseteq \mathcal{O}_0$ corresponds to the curve E_0/I (where we identify $I \subseteq \mathrm{End}\, E_0$). Moreover, $\mathrm{End}(E_0/I)$ is isomorphic to the right order of I in B.

4.1 Path Finding

Given E_1, if we can find a factored isogeny (i.e. a sequence of 2-isogenies) $E_0 \to E_1$ then we can find collisions in the CGL hash function. The equivalent problem in the quaternion setting is the following:

Problem 19 (Path finding problem in quaternion algebras)**.** Fix a maximal order $\mathcal{O}_0 \subseteq B$. Given another maximal order $\mathcal{O}_1$, find a sequence of ideals $I_1 \supseteq \cdots \supseteq I_n$ such that $[I_i : I_{i+1}] = \ell$ the right order of I_n is $\mathcal{O}_1$.

Theorem 20 ([KLPT14, GPS17])**.** *We can solve the problem above in heuristic polynomial (in $\log p$) time.*

This does not directly break the CGL hash function. To do that, we would need to transfer an instance of the isogeny path finding problem to the quaternion setting, find a solution, and then transfer the solution back to isogenies. Being able to move between the elliptic curve setting and the quaternion setting is roughly equivalent to computing endomorphism rings.

Theorem 21 ([EHL+18])**.** *Solving the path finding problem between supersingular elliptic curves is essentially equivalent to computing endomorphism rings.*

The fastest known algorithms for computing endomorphism rings (i.e. isomorphisms $\mathrm{End}\, E$ to maximal orders in B) run in exponential time.

Remark 22. A crucial detail in [EHL+18] is that there is always at least one curve $E_0 \in G$ such that it is easy to write down an explicit isomorphism between $\mathrm{End}\, E_0$ and $\mathcal{O}_0$. For example, suppose $p \equiv 3 \mod 4$, then $E_0 : y^2 = x^3 + x$ (with $j = 0$) is supersingular. We can identify B with $\left(\frac{-1,-p}{\mathbb{Q}}\right)$. The endomorphisms of E_0 given by $(x, y) \mapsto (x^p, y^p)$ and $(x, y) \mapsto (-x, \sqrt{-1}y)$ with j, i respectively. We can then check whether $(1+i)/2 \in \mathcal{O}_0$ by checking whether $1+i$ corresponds to an isogeny whose kernel contains $E[2]$. This gives an efficient way to build an explicit isomorphism between $\mathrm{End}\, E_0$ and a maximal order for this particular curve.

Theorem 23. *Given an isomorphism* $\operatorname{End} E_0 \to \mathcal{O}_0$ *and an* ℓ*-isogeny* $E_0 \to E_1$*, we can compute an isomorphism* $\operatorname{End} E_1 \to \mathcal{O}_1 \subseteq B$ *in polynomial time (polynomial in* ℓ *and* $\log p$*).*

Proof. Let $\psi : \operatorname{End} E_0 \to \mathcal{O}_0$ be the isomorphism and $\varphi : E_0 \to E_1$ be the isogeny. Our algorithm goes as follows:

1. Convert the isogeny $E_0 \to E_1$ to a corresponding left-ideal $I \subseteq \mathcal{O}_0$
2. Compute the right order $\mathcal{O}_1$ of I.
3. Compute an isomorphism map $\mathcal{O}_1 \to \operatorname{End} E_1$.

The first step can be done by solving a DLP in $E_0[\ell]$. That is, pick a $\mathbb{Z}$-basis for $\mathcal{O}_0$, say $\alpha_1, \ldots, \alpha_4$. The action of $\phi_i = \psi^{-1}(\alpha_i)$ on $E_0[\ell]$ can be represented by matrices in $M_2(\mathbb{Z}/\ell\mathbb{Z})$. The kernel of φ corresponds to some 1-dimensional subspace W in $(\mathbb{Z}/\ell\mathbb{Z})^2$, so we choose linear combinations of the ϕ_i whose kernel is W. Repeated enough, and applying a basis reduction algorithm, this should give us a basis for an left-ideal I with reduced norm ℓ of $\mathcal{O}_0$ corresponding to φ.

The second step is a standard algorithm in quaternion algebras, see [KV10] for references.

The isomorphism $\mathcal{O}_1 \to \operatorname{End} E_1$ given by (see [Voi18, Lem. 42.2.9])

$$\alpha \mapsto \frac{1}{\ell^2}\left(\varphi \circ \psi^{-1}(\ell\alpha) \circ \hat{\varphi}\right).$$

Here $\frac{1}{\ell^2}$ means that $\ker \varphi \circ \psi^{-1}(\ell\alpha) \circ \hat{\varphi}$ contains $E_1[\ell^2]$ so it can be factored as $E_1 \xrightarrow{\ell^2} E_1 \xrightarrow{\rho} E_1$, and the map ρ is the output. □

4.2 Short Cycles

Here is another example showing how computing endomorphism rings can break CGL.

In Section 4.2 of `https://eprint.iacr.org/2017/962.pdf`, Petit and Lauter give an algorithm that, given an isomorphism of $\operatorname{End} E_0$ and a maximal order $\mathcal{O}_0$ of the quaternion algebra $\mathbb{Q}_{p,\infty}$, product a short cycle containing E_0. In their article, they mention this gives rise to a possible backdoor. If the hash function publisher knows an isomorphism $\operatorname{End} E_0 \cong \mathcal{O}_0$ and no one else does, then they could compute collisions.

Remark 24. It would be interesting to find applications of this protocol when the endomorphism ring of the "starting" curve is private and everything else is public. See for example [GPS17].

References

[CJS14] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *J. Math. Cryptol.*, 8(1):1–29, 2014.

[CLG09] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *J. Cryptology*, 22(1):93–113, 2009.

[De 17] Luca De Feo. Mathematics of Isogeny Based Cryptography. *arXiv e-prints*, page arXiv:1711.04062, Nov 2017.

[EHL+18] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In *Advances in cryptology—EUROCRYPT 2018. Part III*, volume 10822 of *Lecture Notes in Comput. Sci.*, pages 329–368. Springer, Cham, 2018.

[GPS17] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In *Advances in cryptology—ASIACRYPT 2017. Part I*, volume 10624 of *Lecture Notes in Comput. Sci.*, pages 3–33. Springer, Cham, 2017.

[KLPT14] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ-isogeny path problem. *LMS J. Comput. Math.*, 17(suppl. A):418–432, 2014.

[KV10] Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM J. Comput.*, 39(5):1714–1747, 2010.

[Voi18] John Voight. Quaternion algebras. Published online, July 7 2018. `https://math.dartmouth.edu/~jvoight/quat-book.pdf`.