Bruhat–Tits trees and supersingular elliptic curves

Jana Sotáková
University of Amsterdam/QuSoft

j.s.sotakova@uva.nl

March 23, 2021

Abstract

These are notes accompanying the two talks I gave about my joint work with Laia Amorós, Annamaria Iezzi, Kristin Lauter, and Chloe Martindale, *Explicit connections between supersingular isogeny graphs and Bruhat–Tits trees*([1]. The first talk was at the Leiden Algebra, Geometry, and Number Theory Seminar on March 15, 2021 and the second at the RTG seminar at Clemson University on March 22, 2021. I would like to thank the organizers of these seminars to give me an opportunity to speak and share my research at their seminars.

Contents

1 Dictionary

Here we give a short dictionary of the terms we will need in the following discussion. Everything happens over $\mathbb{F}_{p^2}$, which is without loss of generality for supersingular elliptic curves.

Finite field $\mathbb{F}_{p^2}$ for a prime $p > 3$, for simplicity also $p \equiv 3 \bmod 4$.

Elliptic curve $E : y^2 = x^3 + ax + b$ for some $a, b \in \mathbb{F}_{p^2}$ with $4a^3 + 27b^2 \neq 0$. We consider elliptic curves up to isomorphism, that is, an invertible change of coordinates $x, y \mapsto ax + by, cx + dy$: the 'same' equations describes the same geometric object.

Rational points $E(\mathbb{F}_{p^2}) = \{P = (u, v) \in \mathbb{F}_{p^2}^2 : v^2 = u^3 + au + b\} \cup \{O_E\}$ form a finite abelian group, with the point at infinity O_E being the neutral element.

supersingular elliptic curves , pick your favorite definition. For instance, $p \mid |E(\mathbb{F}_{p^2})| - 1$. Or $E(\mathbb{F}_{p^2}) \cong \mathbb{Z}/(p+1) \times \mathbb{Z}/(p+1)$ (this is without loss of generality).

j-invariant $j(E) = 1728\frac{4a^3}{4a^3+27b^2}$, used as labels for supersingular elliptic curves, for which $j(E) \in \mathbb{F}_{p^2}$. The j-invariant also gives the smallest field over which an elliptic curve can be defined, so all supersingular elliptic curves can be defined over $\mathbb{F}_{p^2}$.

Torsion subgroups: we denote by $E[\ell^k]$ the points of E of order dividing ℓ^k. In general, these points may not be defined over $\mathbb{F}_{p^2}$ but over higher extensions, but for the SIDH setting, the points we will consider will be defined over $\mathbb{F}_{p^2}$.

Isogeny is a rational map $\varphi : E_A \to E_B$, which is also a group homomorphism. Isogenies work just like quotient maps in groups:

- $\ker \varphi$ is a subgroup of E_A,
- for any subgroup H get a map $E_A \to E_A/H$ with kernel H,
- degree $\deg \varphi = \#\ker \varphi$ for separable isogenies
- isogeny of degree $\ell \longleftrightarrow$ subgroup of size ℓ in $E_A[\ell]$

2 Motivation

We examine the 'extra' information coming from isogeny-based protocols.

SIDH key exchange

In isogeny-based cryptography, Alice and Bob communicate on a public channel (which anyone can listen in on) and want to obtain an elliptic curve E_{AB} using isogenies, in a way that nobody can compute E_{AB} just by looking at their communication.

In the SIDH protocol, Alice and Bob proceed as follows:

1. First, they agree on a starting elliptic curve E.
2. Alice then chooses a secret point $R_A \in E_0[2^a]$ and computes the secret isogeny $\varphi_A : E \to E_A = E/\langle R_A \rangle$, Bob similarly chooses a secret point in $R_B \in E[3^b]$ and computes the secret isogeny $\varphi_B : E \to E_B = E/\langle R_B \rangle$.

3. Alice then sends her elliptic curve E_A to Bob, Bob sends his elliptic curve E_B to Alice.
4. Bob then computes $E_A/\langle R_B \rangle$ and Alice computes $E_B/\langle R_A \rangle$.
5. Now both Alice and Bob share the elliptic curve $E_{AB} = E/\langle R_A, R_B \rangle$

However, to compute E_{AB} from E_A, Bob needs to know $\varphi_A(R_B)$ the image of his point R_B under Alice's isogeny φ_A. Either is a secret which they cannot share directly.

Torsion points

To circumvent this, in the setup, Alice and Bob also agree on public bases for the torsion groups $P_A, Q_A \subset E[2^a]$ and $P_B, Q_B \subset E[3^b]$.

- Bob chooses his secret point as $R_B = P_B + r_B Q_B$;
- Alice publishes not just E_A but also $\varphi_A(P_B), \varphi_A(Q_B)$;
- Bob can compute $\varphi_A(R_B) = \varphi_A(P_B) + r_B \varphi_A(Q_B)$.

Therefore, the SIDH key exchange setup does not just specify elliptic curves from which we compute isogenies, but also bases of (rather large) torsion groups.

Known endomorphism rings

Essentially the only way we can find supersingular elliptic curves is by starting from a known supersingular elliptic curve and computing some isogeny. However, this procedure means that we almost always known the endomorphism ring: not just as an abstract ring but with an explicit representation of the generators and endomorphisms on E.

Certainly for the starting curves in SIKE, we do know the endomorphism ring and by constructing their isogenies, Alice or Bob can also compute the endomorphism rings of their curves E_A, E_B.

Isogeny graphs vs extra information

Bob computes his isogeny

$$\varphi_B : E \to E_B = E/\langle R_B \rangle$$

as a sequence of ℓ-isogenies, hence giving him a secret path of length n in the supersingular isogeny graph $\mathcal{G}_\ell$:

- vertices are elliptic curves, up to isomorphism,
- edges are ℓ-isogenies.

In this note, we argue that Bruhat–Tits trees are better at bookkeeping all this information than just considering supersingular isogeny graphs, which keep track of the elliptic curves.

$$\begin{aligned} \text{Bruhat–Tits tree } \mathcal{T}_\ell &\twoheadrightarrow \text{supersingular isogeny graph } \mathcal{G}_\ell \\ E, \mathrm{End}(E), \langle P_A, Q_A \rangle, &\mapsto E \end{aligned}$$

3 Background

We revisit some of the background information on elliptic curves and endomorphism rings.

3.1 Endomorphism rings

For simplicity, we let $p \equiv 3 \bmod 4$.

Definition 3.1 (The quaternion algebra $B_{p,\infty}$)**.** *The quaternion algebra $B_{p,\infty}$ is the $\mathbb{Q}$-algebra generated by $1, i, j, k$ with $i^2 = -1$ and $j^2 = -p$ and $ij = -ji = k$. Note that also $k^2 = -p$.*

Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve. Then $\mathrm{End}(E)$ can be embedded as a maximal order in the quaternion algebra $B_{p,\infty}$. The identification is as follows:

- an endomorphism φ corresponds to some $\alpha = a + bi + cj + dk$ with $a, b, c, d \in \mathbb{Q}$,
- the dual endomorphism $\hat{\varphi}$ is then $\overline{\alpha} = a - bi - cj - dk$,
- the degree $\deg \varphi$ is the reduced norm $\mathrm{nrd}(\alpha) = a^2 + b^2 + pc^2 + pd^2$.
- composition of isogenies is the multiplication in $B_{p,\infty}$.

Example 3.2. *The endomorphism ring of the curve $E : y^2 = x^3 + x$:*
Define $i, j : E \to E$ as the maps $i(x, y) = (-x, \sqrt{-1}y)$ and $j(x, y) = (x^p, y^p)$. Then we have

$$\mathrm{End}(E) = \left\langle 1, i, \frac{i+j}{2}, \frac{1+k}{2} \right\rangle_{\mathbb{Z}} = \left\langle \frac{1+j}{2}, \frac{i+k}{2}, j, k \right\rangle_{\mathbb{Z}}.$$

So for E, we not only know the endomorphism ring abstractly, but also as explicit maps on E. This is a very strong property that allows one to do more than for just any supersingular elliptic curve. A convincing example of what happens if we do know information about torsion points and understand the endomorphism ring well, is the paper [2]. We also note that knowing both $\mathrm{End}(E)$ and $\mathrm{End}(E')$, we can compute the isogeny $\varphi : E \to E'$ [4].

Theorem 3.3 (Deuring's correspondence, first version)**.** *For every maximal order $\mathcal{O} \subset B_{p,\infty}$ there is a supersingular elliptic curve E with*

$$\mathrm{End}(E) \cong \mathcal{O}.$$

Two elliptic curve E, E' satisfy $\mathrm{End}(E) \cong \mathrm{End}(E')$ *(as maximal orders, that is, they are conjugate by an element in $B_{p,\infty}$) if and only if*

$$j(E) = j(E') \text{ or } j(E) = j(E')^p.$$

Write $\mathcal{O} = \mathrm{End}(E)$ and $\mathcal{O}' = \mathrm{End}(E')$. An isogeny $\varphi : E \to E'$ gives us an ideal $I \subset \mathcal{O}$:

$$I = \{\alpha \in \mathcal{O} : \alpha(P) = O_{E'} \text{ for all } P \in \ker \varphi\}$$

This ideal satisfies

$$O_R(I) = \mathcal{O}, \quad O_L(I) = \mathcal{O}'$$

Such ideals are called connecting ideals for the maximal orders $\mathcal{O}, \mathcal{O}'$.

In simpler words, the correspondence between isogenies and ideals is as follows:

$\varphi : E \to E'$ an isogeny	right ideal $I \subset \mathcal{O}$
$\ker \varphi$	torsion subgroup $E[I]$
post-composing with $\beta \in \mathrm{End}(E')$:	multiplying $\beta \cdot I$
$\beta \circ \varphi : E \to E' \to E'$	$\beta \cdot I \sim I$

By post-composing with an endomorphism on E', we do not change the kernel of φ. 'Morally speaking', the isogeny is 'the same'. This vague reasoning can be made precise:

We define the class set $\mathrm{Cl}(\mathcal{O}) = \{[I] : I \sim J \text{ if and only if } I = \alpha J \text{ for } \alpha \in B_{p,\infty}^\times\}$.

Theorem 3.4 (Deuring's correspondence, second version)**.** *Let E be a supersingular elliptic curve and identify* $\mathrm{End}(E) = \mathcal{O} \subset B_{p,\infty}$. *Consider the class set* $\mathrm{Cl}(\mathcal{O})$ *of invertible right-ideals of $\mathcal{O}$.*

The set $\mathrm{Cl}(\mathcal{O})$ *is in a bijection with supersingular j-invariants in $\mathbb{F}_{p^2}$.*

The bijection induces the following translations between isogenies and ideals:

elliptic curve E	order $\mathcal{O} = \mathrm{End}(E)$
isogeny $\varphi : E \to E'$	ideal $I \subset \mathcal{O}$
ℓ-isogeny	ideal of norm ℓ
$\mathrm{End}(E')$	$O_L(I) = \{\beta \in B_{p,\infty} : \beta \cdot I \subset I\}$.

For two elliptic curves E, E', we briefly explain how to use the knowledge of endomorphism rings to obtain an isogeny between them.

1. If we know $\mathrm{End}(E) = \mathcal{O}$ and $\mathrm{End}(E') = \mathcal{O}'$, it is easy to find a *connecting ideal*: for instance, if $M = [\mathcal{O} : \mathcal{O} \cap \mathcal{O}']$, then

   $$I(\mathcal{O}, \mathcal{O}') = \{\alpha \in B_{p,\infty} : \alpha \mathcal{O}' \bar{\alpha} \subset M\mathcal{O}'\}$$

   is a connecting ideal;

2. it is 'easy' to find an equivalent ideal of smooth norm, using the KLPT algorithm [4];
3. from an ideal of smooth norm, it is easy to reconstruct an isogeny $E \to E'$, as a sequence of isogenies of small norm.

Finally, we need the following information about localizing the endomorphims rings of supersingular elliptic curves. Recall the notation: $E/\mathbb{F}_{p^2}$ supersingular elliptic curve, identify $\mathrm{End}(E) \cong \mathcal{O}$ maximal order in the quaternion algebra $B_{p,\infty}$.

Lemma 3.5 (Localizing). *For $\ell \neq p$:*

1. *we have $B_{p,\infty} \otimes \mathbb{Z}_\ell = M_2(\mathbb{Q}_\ell)$,*
2. *maximal order $\mathcal{O} \subset B_{p,\infty}$ embeds as a maximal order $\mathcal{O}_\ell = \mathcal{O} \otimes \mathbb{Z}_\ell \subset M_2(\mathbb{Q}_\ell)$,*
3. *all maximal orders in $M_2(\mathbb{Q}_\ell)$ are conjugate to $M_2(\mathbb{Z}_\ell)$.*

One way to think about this: $M_2(\mathbb{Z}_\ell)$ is then endomorphism ring of the standard lattice $\mathbb{Z}_\ell \times \mathbb{Z}_\ell$, which we write as the identity matrix. If the maximal order $\mathcal{O}_\ell$ is conjugate to $M_2(\mathbb{Z}_\ell)$ by A, then $\mathcal{O}_\ell$ is the endomorphism ring of the lattice A.

Next, we identify this lattice A as the Tate module of E.

3.2 Tate module

Let $\ell \neq p$ be a prime. For any n, we know that (as abelian groups), we have

$$E[\ell^n] \cong \mathbb{Z}/\ell^n\mathbb{Z} \times \mathbb{Z}/\ell^n\mathbb{Z}.$$

Moreover, there are connecting maps $E[\ell^n] \to E[\ell^{n-1}]$, which are multiplications by $[\ell]$, and are surjective. By taking a projective limit over the system of $E[\ell^n]$ with the connecting maps, we get the Tate module:

$$T_\ell(E) = \varprojlim E[\ell^n].$$

Abstractly, we have $T_\ell(E) \cong \mathbb{Z}_\ell \times \mathbb{Z}_\ell$, a $\mathbb{Z}_\ell$-lattice of rank 2.

Because the torsion subgroups $E[\ell^n]$ may only be defined over an extension of $\mathfrak{f}_{p^2}$, in general, there is also a Galois action on $T_\ell(E)$. However, for supersingular elliptic curves (over $\mathbb{F}_{p^2}$), the Frobenius action is by a scalar. In particular, the Frobenius action commutes with everything, so all maps are Galois-equivariant maps.

Theorem 3.6 (Endomorphisms and Tate modules). *For supersingular elliptic curve $E, E'/\mathfrak{f}_{p^2}$ we have*

$$\mathrm{End}(T_\ell(E)) \cong \mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Z}_\ell$$
$$\mathrm{Hom}(E, E') \otimes_Z \mathbb{Z}_\ell \cong \mathrm{Hom}(T_\ell(E), T_\ell(E'))$$

This means that isogenies of degree ℓ^k correspond to maps of Tate modules, which are just linear maps over $\mathbb{Z}_\ell$. Moreover, remember that $\mathrm{End}(E) \otimes \mathbb{Z}_\ell \cong M_2(\mathbb{Z}_\ell)$.

We note that despite Tate modules being ℓ-adic, having a basis $(P_n, Q_n) \subset E[\ell^n]$ allows us to reduce everything to $M_2(\mathbb{Z}/\ell^n\mathbb{Z})$, which we can represent by integer matrices. This has the following important consequence: isogenies of degree dividing ℓ^k with $k \leq n$ give matrices in $M_2(\mathbb{Z})$ with determinant ℓ^k.

4 Bruhat–Tits tree

Remember that the localization of the endomorphism ring $\mathrm{End}(E) \otimes \mathbb{Z}_\ell$ is the endomorphism ring of the $\mathbb{Z}_\ell$-lattice $T_\ell(E)$. We look deeper into $\mathbb{Z}_\ell$-lattices.

Consider homothethy classes of $\mathbb{Z}_\ell$-lattices (of rank 2) $\Lambda \subset \mathbb{Q}_\ell^2 = V$:

$$\begin{aligned} \Lambda &= \mathbb{Z}_\ell u + \mathbb{Z}_\ell v \quad \text{for some } u, v \in V \text{ linearly independent} \\ &= \langle u, v \rangle \\ \Lambda \sim \Lambda' \longleftrightarrow \Lambda &= \alpha \Lambda' \text{ for some } \alpha \in \mathbb{Q}_\ell^* \end{aligned}$$

Definition 4.1 (Adjacent lattices)**.** *The homothethy classes of lattices Λ, Λ' are adjacent if for some representatives*

$$\ell\Lambda \subsetneq \Lambda' \subsetneq \Lambda.$$

Being adjacent is a symmetric relation (multiply all the lattices in the definition by ℓ and remember that $\Lambda \sim \ell\Lambda$). Because $[\Lambda : \ell\Lambda] = \ell^2$, being adjacent means that we can choose representatives Λ and Λ' such that $\Lambda' \subset \Lambda$ with exact index ℓ.

Definition 4.2 (The Bruhat–Tits tree)**.** *The Bruhat–Tits tree $\mathcal{T}_\ell$ is the graph:*

- *with vertices given by homothethy classes of lattices,*
- *edges between adjacent classes of lattices.*

Theorem 4.3. *The Bruhat–Tits tree $\mathcal{T}_\ell$ is a $(\ell+1)$-regular infinite tree for every ℓ.*

We take another look at the definition, which we can compare to a picture of the Bruhat–Tits tree in Figure 1.

Let $\Lambda = \langle u, v \rangle = \mathbb{Z}_\ell u + \mathbb{Z}_\ell v$ represent a vertex in $\mathcal{T}_\ell$. We can write it as a matrix (using column vectors)

$$M = (u|v) \in GL_2(\mathbb{Q}_\ell).$$

Then $GL_2(\mathbb{Q}_\ell)$ acts on the lattices from the left. One can show that $GL_2(\mathbb{Q}_\ell)$ acts transitively on the set of lattices, so by computing the stabilizer of any vertex, we can compute:

$$\text{vertices of } \mathcal{T}_\ell = GL_2(\mathbb{Q}_\ell) / \left(\mathbb{Q}_\ell^* GL_2(\mathbb{Z}_\ell)\right).$$

This is the same as simply consider matrices up to scaling and a change of basis.

The adjacent lattices are given by

$$\langle \ell u, v \rangle \quad \text{and} \quad \langle u + iv, \ell v \rangle \text{ for } i = 0, \ldots, \ell - 1$$

(this is one of the choices and it depends on the choice of a particular basis (u, v)). In matrix form, we can get all the neighboring vertices as

$$M \cdot \left(\begin{smallmatrix} \ell & 0 \\ 0 & 1 \end{smallmatrix}\right) \quad \text{and} \quad M \cdot \left(\begin{smallmatrix} 1 & 0 \\ i & \ell \end{smallmatrix}\right), \text{ for } i = 0, \ldots, \ell - 1.$$

We call the first matrix the direction infinity and the latter directions i. Note that a path in the Bruhat–Tits tree is then a sequence of these direction, which we can either represent by the directions, individual matrices, or a product of the matrices.

The following is the essential result about Bruhat–Tits trees for isogeny cryptographers:

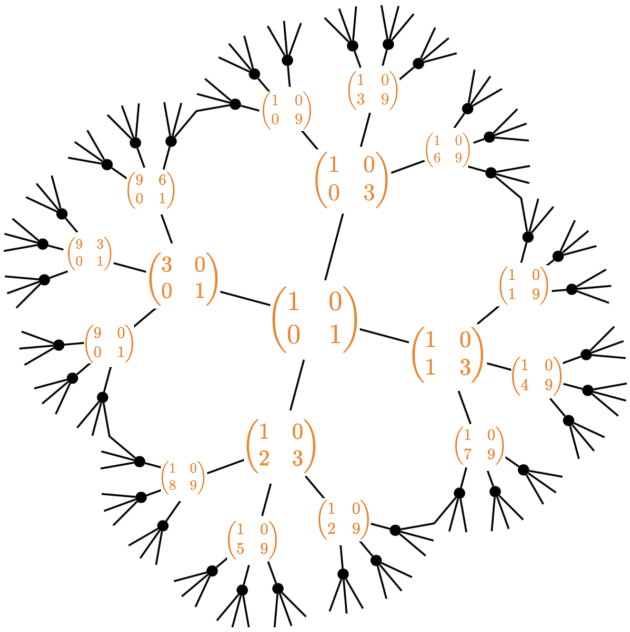

Figure 1: The Bruhat–Tits tree $\mathcal{T}_3$ for $\ell = 3$. Picture from [1].

Theorem 4.4. *For a certain $\Gamma \subset GL_2(\mathbb{Q}_\ell)$, there are bijections*

$$\Gamma\backslash\mathcal{T}_\ell \longleftrightarrow \mathrm{Cl}(\mathcal{O}) \longleftrightarrow \textit{supersingular isogeny graph } \mathcal{G}_\ell.$$

We make this Γ explicit in a moment in Section 4.2.

4.1 Mapping elliptic curves to the Bruhat–Tits tree

Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve. To map E to the Bruhat–Tits tree, we map it to its Tate module T_ℓ, considered as a lattice in $\mathbb{Q}_\ell^2$. Isogenies of elliptic curves map to linear map of lattices.

This translation is rather explicit. For instance, if $P, Q \in T_\ell(E)$ are a basis of the Tate module and $P_1, Q_1 \in E[\ell]$ the projections onto the ℓ-torsion, then the isogeny $E \to E/\langle P + iQ\rangle$ corresponds to the edge $\left(\begin{smallmatrix} 1 & 0 \\ i & \ell \end{smallmatrix}\right)$.

Moreover, if we only have a basis of $E[\ell^n]$, say P_n, Q_n, we know that it can be extended to some basis $P, Q \subset T_\ell(E)$, which identifies

$$T_\ell(E) \longleftrightarrow \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \in \mathcal{T}_\ell$$
$$\mathrm{End}(E) \otimes \mathbb{Z}_\ell \longleftrightarrow M_2(\mathbb{Z}_\ell)$$

But if we only have P_n, Q_n, we can still recover the tree $\mathcal{T}_\ell$ up to distance n. We call this the truncated tree.

Let us now specialize to the SIKE world: take prime $p = 2^a \cdot 3^b - 1$. This gives us naturally two large torsion subgroups $E[\ell^n]$: the groups $E[2^a]$ and $E[3^b]$ (the discussion is the same for either).

We start with the elliptic curve $E : y^2 = x^3 + x/\mathbb{F}_{p^2}$ (in fact, $y^2 = x^3 + 6x^2 + x$, but the discussion is the same).

We are given a basis $P_n, Q_n \subset E[\ell^n]$, so we can construct[1] the truncated Bruhat–Tits tree with $T_\ell(E) \leftrightarrow \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$.

When generating secret points, Alice/Bob choose $0 \le m < \ell^n$ and compute

$$R = P_n + mQ_n.$$

This is the same as ignoring the branch from E that goes in the direction ∞.

Then

$$E/\langle R\rangle \longleftrightarrow \left(\begin{smallmatrix} 1 & 0 \\ m & \ell^n \end{smallmatrix}\right).$$

If we write $m = \sum_j i_j \ell^j$, then the path $E \to E/\langle R\rangle$ corresponds to walking on the Bruhat–Tits tree by following directions $(i_0, \ldots, i_{n-1})$.

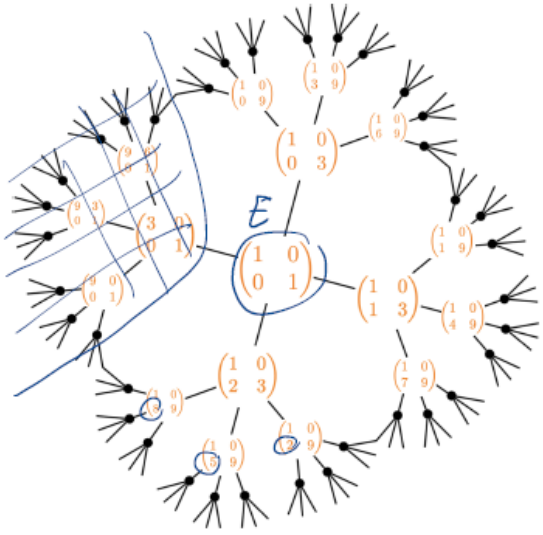

Figure 2: The *SIKE-tree* for $\ell = 3$.

We give the vague name SIKE-tree (cf. Figure 4.1) to the subtree of the Bruhat–Tits tree spanned by the root E and all the possible public keys $E/\langle P_n + mQ_n\rangle$.

Remember that going from the Bruhat–Tits tree $\mathcal{T}_\ell$ to the supersingular isogeny graph $\mathcal{G}_\ell$, we need to quotient by some group Γ. However:

Theorem 4.5 ([5])**.** *The SIKE-tree as a subtree of $\mathcal{T}_\ell$ is basically the same as the corresponding subgraph of $\mathcal{G}_\ell$.*

That is, the when we map the SIKE-tree to $\mathcal{G}_\ell$, we obtain an almost-tree: typically only a few (0 or 2) pairs of leaves will be glued.

4.2 Quotients of the Bruhat–Tits tree

Fix an embedding $\Phi : B_{p,\infty} \to M_2(\mathbb{Q}_\ell)$. Typically we take an embedding induced by the identification $\mathrm{End}(E) \cong \mathcal{O} \hookrightarrow M_2(\mathbb{Z}_\ell)$ because we want to keep track of the endomorphism rings.

[1] That is, describe, set up, we can't enumerate the tree for cryptographic sizes.

Then we can describe the supersingular isogeny graph as a quotient of the Bruhat–Tits tree as:

$$\Gamma\backslash\mathcal{T}_\ell = \Gamma\backslash GL_2(\mathbb{Q}_\ell)/\mathbb{Q}_\ell^* GL_2(\mathbb{Z}_\ell)$$

for $\Gamma = \Phi(\mathcal{O}[1/\ell]^\times)$.

We can replace Γ by $\bigcup_k \Phi(U_k)$ with $U_k = \{\alpha \in \mathcal{O} : \text{norm}(\alpha) = \ell^k\}$. If we only want the tree up to distance n, we only need to consider the first n such sets.

Moreover, replacing Γ by $\Gamma_+ = \bigcup_k \Phi(U_{2k})$, the graph $\Gamma_+\backslash\mathcal{T}_\ell$ is the graph of the special fiber of a certain Shimura curve. Sage can compute with these in the BTQuotient module based on [3]. This code works by first computing the whole quotient $\Gamma_+\backslash\mathcal{T}_\ell$, which is of course impossible for cryptographic sizes, but that is easily adaptable.

Moreover, because we start with the embedding $\mathcal{O} \hookrightarrow M_2(\mathbb{Z}_\ell)$, we have enough tools to translate the local information we get from the Bruhat–Tits tree (such as changing the endomorphism rings of the lattices, etc) to global information, like connecting ideals, endomorphism rings of elliptic curves, norm forms.

5 Conclusions

These are our main selling points:

- Isogeny problems naturally come with information about torsion points and endomorphism rings, and unlike supersingular isogeny graphs, Bruhat–Tits trees can keep track of this information.

- There are bijections:

  $$\Gamma\backslash\mathcal{T}_\ell \cong \text{Cl}(\mathcal{O}) \cong \text{supersingular } \ell\text{-isogeny graph } \mathcal{G}_\ell.$$

  The latter is well-known to isogeny-mathematicians/cryptographers and has yielded amazing results. We propose going to Bruhat–Tits trees and study this amazing object which we believe will yield more information about isogenies.

  The first bijection is also known to people who study quaternion algebras and Shimura curves: if we ask them what they know, we might obtain yet new perspectives on our isogeny problems.

- It is easy to compute with the Bruhat–Tits tree: everything in terms of matrices in $M_2(\mathbb{Z}_\ell)$; can be lifted to integer matrices (with some precision). But this is the same as building the tree from a basis of $E[\ell^n]$, which is exactly what we get in isogeny-problems, so no information is lost.

- Sage can compute with Bruhat–Tits trees in the BTQuotient module. This code is easy to extend to be able to experiment with examples of cryptographic size.

- Because of our setup, we can translate local information to global information. Walking on the Bruhat–Tits tree, we can keep track of arithmetic information, e.g. endomorphism rings and norm forms. This is the same as using quaternion algebras, bu Bruhat–Tits trees can also take torsion points into account.

- (Our most speculative application/current promising project): Bruhat–Tits trees come with directions: $\left(\begin{smallmatrix} 1 & 0 \\ m & \ell^n \end{smallmatrix}\right) \longleftrightarrow (i_0, \ldots, i_{n-1})$ for $m = \sum_j i_j \ell^j$. These are dependent on various choices, however, *these choices have already been made in the setup of SIDH/SIKE protocols*. Our goal: are elliptic curves in one direction different from other directions? Can we predict how the arithmetic invariants of elliptic curves will change depending on which directions we take (that is, without computing it first)?

We are continuing to work in this amazing new direction and hope to give you more results soon. In any case, we do believe Bruhat–Tits trees are worth learning about and will lead to new perspectives on isogeny graphs.

References

[1] Laia Amorós, Annamaria Iezzi, Kristin Lauter, Chloe Martindale, and Jana Sotáková. Explicit connections between supersingular isogeny graphs and bruhat–tits trees. Cryptology ePrint Archive, Report 2021/372, 2021. `https://eprint.iacr.org/2021/372`.

[2] Victoria de Quehen, Péter Kutas, Chris Leonardi, Chloe Martindale, Lorenz Panny, Christophe Petit, and Katherine E. Stange. Improved torsion point attacks on sidh variants. Cryptology ePrint Archive, Report 2020/633, 2020. `https://eprint.iacr.org/2020/633`.

[3] Cameron Franc and Marc Masdeu. Computing fundamental domains for the Bruhat-Tits tree for $\mathrm{GL}_2(\mathbb{Q}_p)$, p-adic automorphic forms, and the canonical embedding of Shimura curves. *LMS Journal of Computation and Mathematics*, 17(01):1–23, 2014.

[4] 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.

[5] Hiroshi Onuki, Yusuke Aikawa, and Tsuyoshi Takagi. The existence of cycles in the supersingular isogeny graphs used in sike. Cryptology ePrint Archive, Report 2020/439, 2020. `https://eprint.iacr.org/2020/439`.