Cycles and Cuts in Supersingular L-Isogeny Graphs

Sarah Arpin[a], Ross Bowden[b], James Clements[b], Wissam Ghantous[c], Jason T. LeGrow[a], Krystal Maughan[d]

[a]*Department of Mathematics, Virginia Tech, Blacksburg, Virginia, USA*
[b]*School of Computer Science, University of Bristol, Bristol, United Kingdom*
[c]*Department of Mathematics, University of Central Florida, Orlando, Florida, USA*
[d]*Department of Computer Science, University of Vermont, Burlington, Vermont, USA*

Abstract

Supersingular elliptic curve isogeny graphs underlie isogeny-based cryptography. For isogenies of a single prime degree ℓ, their structure has been investigated graph-theoretically. We generalise the notion of ℓ-isogeny graphs to L-isogeny graphs (studied in the prime field case by Delfs and Galbraith), where L is a set of small primes dictating the allowed isogeny degrees in the graph. We analyse the graph-theoretic structure of L-isogeny graphs. Our approaches may be put into two categories: cycles and graph cuts.

On the topic of cycles, we provide: a count for the number of cycles in the L-isogeny graph with cyclic kernels using traces of Brandt matrices; an efficiently computable estimate based on this approach; and a third ideal-theoretic count for a certain subclass of L-isogeny cycles. We provide code to compute each of these three counts.

On the topic of graph cuts, we compare several algorithms to compute graph cuts which minimise a measure called the *edge expansion*, outlining a cryptographic motivation for doing so. Our results show that a *greedy neighbour* algorithm out-performs standard spectral algorithms for computing optimal graph cuts. We provide code and study explicit examples.

Furthermore, we describe several directions of active and future research.

Keywords: 14K02, 14H52, 11T71, 11S31

1. Introduction

Isogeny-based cryptography has spurred significant research into supersingular elliptic curve isogeny graphs. Research advancements on the underlying structures are increasingly critical as cryptographic protocols evolve. Supersingular elliptic curve ℓ-isogeny graphs are $(\ell+1)$-regular, Ramanujan [1] and individual isogeny steps in the graph have a low computational cost, making these graphs desirable for cryptographic applications. The first supersingular isogeny-based cryptographic protocol [2] was developed quite recently, and mathematicians have only been studying these graphs with cryptographic applications in mind for the past two decades.

In this work, we consider supersingular elliptic curve L-isogeny graphs $\mathcal{G}(p, L)$ over $\overline{\mathbb{F}}_p$ where $L = \{\ell_1, \ldots, \ell_r\}$ is a collection of allowed prime isogeny degrees. This work is motivated by [3] who studied such graphs over $\mathbb{F}_p$. In isogeny-based cryptographic protocols where multiple degrees are allowed, an L-isogeny graph is the structure underlying the security of the protocol. These graphs are $(\ell_1 + \cdots + \ell_r + r)$-regular, but no longer satisfy the Ramanujan property. Having more edges means these graphs have more *cycles*. Cycles in ℓ-isogeny graphs have been well-studied [4, 5, 6, 7]. A graph cycle starting at a vertex corresponding to the elliptic curve E gives an endomorphism of E. Computing the endomorphism ring of a supersingular elliptic curve is an active area of research and a question which underlies almost all of isogeny-based cryptography.

1.1. Our contributions

Our results concern two graph-theoretic structures in L-isogeny graphs: isogeny cycles (as defined in Definition 8) and cuts.

Cycles. We provide explicit counts of cycles in L-isogeny graphs in Section 4: one count using the theory of Brandt matrices in Section 4.1 and the other count using the theory of embeddings in quaternion algebras in Section 4.2. These two sections count slightly different classes of cycles.

Cuts. We study edge expansion in the L-isogeny graph using methodology inspired by Fiedler cuts in Section 5. We provide a comparison with alternative algorithms for finding cuts with good edge expansion, and show heuristically that the greedy neighbour algorithm out-perform Fiedler's algorithm.

We provide code in SageMath [8] at: `https://github.com/jtcc2/cycles-and-cuts`. This includes implementations of our counting arguments, clustering algorithms, figures, and examples.

1.2. Related work

While this work was in preparation, the authors were made aware of [9], which also studies cycles in L-isogeny graphs. The authors [9] take a computational approach with the goal of computing endomorphism rings using cycles found in L-isogeny graphs, applying an approach similar to [5]. Their study of cycles restricts to a very special class, aligning with the work of [5]. Our work is a complementary extension: we consider the general class of isogeny cycles and provide explicit cycle counts.

1.3. Acknowledgements

This work was completed as part of the Banff International Research Station (BIRS) Isogeny Graphs in Cryptography Workshop in 2023. This workshop was held in-person at BIRS in Banff, Alberta, Canada as well as in-person in Bristol, UK. We would like to express our gratitude to the organisers, Victoria de Quehen, Chloe Martindale, and Christophe Petit. This particular working

group, led by Sarah Arpin, Giulio Codogni, and Travis Morrison, had participants in both Banff and Bristol. We received a lot of interest in this topic, and broke into smaller project groups. The authors are indebted especially to contributions by Giulio Codogni. The first and fifth authors are supported in part by faculty fellowships from the Commonwealth Cyber Initiative (CCI), an investment in the advancement of cyber R&D, innovation, and workforce development. For more information about CCI, visit www.cyberinitiative.org. The authors further express their gratitude to their fellow group members for extremely helpful conversations, including: Kirsten Eisenträger, Jun Bo Lau, William Mahaney, Travis Morrison, Eli Orvis, James Rickards, Maria Sabitova, Gabrielle Scullard, and Lukas Zobernig. We extend our thanks to the anonymous reviewers for their careful reading of this manuscript and many suggestions to improve this work.

2. Preliminaries

2.1. Supersingular elliptic curves

We recall some standard facts about supersingular elliptic curves. For more detail, see [10, 11]. An elliptic curve over a finite field $\mathbb{F}_q$ is supersingular if and only if the geometric endomorphism ring of E, $\mathrm{End}(E) := \mathrm{End}_{\overline{\mathbb{F}}_q}(E)$, is a maximal order in a quaternion algebra. The classical Deuring correspondence [12] gives a categorical equivalence through which we can interpret isogenies of elliptic curves as left ideals of maximal orders in a quaternion algebra. To a given left ideal I of $\mathrm{End}(E)$, we associate an isogeny φ_I with kernel given by the scheme-theoretic intersection

$$\ker \varphi_I = \bigcap_{\alpha \in I} \ker \alpha.$$

When I is a principal ideal, φ_I is an endomorphism. If two left ideals I, J of $\mathrm{End}(E)$ are in the same ideal class, the codomains of φ_I and φ_J are isomorphic. We almost have a group action: however, the set of left ideal classes of $\mathrm{End}(E)$ is not a group.

If E is a supersingular elliptic curve defined over a prime field $\mathbb{F}_p$, the $\mathbb{F}_p$-part of its endomorphism ring, $\mathrm{End}_{\mathbb{F}_p}(E)$, is an imaginary quadratic order. In this case, we have a true action of the class group of ideals: the invertible ideals of the imaginary quadratic order $\mathrm{End}_{\mathbb{F}_p}(E)$ act freely on the set of supersingular elliptic curves with endomorphism ring isomorphic to $\mathrm{End}_{\mathbb{F}_p}(E)$.

The $\overline{\mathbb{F}}_p$-automorphism groups of supersingular elliptic curves are generally $\{[\pm 1]\}$. For precisely two $\overline{\mathbb{F}}_p$-isomorphism classes of elliptic curves are these automorphism groups larger: $j = 0, 1728$ [13]. Curves E with $j(E) = 0$ or 1728 are said to have *extra automorphisms*.

2.2. ℓ-isogeny graphs

Definition 1 (ℓ-isogeny graph)**.** *For two primes p and ℓ, the ℓ-isogeny graph, denoted $\mathcal{G}(p, \ell)$ is the graph with vertices given by isomorphism classes of su-*

persingular elliptic curves over $\overline{\mathbb{F}}_p$, with edges given by degree-ℓ isogenies, up to post-composing with an automorphism.

A priori, an isogeny has a unique dual so we should be able to identify the edge corresponding to an isogeny with the edge corresponding to its dual and create an undirected graph. However, the definition of the edges in Definition 1 opens the door to some issues for vertices with automorphism groups larger than $[\pm 1]$: if $\varphi_1 : E_{1728} \to E$ and $\varphi_2 : E_{1728} \to E$ have the same domain with j-invariant 1728 and the same codomain, it may happen that $\widehat{\varphi}_1 = [i] \circ \widehat{\varphi}_2$. In this case, there are two edges from the vertex E_{1728} to the vertex E, but only one edge from E to E_{1728}. If we wish to create an undirected graph, we sacrifice the regularity of the graph. The issue can only arise adjacent to vertices with $j = 0$ or $j = 1728$, so given a fixed ℓ, this problem can be avoided completely by choosing p according to certain congruence conditions.

If we associate to each edge an isogeny with explicit rational equations, then a cycle in the ℓ-isogeny graph can be composed into an endomorphism. However, if we only specify kernels and one of the vertices in the cycle is an isomorphism class with extra automorphisms, there is no canonical way of identifying this cycle with an endomorphism: composing with an extra automorphism may completely change the discriminant of the endomorphism obtained (where the discriminant of an endomorphism is defined as in [5, Section 2C], via the Deuring correspondence [12]). There are finitely many extra automorphisms, so the set of possibilities is finite - the issue is identifying them. To handle this issue, the notion of *arbitrary assignment* was developed in [14].

Definition 2 (Arbitrary assignment, [14, Definition 3.3])**.** *Given an ℓ-isogeny graph $\mathcal{G}(p, \ell)$, an arbitrary assignment is a choice (up to sign) of equivalence class representative $\pm\varphi$ for every edge $[\varphi]$.*

For edges yielding isogenies where the codomain j-invariants are not equal to 0 or 1728, an arbitrary assignment is automatic and does not involve any choice. When $p \equiv 1 \pmod{12}$, such an assignment is not necessary as $[-1]$ commutes with every isogeny, so this can always be a post-composition, which is already an equivalence we place on edges.

2.3. L-isogeny graphs

Notation. Let $p, \ell_1, \ldots, \ell_r$ be distinct primes. Let $L := \{\ell_1, \ldots, \ell_r\}$. For convenience and without loss of generality, suppose $\ell_1 < \ell_2 < \cdots < \ell_r$. In what follows, $e_i \in \mathbb{Z}_{\geq 0}$ for all $i = 1, \ldots, r$.

In [15], the author studies two isogeny graphs $\mathcal{G}(p, \ell_1), \mathcal{G}(p, \ell_2)$ simultaneously. We now consider a generalisation to an arbitrary number of isogeny graphs $\mathcal{G}(p, \ell_i)$, for $i = 1, \ldots, r$.

Definition 3 (L-isogeny graph)**.** *The L-isogeny graph, denoted $\mathcal{G}(p, L)$ is the graph with vertices given by isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$, with edges given by degree ℓ_i-isogenies (up to post-composing with an automorphism) for $i = 1, \ldots, r$.*

As defined, L-isogeny graphs are directed graphs. As in the ℓ-isogeny case described in Section 2.2, they are undirected for appropriate congruence conditions on p. See Figure 4 for the $\{2,3\}$-isogeny graph of supersingular elliptic curves over $\overline{\mathbb{F}}_{61}$.

2.4. Brandt matrices

Let $\mathcal{B}_p$ denote the (unique up to isomorphism) quaternion algebra over $\mathbb{Q}$ ramified at p and ∞. Two left ideals I, J of an order in $\mathcal{B}_p$ are *equivalent* if there exists $\beta \in \mathcal{B}_p^\times$ such that $J = I\beta$. The class set $cl(\mathfrak{O}) = \{I_1, I_2, \ldots, I_n\}$, for a maximal order $\mathfrak{O}$ in $\mathcal{B}_p$, is the set of representatives of equivalence classes of left $\mathfrak{O}$-ideals. We let $I_1 = \mathfrak{O}$ by convention. The class group is finite and its cardinality n is called the *class number* of $\mathcal{B}_p$ (it is the same for every maximal order $\mathfrak{O}$ in $\mathcal{B}_p$). Let $O_R(I)$ denote the right order of the ideal I. The set $\Gamma_i := O_R(I_i)^\times/\mathbb{Z}^\times$ is finite, as it is a discrete subgroup of the compact Lie group $(\mathcal{B}_p \otimes \mathbb{R})^\times/\mathbb{R}^\times \cong \mathrm{SO}_3(\mathbb{R})$. Let w_i be its cardinality. See [11].

Following [16], we now introduce theta series and Brandt matrices. For an ideal I_i, define the inverse ideal I_i^{-1} as $I_i^{-1} := \{\beta \in \mathcal{B}_p : I_i\beta I_i \subseteq I_i\}$, and let $M_{ij} := I_j^{-1}I_i = \{\sum_{k=1}^N a_k b_k : N \in \mathbb{Z}_{\geq 1}, a_k \in I_j^{-1}, b_k \in I_i\}$. We define the reduced norm $\mathrm{Nm}(M_{ij})$ of M_{ij} to be the unique positive rational number such that all quotients $\mathrm{Nm}(a)/\mathrm{Nm}(M_{ij})$, for $a \in M_{ij}$, form a coprime set of integers. We now define the following *theta series* θ_{ij}:

$$\theta_{ij}(\tau) := \frac{1}{2w_j} \sum_{a \in M_{ij}} q^{\frac{\mathrm{Nm}(a)}{\mathrm{Nm}(M_{ij})}} = \sum_{m \geq 0} B_{ij}(m) q^m,$$

where $q := e^{2\pi i \tau}$ and $\tau \in \{z \in \mathbb{C} : \mathrm{Im}(z) > 0\}$. Finally, we define the *Brandt matrix* $B(m)$ to be $B(m) := \left[B_{ij}(m)\right]_{1 \leq i,j \leq n}$, for $m \in \mathbb{Z}_{\geq 1}$. Note that $B(1)$ is the identity matrix. We have the following properties for Brandt matrices.

Proposition 4 ([16, Proposition 2.7])**.**

(i) If $m \geq 1$, then $B_{ij}(m) \in \mathbb{Z}_{\geq 0}$, and for all $1 \leq i \leq n$,

$$\sum_{j=1}^{n} B_{ij}(m) = \sum_{\substack{d \mid m \\ \gcd(d,p)=1}} d.$$

(ii) If m and m' are coprime, then $B(mm') = B(m)B(m')$.

(iii) If $\ell \neq p$ is a prime, then

$$B(\ell^k) = B(\ell^{k-1})B(\ell) - \ell B(\ell^{k-2}), \tag{1}$$

for all $k \geq 2$.

(iv) $w_j B_{ij}(m) = w_i B_{ji}(m)$ for all m and for all $1 \leq i, j \leq n$.

The main reason for considering Brandt matrices in our study of isogeny graphs is the fact that $B_{ij}(m)$ is the number of equivalence classes of isogenies of degree m between the elliptic curves corresponding to I_i and I_j (cf. [16, Proposition 2.3]). Therefore, the Brandt matrix $B(\ell)$ is the adjacency matrix of the supersingular isogeny graph $\mathcal{G}(p, \ell)$. Hence, in order to study cycles of length r in the graph $\mathcal{G}(p, \ell)$, one can alternatively look at the entries $B(\ell^r)$ to count those closed walks based at I_i, or at $\mathrm{Tr}(B(\ell^r))$ to count them all (counting all starting points). See [15, 17] for more details on this approach, in the case of two primes.

Now, computing the entries of $B(\ell^r)$ can be done using the recursion relation of Equation (1), but this can be very inefficient for larger powers of ℓ or large primes p. However, there is a way of computing certain traces of Brandt matrices, by re-expressing them in terms of modified Hurwitz class numbers, which we define below.

Given an order $\mathcal{O}$ in an imaginary quadratic field, let d be its (negative) discriminant, $h_{\mathcal{O}}$ the size of the class group $cl(\mathcal{O})$, and $u_{\mathcal{O}} := |\mathcal{O}^\times/\mathbb{Z}^\times|$. To emphasize the discriminant d of an order, we will denote the order $\mathcal{O}_d$. Fix $D > 0$ for which there exists an order $\mathcal{O}_{-D}$ of discriminant $-D$. The *Hurwitz Class Number* $H(D)$ is

$$H(D) = \sum_{\substack{d \in \mathbb{Z}, \mathfrak{f} > 0 \\ d \cdot \mathfrak{f}^2 = -D}} \frac{h_{\mathcal{O}_d}}{u_{\mathcal{O}_d}}.$$

and the *modified Hurwitz Class Number* $H_p(D)$, for a prime p, is

$$H_p(D) := \begin{cases} 0 & \text{if } p \text{ splits in } \mathcal{O}_{-D}; \\ H(D) & \text{if } p \text{ is inert in } \mathcal{O}_{-D}; \\ \frac{1}{2}H(D) & \text{if } p \text{ is ramified in } \mathcal{O}_{-D} \\ & \text{but does not divide the conductor of } \mathcal{O}_{-D}; \\ H(\frac{D}{p^2}) & \text{if } p \text{ divides the conductor of } \mathcal{O}_{-D} \end{cases} \tag{2}$$

For $D = 0$, we set $H(0) = -1/12$ and $H_p(0) := \frac{p-1}{24}$. Note that $H_p(D) \leq H(D)$. We conclude this section with the following two results on Hurwitz class numbers. The first connects Hurwitz class numbers to the trace of Brandt matrices. The second gives us a way to compute certain sums of Hurwitz class numbers.

Theorem 5 ([16, Prop. 1.9])**.** *For all integers $m \geq 0$,*

$$\mathrm{Tr}(B(m)) = \sum_{\substack{s \in \mathbb{Z} \\ s^2 \leq 4m}} H_p(4m - s^2).$$

Theorem 6 ([18, Section 7])**.** *For all integers $m \geq 1$,*

$$\sum_{\substack{s \in \mathbb{Z} \\ s^2 \leq 4m}} H(4m - s^2) = 2 \sum_{\substack{d \mid m \\ d > 0}} d - \sum_{\substack{d \mid m \\ d > 0}} \min\{d, m/d\}.$$

2.5. Matrix representations of the graph

Let $G = (V, E)$ be a graph. Fix an ordering of the vertices $V = (v_1, \ldots, v_n)$. The *adjacency matrix* of G with respect to the ordering $(v_1, \ldots, v_n)$ is the $n \times n$ matrix A whose entries a_{ij} are the number of edges from v_i to v_j. If the graph is undirected, the adjacency matrix is symmetric. If the graph is directed, this may or may not be the case.

For an undirected graph, the degree of a vertex is the number of edges incident to that vertex. If the graph is directed, we define the in- and out-degree of a vertex to be the number of in- and out-edges, respectively. Given an ordering of the vertices $V = (v_1, \ldots, v_n)$, the (in- or out-)degree matrix is the diagonal matrix D whose nonzero entries d_{ii} are the (in- or out-)degrees of the vertices v_i.

The Laplacian of an undirected graph is the matrix $L := D - A$. If G is directed, we can define the in- or out-Laplacian where D is the in- or out-degree matrix, respectively.

2.6. Edge Expansion

A *cut* of a graph $G = (V, E)$ is an ordered partition of the vertices $V = V_1 \sqcup V_2$ with $V_1, V_2 \neq \varnothing$. The sets V_1 and V_2 are called the *sides* of the cut. The directed edges $e = (u, v) \in E$ which have $u \in V_1$ and $v \in V_2$ (or vice versa), and undirected edges with one endpoint in each of V_1 and V_2, are said to *cross* the cut, and the set of such edges is called the *cut set* of the cut, denoted $E(V_1, V_2)$. The cuts which are usually of most interest in graph theory are those which are relatively balanced (i.e., in which $|V_1| \approx |V_2|$) and relatively *sparse* (i.e., they have relatively few edges which cross the cut). These desirable properties are incorporated in the *edge expansion*, defined in Definition 7.

Definition 7 (Edge expansion)**.** *Let $G = (V, E)$ be a graph. The* edge expansion *of a cut $V = C \sqcup (V \setminus C)$ is*

$$h_G(C) := \frac{|E(C, V \setminus C)|}{\min\{\mathrm{Vol}(C), \mathrm{Vol}(V \setminus C)\}},$$

where $\mathrm{Vol}(C)$ is the sum of the vertex degrees in C, which means for d-regular graphs, $\mathrm{Vol}(C) = d \cdot |C|$. The edge expansion of G is $h(G) := \min_{\varnothing \subsetneq T \subsetneq V} h_G(T)$.

3. Cycles in the L-isogeny graph

Cycles in isogeny graphs are of interest for two main reasons:

1. They correspond to collisions of walks in the isogeny graph [2];
2. They correspond to endomorphisms of elliptic curves and can be used to compute endomorphism rings [4, 5, 6, 9].

In this section and the following section, we restrict ourselves to the case $p \equiv 1 \pmod{12}$, which allows us to treat $\mathcal{G}(p, \ell)$ as an undirected graph by identifying each isogeny with its dual, thus avoiding counting issues arising from extra automorphisms. Moreover, we assume that L never contains p.

3.1. Setup

Definition 8 (Isogeny cycle)**.** *An $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycle*[1] *is a closed walk in $\mathcal{G}(p, \{\ell_1, \ldots, \ell_r\})$ where the set of edges in the walk consists of e_i degree-ℓ_i isogenies for $i = 1, \ldots, r$, with $e_i > 0$. We denote the set of all $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycles in $\mathcal{G}(p, \{\ell_1, \ldots, \ell_r\})$ by $\mathrm{Cycles}_p(\ell_1^{e_1}, \ldots, \ell_r^{e_r})$.*

Remark 9. *It is possible to discuss $\{\ell_i^{e_i}\}_{i=1}^r$-isogeny cycles in an L-isogeny graph $\mathcal{G}(p, L)$ where $\{\ell_1, \ldots, \ell_r\} \subsetneq L$, but since the cycle does not contain edges of degree ℓ for any $\ell \in L \setminus \{\ell_1, \ldots, \ell_r\}$, we may simply consider this isogeny cycle as belonging to the graph $\mathcal{G}(p, \{\ell_1, \ldots, \ell_r\})$. Going forward, when we discuss $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycles, we automatically define $L := \{\ell_1, \ldots, \ell_r\}$.*

Given a walk (resp. closed walk) in $\mathcal{G}(p, L)$ as a sequence of edges, we may successively compose compatible isogeny representatives corresponding to these edges to obtain an isogeny (resp. endomorphism) of elliptic curves. Each vertex represents an isomorphism class of elliptic curves, so the representative of each edge source is successively chosen in a compatible way.[2] The reverse process involves decomposing an isogeny into a sequence of prime degree isogenies.

Definition 10. *Let $\varphi : E \to E'$ be a separable isogeny of degree divisible only by primes contained in L. A (prime) decomposition of φ is a factorisation of φ into a sequence of (prime degree) isogenies*

$$\varphi = \varphi_k \circ \varphi_{k-1} \circ \cdots \circ \varphi_2 \circ \varphi_1 =: (\varphi_i)_{i=1}^k.$$

Define $\mathcal{D}(\varphi)$ to be the set of prime decompositions of φ.

A prime decomposition yields a graphical representation of an isogeny as a walk in $\mathcal{G}(p, L)$, since each φ_i corresponds to an edge. In the context of the discussion above, we treat prime decompositions as equivalent if they yield the same graph-theoretic walk. Since we are assuming $p \equiv 1 \mod 12$, every elliptic curve has $\mathrm{Aut}(E) = \{\pm 1\}$, so two prime decompositions are equivalent when they differ by replacing φ_i with $-\varphi_i$ for any subset of indices i. Even up to this equivalence, a prime decomposition need not be unique.

Lemma 11. *Let $\varphi : E \to E'$ be an isogeny with kernel G such that $\deg \varphi$ is divisible only by primes contained in L. Let $(P, \lhd)$ be the poset of subgroups of G, ordered by inclusion. There is a 1-1 correspondence between the prime decompositions of φ (elements of $\mathcal{D}(\varphi)$) and the maximal chains in $(P, \lhd)$.*

Proof. Given a prime decomposition $\varphi = (\varphi_i)_{i=1}^k$, define the sequence of groups

$$G_0 = \{O_E\}, G_i = \ker(\varphi_i \circ \varphi_{i-1} \circ \cdots \circ \varphi_1).$$

[1] We note that this is different from a cycle in the typical graph-theoretic sense, which is a closed *path*—i.e., a closed walk with no repeated vertices.

[2] This is possible as no additional automorphisms exist, so pre- and post-composition of an isogeny with automorphisms will always correspond to the same edge.

This sequence forms a chain in P as each φ_i is a group homomorphism. Suppose there exists $H \in P$ such that $G_{i-1} \lhd H \lhd G_i$. Comparing indices, we have $\deg \varphi_i = |G_i/G_{i-1}| = |G_i/H||H/G_{i-1}|$, and as $\deg \varphi_i$ is prime, either $H = G_i$ or $H = G_{i-1}$. As $G_k = G$, we conclude that this chain is maximal.

For the reverse direction, let $\{O_E\} = G_0 \lhd G_1 \lhd \cdots \lhd G_k = G$ be a maximal chain in P. For each $1 \leq i \leq k$, there exists a unique curve E_i (up to isomorphism) and isogeny $\psi_i : E \to E_i$ (up to sign) with $\ker \psi_i = G_i$ ([13, Chapter III, Proposition 4.11]). Define $\varphi_1 = \psi_1$, and note that $\psi_k = \varphi$. For $i > 1$, define φ_i to be the unique isogeny satisfying $\psi_i = \varphi_i \circ \psi_{i-1}$ ([13, Chapter III, Corollary 4.11]). This procedure yields a decomposition $\varphi = \psi_k = \varphi_k \circ \varphi_{k-1} \circ \cdots \circ \varphi_1$ where $\deg \varphi_i = |G_i/G_{i-1}|$. Each of these indices is prime as the G_i's form a maximal chain. Finally, these correspondences are inverse of each other up to equivalence of decompositions. □

As isogenies have finite kernels, $\mathcal{D}(\varphi)$ is always a finite set. We can compute $|D(\varphi)|$ from the prime power decomposition of $\deg \varphi$ together with knowledge of the largest integer n such that $\varphi(E[n]) = O_{E'}$. As a special case, suppose the isogeny φ is primitive ([4, Definition 4.1]; equivalently, has cyclic kernel); then, we have $n = 1$ and writing $\deg \varphi = \prod_{i=1}^{r} \ell_i^{e_i}$ we have

$$|\mathcal{D}(\varphi)| = \binom{\sum_{i=1}^{r} e_i}{e_1, e_2, \ldots, e_r} = \frac{(\sum_{i=1}^{r} e_i)!}{\prod_{i=1}^{r} e_i!} \tag{3}$$

using a simple combinatorial argument [19, Section 5.3, Theorem 1]. In particular, if the number of distinct primes dividing $\deg \varphi$ is greater than 1, then $|D(\varphi)| > 1$: such isogenies do not have a unique prime decomposition.

3.2. Backtracking

A walk in the ℓ-isogeny graph composes to an isogeny with non-cyclic kernel if and only if the walk contains backtracking [4, Definition 4.3]. In particular, there is no canonical choice of walk in the ℓ-isogeny graph corresponding to the multiplication-by-ℓ isogeny of a particular starting vertex.

Definition 12 (Backtracking)**.** *An isogeny cycle given by the composition* $\varphi_k \circ \varphi_{k-1} \circ \cdots \circ \varphi_1$ *has* backtracking *if* $\widehat{\varphi}_i \in \mathrm{Aut}(E_i)\varphi_{i+1}$ *for some* $i \in \{1, 2, \ldots, k-1\}$ *where* $\varphi_i : E_i \to E_{i+1}$ *and* $\varphi_{i+1} : E_{i+1} \to E_{i+2}$.

In $\mathcal{G}(p, \ell)$, a cycle has backtracking if and only if the endomorphism obtained by composing (compatible isogeny representatives of) edges in the cycle has non-cyclic kernel [4, Proposition 4.5]. With L-isogeny graphs, the kernel of an isogeny which is the composition of a non-backtracking walk can still be non-cyclic, as demonstrated in Example 13.

Example 13 (A Non-Cyclic Kernel Arising from an Isogeny Diamond)**.** Let E_1 be an elliptic curve, and suppose $\varphi\colon E_1 \to E_2, \psi\colon E_1 \to E_3$ are two isogenies of distinct prime degrees, each with domain E_1. Define isogenies φ' from E_3 and ψ' from E_2 with $\ker \varphi' = \psi(\ker \varphi)$ and $\ker \psi' = \varphi(\ker \psi)$. The diagram in

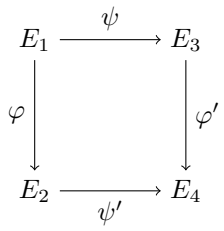

Figure 1: An isogeny diamond.

Figure 1 commutes (up to post-composition by an automorphism) [13, Corollary III.4.11].

The isogeny diamond gives a $\{(\deg\varphi)^2, (\deg\psi)^2\}$-isogeny cycle in the $\{\deg\psi, \deg\varphi\}$-isogeny graph, namely $\widehat{\varphi} \circ \widehat{\psi'} \circ \varphi' \circ \psi =: \alpha \in \mathrm{End}(E_1)$. We see that α and $[\deg\psi] \circ [\deg\varphi]$ are equivalent by computing $\alpha(E[\deg\psi]) = \alpha(E[\deg\varphi]) = O_E$.

Indeed, the isogeny cycle $\widehat{\varphi} \circ \widehat{\psi'} \circ \varphi' \circ \psi \in \mathrm{End}(E_1)$ from Example 13 does not have backtracking. By refactoring the cycle as $[\deg\psi] \circ [\deg\varphi]$ so that isogenies of the same degree are grouped, we see that it has non-cyclic kernel.

Lemma 14 (Isogeny swapping)**.** *Suppose we are given an isogeny as the composition $\varphi' \circ \psi$ of two isogenies of coprime degrees. Define an isogeny ψ' with $\ker\psi' := \varphi(\ker\psi)$ and an isogeny φ with $\ker\varphi := \widehat{\psi}(\ker\varphi')$. Then,*

$$\varphi' \circ \psi = \alpha \circ \psi' \circ \varphi,$$

where α is an automorphism of the codomain of φ'.

Proof. The result follows directly from the definitions of φ and ψ': the resulting diagram is an isogeny diamond (see Figure 1). □

3.3. Canonical Decomposition

To handle the issue of multiple cycles composing to the same endomorphism, we provide a canonical means of refactoring a given isogeny cycle: we refactor the isogeny using Lemma 14, grouping the isogenies by increasing prime degree ($\ell_1 < \ell_2 < \cdots < \ell_r$, as in Notation 2.3). Consecutive isogenies of the same prime degree cannot be reordered; however they may compose to a scalar multiplication map, which may be reordered arbitrarily, so we obtain a canonical decomposition which pushes scalar multiplication maps together.

Definition 15 (Canonical decomposition)**.** *Let $\ell_1 < \ell_2 < \cdots < \ell_r$ be distinct primes. Given a fixed $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycle, let α denote the endomorphism given by the composition of the isogenies in this isogeny cycle.*

There is a largest positive integer n such that $\alpha(E[n]) = O_E$. Write $\frac{\deg(\alpha)}{n^2} = \ell_{i_1}^{f_1} \cdots \ell_{i_m}^{f_m}$ where $i_1, \ldots, i_m$ is a subsequence of $1, \ldots, r$ such that each f_i is non-zero.

The canonical decomposition *of α is the decomposition:*

$$\alpha = [n] \circ (\varphi_{m,f_m} \circ \cdots \circ \varphi_{m,1}) \circ \cdots \circ (\varphi_{2,f_2} \circ \cdots \circ \varphi_{2,1}) \circ (\varphi_{1,f_1} \circ \cdots \circ \varphi_{1,1}), \quad (4)$$

where $\deg \varphi_{j,k} = \ell_{i_j}$ *for all* k.

We define the L-isogeny graph walk associated to the canonical decomposition *of α to be the walk given by the list of edges:*

$$\varphi_{m,f_m}, \cdots, \varphi_{m,1}, \cdots, \varphi_{2,f_2}, \cdots, \varphi_{2,1}, \varphi_{1,f_1}, \cdots, \varphi_{1,1}.$$

Remark 16. *We do note associate any isogeny walk in the L-isogeny graph to the $[n]$ that appears in Equation (4), since there are many walks in the graph whose associated isogenies compose to $[n]$ and there is no canonical way to choose one. For example, there are three walks in the 2-isogeny graph corresponding to $[2]$ at a particular starting vertex, but $[2]$ itself is not an edge in the graph.*

However, the endomorphism $(\varphi_{m,f_m} \circ \cdots \circ \varphi_{m,1}) \circ \cdots \circ (\varphi_{2,f_2} \circ \cdots \circ \varphi_{2,1}) \circ (\varphi_{1,f_1} \circ \cdots \circ \varphi_{1,1})$ corresponds to a uniquely-determined $\{\ell_{i_j}^{f_j}\}$-isogeny cycle.

This motivates the following algorithm for finding the canonical decomposition of an isogeny cycle.

Algorithm 1: Finding a canonical decomposition of an endomorphism

Input: An endomorphism α of a supersingular elliptic curve E given as the composition of a list of (known) prime-degree isogenies: $(\varphi_1, \ldots, \varphi_N)$

Output: An ordered list $(\psi_1, \ldots, \psi_M)$ of isogenies of non-decreasing prime degree, and an integer n such that the composition $[n]_E \circ \psi_M \circ \cdots \circ \psi_1$ is equal to α up to sign.

```
Apply BubbleSort[3] to the list Φ := (φ_1, ..., φ_N) to sort by increasing
  degree, using the isogeny swapping procedure in Lemma 14
n_0 ← 1
i ← 1
While i ≤ len(Φ) do
    Let E_i be the domain of Φ[i]
    if Φ[i+1] ∘ Φ[i] ∈ Aut(E_i)[deg Φ[i]] then
        n_0 ← n_0 · deg Φ[i]
        Remove Φ[i], Φ[i+1] from Φ
        i ← max(1, i − 1)
    else
        i ← i + 1
return Φ, n_0.
```

[3]Lemma 14 allows us to swap *adjacent* isogenies of different degrees in an isogeny cycle. Iteratively applying this idea, we can sort the isogenies in increasing degree by applying any sorting algorithm that only makes adjacent swaps. BubbleSort is one such algorithm [20].

Proposition 17. *Algorithm 1 computes a canonical decomposition of an isogeny cycle.*

Proof. Let α be the endomorphism obtained from an $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycle represented as a prime decomposition $(\varphi_1, \ldots, \varphi_N)$. Let $(\tilde{\Phi}, \tilde{n})$ be the output of Algorithm 1 on input α, where $\tilde{\Phi} = (\psi_1, \ldots \psi_M)$. Let $\beta = \psi_M \circ \cdots \circ \psi_1$, and let n be the largest integer such that $E[n] \subseteq \ker \alpha$. The list Φ obtained after line 1 gives a prime decomposition of α, with the isogenies ordered by increasing degree. This property is preserved by the loop, as well as the invariant $\alpha = [n_0]_E \circ \Phi[\mathsf{len}(\Phi)] \circ \cdots \circ \Phi[1]$. Hence the identity $\alpha = [\tilde{n}] \circ \beta$ is satisfied. The internal loop variable n_0 is non-decreasing, hence we have $\tilde{n} \leq n$. To show equality, note that if β factors through the scalar map $[\ell]$ for some $\ell \in L$, then $\psi_b \circ \cdots \circ \psi_a$ must factor through $[\ell]$ where $\psi_a, \ldots \psi_b$ are the elements of $\tilde{\Phi}$ of degree ℓ. In turn, an $\ell^\star$-isogeny factors through $[\ell]$ if and only if it is backtracking, i.e., there exists an index c with $a \leq c < b$ such that $\psi_{c+1} = \pm\widehat{\psi_c}$, because $p \equiv 1 \pmod{12}$. This condition is prohibited in the output by the conditional on line 6.

To see that the canonical decomposition is unique, note that the degree determines the ordering of the prime-power parts of the walk. The decomposition of a prime-power degree isogeny is unique up to composition with automorphisms at the vertices. Thus, if the vertices in the walk have only $[\pm 1]$ automorphisms, the canonical decomposition is unique up to $[\pm 1]$. □

Corollary 18. *The L-isogeny graph walk corresponding to the canonical decomposition of an isogeny cycle contains no backtracking and the corresponding endomorphism has cyclic kernel.*

Proof. The scalar portion of the canonical decomposition is not part of the walk of the graph, see Remark 16. Since the canonical decomposition of an isogeny cycle groups together isogenies of the same prime degree, it contains backtracking if and only if it contains a walk of degree-ℓ_i isogenies which contains backtracking [4, Prop. 4.5]. □

Every isogeny can be written as $[n] \circ \phi$, where ϕ has cyclic kernel. If $[n] = [1]$, we will call the isogeny cycle *principal*.

Definition 19 (Principal isogeny cycles)**.** *An $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycle is principal if its canonical decomposition has scalar factor $[1]$. Let $\mathrm{Cycles}_p(\ell_1^{e_1}, \ldots, \ell_r^{e_r})^{\mathrm{prin}}$ denote the subset of principal cycles.*

The endomorphism obtained from an isogeny cycle in $\mathrm{Cycles}_p(\ell_1^{e_1}, \ldots, \ell_r^{e_r})^{\mathrm{prin}}$ has cyclic kernel, by Corollary 18.

Example 20 (Isogeny cycle)**.** Consider the abstract $\{2^2, 3^3\}$-isogeny cycle beginning at vertex j_1 depicted in Figure 2. Let $\theta \in \mathrm{End}(E_{j_1})$ denote the endomorphism obtained by composing (compatible isogeny representatives of) the edges of this cycle. The endomorphism θ can be written as a composition of prime degree isogenies in a total of $\binom{5}{2} = 10$ ways, each with two degree-2 isogenies and

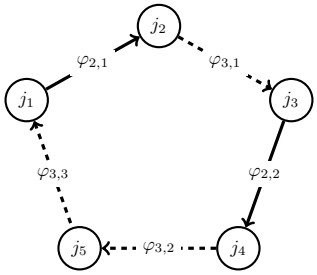

Figure 2: A $\{2^2, 3^3\}$-isogeny cycle in the $\{2, 3\}$-isogeny graph. The (solid line) isogenies $\varphi_{2,1}$ and $\varphi_{2,2}$ are degree-2 and the (dashed line) isogenies $\varphi_{3,1}, \varphi_{3,2}, \varphi_{3,3}$ are degree-3. This $\{2^2, 3^3\}$-isogeny cycle can be specified by the tuple $(\varphi_{3,3}, \varphi_{3,2}, \varphi_{2,2}, \varphi_{3,1}, \varphi_{2,1})$, starting at the vertex j_1.

three degree-3 isogenies. The canonical decomposition of this endomorphism will be of the form $\psi_{3,3} \circ \psi_{3,2} \circ \psi_{3,1} \circ \psi_{2,2} \circ \psi_{2,1}$, where the ψ_{i_j} isogenies are chosen so that the kernel of the composition remains unchanged. In particular,

$$\ker \psi_{2,2} = \ker(\varphi_{3,3} \circ \varphi_{3,2} \circ \varphi_{2,2} \circ \varphi_{3,1}) \cap E(j_2)[2],$$

where $E(j_2)$ is the codomain of $\varphi_{2,1}$.

Example 21 (Isogeny cycle which is not principal)**.** Here, we give an explicit example of an isogeny cycle in non-canonical form which does not contain backtracking, but whose canonical form does contain backtracking (it has a scalar factor [2]). See Figure 3.

Let $p = 2689$. We consider a non-backtracking $\{2^2, 5, 13\}$-isogeny cycle of a supersingular elliptic curve over $\overline{\mathbb{F}}_p$. The details of this computation can be found in the code file `example_nonprincipalcycle.ipynb`.

We have the following isogenies $\varphi_{i_{(\cdot)}}$, where $\deg \varphi_{i_{(\cdot)}} = i$:

1. $\varphi_{2,1} : E_0 \to E_1$,
2. $\varphi_5 : E_1 \to E_2$,
3. $\varphi_{2,2} : E_2 \to E_3$,
4. $\varphi_{13} : E_3 \to E_0$,

where $E_0 : y^2 = x^3 + 2236x + 1886$, $E_1 : y^2 = x^3 + 732x + 2243$, $E_2 : y^2 = x^3 + 750x + 791$, $E_3 : y^2 = x^3 + 1996x + 1015$.

This isogeny cycle is depicted in Figure 3a. This cycle is non-principal. Let $\eta := \varphi_{13} \circ \varphi_{2,2} \circ \varphi_5 \circ \varphi_{2,1}$. The endomorphism η contains the factor [2], since $\eta(E_0[2]) = \mathcal{O}_{E_0}$. However, the isogeny cycle does not contain backtracking. This is not a contradiction of Corollary 18, because the isogeny cycle has not been presented in canonical form.

To refactor η into canonical form, we swap φ_5 and $\varphi_{2,2}$ in the order of compositions. Define a new $\psi_{2,2} : E_1 \to E_0$ with $\ker \psi_{2,2} = \ker \widehat{\varphi_{2,1}}$. In fact, $\psi_{2,2} = \pm\widehat{\varphi_{2,1}}$, and here we see the backtracking. Next, compute $\psi_{2,2}(\ker \varphi_5) = \widehat{\varphi_{2,1}}(\ker \varphi_5)$ to define $\psi_5 : E_0 \to E_3$. This completes the refactorisation of η as $\eta := [2] \circ \varphi_{13} \circ \psi_5$, as $\widehat{\varphi_{2,1}} \circ \varphi_{2,1} = [2]$. See Figure 3b. In this decomposition, we see

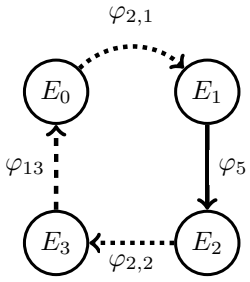

(a) The original $\{2^2, 5, 13\}$-isogeny cycle whose edges compose to the endomorphism η described in Example 21.

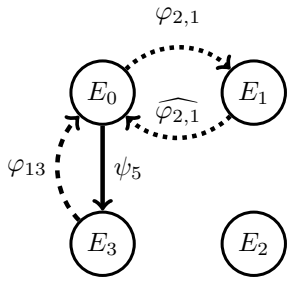

(b) The endomorphism η refactored in canonical decomposition. η is the composition of a principal $\{5, 13\}$-isogeny cycle and the scalar multiplication [2].

Figure 3: Isogeny cycle decompositions.

that η is the concatenation of two isogeny cycles: one principal $\{5, 13\}$-isogeny cycle and one non-principal $\{2^2\}$-isogeny cycle. In particular, once the prime-degree isogenies in the cycle have been sorted in order of degree, the resulting walk has backtracking, and the canonical decomposition has the scalar factor [2].

3.4. Equivalence classes of cycles

In undirected graphs, a graph-theoretic cycle beginning at a base point vertex V can be traversed in two possible directions from that base point V: the reversing of direction is the effect of taking the dual isogeny. For isogeny cycles whose degrees have more than one prime factor, the dual isogeny cycle will have a distinct canonical decomposition. To count isogeny cycles in Section 4, we will really be counting endomorphisms. When $|L| > 1$ and $p \equiv 1 \pmod{12}$, an endomorphism uniquely (up to sign) determines the canonical decomposition of an isogeny-cycle. We define an equivalence relation on the set of isogeny cycles by declaring two cycles to be equivalent if they have the same canonical decomposition.

4. Counting cycles

As in the previous section, we continue to assume that $p \equiv 1 \mod 12$. We now give two methods to count equivalence classes of principal isogeny cycles in L-isogeny graphs. First using Brandt matrices, and secondly using ideal relations in quadratic orders. Implementations of these counting methods, along with the examples given are in code file `cycle_counts.ipynb`.

4.1. Counting isogeny cycles using Brandt matrices

As briefly alluded to in Section 2.4, one way of counting cycles, in the isogeny graph $\mathcal{G}(p, \ell)$, is by relating them to diagonal elements on the Brandt matrix $B(\ell^k)$ for $k \geq 1$ associated to the isogeny graph in question. Indeed, $B(\ell)$

can be seen as the adjacency matrix of the graph $\mathcal{G}(p,\ell)$. One can also obtain a formula for the total number of cycles by considering the trace of certain Brandt matrices, which can be computed efficiently by re-expressing them as sums of modified Hurwitz class numbers (see theorems 5 and 6). This method is detailed thoroughly in [15]. One of the advantages of this method is that one can very efficiently obtain upper bounds as well as estimates for the total number of cycles in the graph.

For a prime p and a set $L := \{\ell_1, \ldots, \ell_r\}$, we are interested in computing the number of principal isogeny cycles of degree $\ell_1^{e_1} \cdots \ell_r^{e_r}$, at any base point E, in the graph $\mathcal{G}(p, L)$. For a fixed base point E, these cycles are denoted $\mathrm{Cycles}_p(\ell_1^{e_1}, \ldots, \ell_r^{e_r})_E^{\mathrm{prin}}/\sim$, where $\sim$ is equivalence up to refactoring. As an extension of the above, we would also like to compute the number of principal isogeny cycles, at any base point E, of length R in $\mathcal{G}(p, L)$, i.e., of degree $\ell_1^{e_1} \cdots \ell_r^{e_r}$ with $\sum_i e_i = R$. In the case $L = \{\ell_1, \ell_2\}$, this is very similar to [15, Section 4.2].

Let $\{E_1, \ldots, E_n\}$ be a set of representatives for the vertices of $\mathcal{G}(p, L)$. Then the number of principal isogeny cycles of degree $\ell_1^{e_1} \cdots \ell_r^{e_r}$, at base point E_i, in $\mathcal{G}(p, L)$ is equal to $B_{ii}(\ell_1^{e_1} \cdots \ell_r^{e_r})$ minus the number of non-principal cycles (endomorphisms) of E_i of degree $\ell_1^{e_1} \cdots \ell_r^{e_r}$.

For any non-principal cycle, the scalar part of its canonical decomposition must be a multiple of $[\ell_j]$ for some j. To remove cycles whose scalar part factors through $[\ell_j]$, we can subtract $B_{ii}(\ell_1^{e_1} \cdots \ell_j^{e_j-2} \cdots \ell_r^{e_r})$ from $B_{ii}(\ell_1^{e_1} \cdots \ell_r^{e_r})$, for all j. Of course, one must then account for the cycles whose scalar part factors through multiple $[\ell_j]$'s that have been subtracted more than once. This number can be computed using the inclusion-exclusion principle. In order to simplify our notation, we extend the field of definition of B to include rational numbers by setting $B(m) := 0$ for $m \notin \mathbb{Z}_{\geq 1}$. So, for example, if $L = \{\ell_1, \ell_2\}$, then we have that the number of endomorphisms of E_i of degree $\ell_1^{e_1}\ell_2^{e_2}$ involving scalar factors is given by $B_{ii}(\ell_1^{e_1-2}\ell_2^{e_2}) + B_{ii}(\ell_1^{e_1}\ell_2^{e_2-2}) - B_{ii}(\ell_1^{e_1-2}\ell_2^{e_2-2})$ (cf. [15, Lemma 4.2]). Hence, in the general case of $L = \{\ell_1, \ldots, \ell_r\}$, the number of endomorphisms of E_i of degree $\ell_1^{e_1} \cdots \ell_r^{e_r}$ involving scalar factors is

$$\sum_{\varnothing \subsetneq J \subseteq \{1,\ldots,r\}} (-1)^{|J|-1} B_{ii}\left(\frac{\prod_{i=1}^r \ell_i^{e_i}}{\prod_{j\in J} \ell_j^2}\right).$$

Summing over all vertices of the graph, we get the following formula.

Theorem 22. *The total number of principal isogeny cycles of degree $\ell_1^{e_1} \cdots \ell_r^{e_r}$ and any base point E_i, in $\mathcal{G}(p, L)$, is given by*

$$|\,\mathrm{Cycles}_p(\ell_1^{e_1}, \ldots, \ell_r^{e_r})^{\mathrm{prin}}/\sim\,| = \sum_{J\subseteq\{1,\ldots,r\}} (-1)^{|J|} \operatorname{Tr}\left(B\left(\frac{\prod_{i=1}^r \ell_i^{e_i}}{\prod_{j\in J} \ell_j^2}\right)\right)$$

Thanks to Theorem 5, we can make the above formula more tractable, by relating each trace to sums of modified Hurwitz class numbers. We will avoid substituting the formula from Theorem 5 into that of Theorem 22 to avoid

writing out a very cumbersome expression. However, if one wanted to compute the quantity $|\operatorname{Cycles}_p(\ell_1^{e_1},\dots,\ell_r^{e_r})^{\text{prin}}/\sim|$ precisely, this substitution could be carried out. As we will see later, if one was interested in finding bounds (or estimates) for the number of principal cycles, this method can be pushed further (by involving Theorem 6).

We will now consider a more general—and very natural notion of cycle count. Let $\mathcal{C}_{E_i}(L;R)$ denote the number of equivalence classes of principal isogeny cycles, at base point E_i, of length R in $\mathcal{G}(p,L)$. By *length R in $\mathcal{G}(p,L)$*, we naturally mean a cycle in the graph that is composed of R edges. In other words, we are not fixing certain exponents $\{e_1,\dots,e_r\}$ and only considering the isogeny cycles of degree $\ell_1^{e_1}\cdots\ell_r^{e_r}$, but rather any isogeny cycle of degree $\ell_1^{e_1}\cdots\ell_r^{e_r}$, for any $(e_1,\dots,e_r)$ satisfying $\sum_i e_i = R$. This can be seen as a generalisation of the *bi-route number* defined in [15]. We then have

$$\sum_{i=1}^{n}\mathcal{C}_{E_i}(L;R)=\sum_{i=1}^{n}\sum_{e_1+\cdots+e_r=R}|\operatorname{Cycles}_p(\ell_1^{e_1},\dots,\ell_r^{e_r})_{E_i}^{\text{prin}}/\sim|. \tag{5}$$

Let us now try to re-express the above quantity. Switching the order of summation in Equation (5) yields the following result.

Theorem 23. *The total number of equivalence classes of principal isogeny cycles, for any base point E_i, of length R in $\mathcal{G}(p,L)$ is given by*

$$\sum_{i=1}^{n}\mathcal{C}_{E_i}(L;R)=\sum_{i=0}^{r}(-1)^i\binom{r}{i}\sum_{\substack{e_1,e_2,\dots,e_r\in\mathbb{Z}_{\geq 0}\\ e_1+\cdots+e_r=R-2i}}\operatorname{Tr}(B(\ell_1^{e_1}\cdots\ell_r^{e_r})). \tag{6}$$

We observe in Equation (6), that for each $i\in\{0,...,r\}$ and each r-tuple of integers $(e_1,...,e_r)$ with $e_1+\cdots+e_r=R-2i$, the binomial coefficient $\binom{r}{i}$ counts the number of tuples $(f_1,...,f_r)$ there are with $f_1+...+f_r=R$ and such that $f_j=e_j+2$ for exactly i of the terms and $f_j=e_j$ for the rest. For each of these tuples, when counting cycles of degree $\ell_1^{e_1}\cdots\ell_r^{e_r}$, Theorem 22 gives us one contribution of the form $(-1)^i\operatorname{Tr}(B(\ell_1^{e_1}\cdots\ell_r^{e_r}))$.

Similarly to Theorem 22, one can compute the quantity in Theorem 23 precisely, using Theorem 5. If one is looking for a bound (or an approximation) of the total number of principal cycles of length R, then one can go further than Equation (6), by applying the methods of [15, 17]. Indeed, noting that $H_p(D)\leq H(D)$, as in Section 2.4, and combining Theorems 5 and 6, we have an upper bound

$$\operatorname{Tr}(B(\ell_1^{e_1}\cdots\ell_r^{e_r}))\leq 2\prod_{i=1}^{r}\frac{\ell_i^{e_i+1}-1}{\ell_i-1} \tag{7}$$

and $\operatorname{Tr}(B(\ell_1^{e_1}\cdots\ell_r^{e_r}))$ is in $O(\ell_1^{e_1}\cdots\ell_r^{e_r})$. Hence, we get the following result on the total number of cycles or length R.

Corollary 24. *There exists a computable constant D, depending on r, such that number of equivalence classes of principal isogeny cycles of length R in the generalised graph $\mathcal{G}(p,L)$ is bounded above by $D\binom{R+r-1}{r-1}\ell_r^R\log^2(\ell_r)$.*

Proof. Noting that Equation (7) trivially holds when we have fractional coefficients, we can write

$$\begin{aligned}
\sum_{i=1}^{n} \mathcal{C}_{E_i}(L;R) &\leq 2 \sum_{e_1+\cdots+e_r=R} \sum_{J\subseteq\{1,\ldots,r\}} \frac{\left(\prod_{i=1}^{r} \ell_i^{e_i+1}\right) / \left(\prod_{j\in J} \ell_j^2\right)}{\prod_{i=1}^{r} \ell_i - 1} \\
&= 2 \sum_{e_1+\cdots+e_r=R} \prod_{i=1}^{r} \frac{\ell_i^{e_i+1}}{\ell_i - 1} \sum_{J\subseteq\{1,\ldots,r\}} \prod_{j\in J} \frac{1}{\ell_j^2} \\
&= 2 \sum_{e_1+\cdots+e_r=R} \prod_{i=1}^{r} \frac{\ell_i^{e_i+1}}{\ell_i - 1} (1 + \ell_i^{-2}) \\
&= 2 \sum_{e_1+\cdots+e_r=R} \prod_{i=1}^{r} \ell_i^{e_i} \left(1 + \frac{1}{\ell_i} + \frac{2}{\ell_i(\ell_i-1)}\right) \\
&\leq 2c_1 \sum_{e_1+\cdots+e_r=R} \prod_{i=1}^{r} \ell_i^{e_i} \left(1 + \frac{2}{\ell_i}\right),
\end{aligned}$$

where $c_1 \leq 4$ depends on r; for instance if $r \geq 6$, we can take $c_1 = 1$. Furthermore, using known bounds on the sum of reciprocals of primes (cf. Equation (3.20) in [21]), we have

$$\begin{aligned}
\prod_{i=1}^{r} \left(1 + \frac{2}{\ell_i}\right) &\leq \exp\left[\sum_{i=1}^{r} \frac{2}{\ell_i}\right] \\
&\leq \exp\left[2\left(\log\log(\ell_r) + 0.2615 + \frac{1}{\log^2(\ell_r)}\right)\right] \\
&\leq c_2 \log^2(\ell_r),
\end{aligned}$$

with $c_2 \leq 3$ if $r \geq 4$. Letting $D := 2c_1c_2$, and bounding $\prod_{i=1}^{r} \ell_i^{e_i}$ by ℓ_r^R, we obtain the desired result. □

Of course, one could get a better upper bound than in Inequality (7) and thus a better upper bound for $\sum_{i=1}^{n} \mathcal{C}_{E_i}(L;R)$, depending on one's intended application.

For a heuristic *estimate*, rather than an upper bound, we can argue (as in [17, Section 5]) that we expect to have $H_p(D) \approx \frac{1}{2}H(D)$ on average. This is because the prime p in the definition of $H_p(D)$ in Equation (2) splits half the time and remains inert the other half. This heuristic assumption gives us the estimate

$$\mathrm{Tr}(B(\ell_1^{e_1}\cdots\ell_r^{e_r})) \approx \prod_{i=1}^{r} \frac{\ell_i^{e_i+1}-1}{\ell_i - 1} - \frac{1}{2} \sum_{d \mid \ell_1^{e_1}\ldots\ell_r^{e_r}} \min\left\{d, \frac{\ell_1^{e_1}\cdots\ell_r^{e_r}}{d}\right\}, \tag{8}$$

which can then be substituted into Equation (6) of Theorem 23. Again, we will not perform this substitution, to avoid writing a very long and cumbersome

expression for the approximation of $\sum_{i=1}^{n} \mathcal{C}_{E_i}(L;R)$, but this can very easily be done (by hand or on a computer), especially once the number of primes in the set L is known. The advantage of this approach is that it is much faster to compute the right-hand side of Equation (8) than the actual trace of $B(\ell_1^{e_1} \cdots \ell_r^{e_r})$.

4.2. An ideal interpretation of certain isogeny cycles

In this section, we restrict to $\{\ell_1, \ldots, \ell_r\}$-isogeny cycles, where $p \neq \ell_i$ and $\ell_i < \ell_{i+1}$ for all i, to give an ideal-theoretic interpretation.

As isogeny cycles correspond to endomorphisms of elliptic curves, there is a connection to the theory of embeddings into quaternion orders. For a prime ℓ, the Deuring correspondence [12] gives a bijection between ℓ-isogenies of supersingular elliptic curves over $\overline{\mathbb{F}}_p$ and integral left ideals of norm ℓ, of maximal orders in $\mathcal{B}_p$, up to isomorphisms of curves/orders. A non-scalar element α of a maximal quaternion order (associated to a non-scalar endomorphism of a supersingular elliptic curve) generates an imaginary quadratic order over $\mathbb{Z}$. In this way, we view isogeny cycles as embeddings of imaginary quadratic orders into the maximal quaternion order. We count isogeny cycles by relating them to these embeddings, extending work of [22, 23, 14] to the L-isogeny graph setting. We recall some of the results and definitions from these works briefly here. For the purposes of this discussion, $\mathcal{O}$ will always denote an order in an imaginary quadratic field.

Definition 25 (Primitive $\mathcal{O}$-embedding [23, Def. 3.1, 3.3])**.** *Let $\mathcal{O}$ denote an imaginary quadratic order, $K := \mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q}$, and let E be a supersingular elliptic curve. An embedding $\iota : K \hookrightarrow \mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}$ is called a* primitive $\mathcal{O}$*-embedding if $\iota(K) \cap \mathrm{End}(E) = \iota(\mathcal{O})$.*

When focusing on the imaginary quadratic order rather than the field, we may denote a primitive $\mathcal{O}$-embedding ι by $\iota : \mathcal{O} \hookrightarrow \mathrm{End}(E)$.

Definition 26 ([23, Sec. 3.1])**.** *Let $\mathrm{SS}_{\mathcal{O}}^{pr}$ denote the set of isomorphism classes of supersingular elliptic curves E together with a choice of primitive $\mathcal{O}$-embedding into $\mathrm{End}(E)$.*

Recall that we let $h_{\mathcal{O}}$ denote the class number of the imaginary quadratic order $\mathcal{O}$.

Proposition 27. *The class group of $\mathcal{O}$ acts freely on $\mathrm{SS}_{\mathcal{O}}^{pr}$ and has one or two orbits, depending on if p is ramified or inert (resp.) in the imaginary quadratic field containing $\mathcal{O}$. In particular, $|\mathrm{SS}_{\mathcal{O}}^{pr}| = h_{\mathcal{O}}$ if p is ramified and $|\mathrm{SS}_{\mathcal{O}}^{pr}| = 2h_{\mathcal{O}}$ if p is inert.*

If $\mathfrak{l}_i$ are ideals above ℓ_i in $\mathcal{O}$ such that $\prod_{i=1}^{r} \mathfrak{l}_i$ is principal, then the endomorphisms of the elliptic curves in $\mathrm{SS}_{\mathcal{O}}^{pr}$ which correspond to $\prod_{i=1}^{r} \mathfrak{l}_i$ have canonical decompositions which are principal $\{\ell_1, \ldots, \ell_r\}$-isogeny cycles.

Proof. This proposition condenses several results in [23, 14], and applies these results to the case of $\{\ell_1, \ldots, \ell_r\}$-isogenies. □

Definition 28. *Define the set $\mathcal{I}_p$ of imaginary quadratic orders $\mathcal{O}$ such that:*

- *p is inert in $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q}$ and*
- *p does not divide the conductor of $\mathcal{O}$.*

The first condition guarantees that $|\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}| = 2h_{\mathcal{O}}$.

Definition 29. *For an imaginary quadratic order $\mathcal{O}$, let $r_{\mathcal{O}}(n)$ denote the number of elements of norm n in $\mathcal{O}$.*

The set $\mathcal{I}_p$ is infinite, but for any fixed n, $\mathcal{I}_p$ contains finitely many orders for which $r_{\mathcal{O}}(n)$ is nonzero. With this classification of orders in mind, we now count isogeny cycles:

Theorem 30. *Let $p, \ell_1, \ldots, \ell_r$ be distinct primes. The number of distinct canonical decompositions of principal $\{\ell_1, \ldots, \ell_r\}$-isogeny cycles in the supersingular L-isogeny graph over $\overline{\mathbb{F}}_p$ is*

$$|\mathrm{Cycles}(\{\ell_1, \ldots, \ell_r\})^{\mathrm{prin}} / \sim| = \sum_{\mathcal{O} \in \mathcal{I}_p} \frac{h_{\mathcal{O}} \cdot r_{\mathcal{O}}(\ell_1 \cdots \ell_r)}{2}. \tag{9}$$

Proof. Let $n = \prod_{i=1}^{r} \ell_i$. For the sake of this proof, we will consider elements $(E, \iota) \in \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ up to complex conjugation. That is, if θ is a generator of the imaginary quadratic order $\mathcal{O}$ and $\iota(\theta) = \iota'(\overline{\theta})$ where $\overline{\theta}$ denotes the complex conjugate of θ, we will consider (E, ι) to be equivalent to (E, ι'). Denote these equivalence classes by $[(E, \iota)] \in (\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim)$. Note that $|\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}| = 2|\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim|$.

Let $R_{\mathcal{O}}(n)$ denote the set of elements of an imaginary quadratic order $\mathcal{O}$ whose norm is equal to n. Denote by $r_{\mathcal{O}}(n) = |R_{\mathcal{O}}(n)|$. Associated to a fixed set $\mathcal{I}_p$, we consider the set

$$\bigsqcup_{\mathcal{O} \in \mathcal{I}_p} (\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim) \times (R_{\mathcal{O}}(n) / \mathcal{O}^{\times}) =$$
$$\{([(E, \iota)], \{\pm\alpha\}) : \mathcal{O} \in \mathcal{I}_p, [(E, \iota)] \in (\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim), \alpha \in \mathcal{O}, N(\alpha) = n\}.$$

The size of this set is

$$\sum_{\mathcal{O} \in \mathcal{I}_p} |\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim| \cdot \frac{r_{\mathcal{O}}(n)}{|\mathcal{O}^{\times}|} = \sum_{\mathcal{O} \in \mathcal{I}_p} \frac{h_{\mathcal{O}} \cdot r_{\mathcal{O}}(n)}{2},$$

as $|\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim| = h_{\mathcal{O}}$ since p is inert in $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q}$ and $|\mathcal{O}^{\times}| = 2$ since $p \equiv 1 \pmod{12}$.

We define a map:

$$\rho : \bigsqcup_{E} (\mathrm{Cycles}(\{\ell_1, \ldots, \ell_r\})_E^{\mathrm{prin}} / \sim) \to \bigsqcup_{\mathcal{O} \in \mathcal{I}_p} (\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} / \sim) \times (R_{\mathcal{O}}(n) / \mathcal{O}^{\times}) \tag{10}$$

as follows. Given an $\{\ell_1, ..., \ell_r\}$-isogeny cycle, let α denote the endomorphism obtained by composing the isogenies in this cycle starting at a vertex with isomorphism class representative E. By definition, $N(\alpha) = n$ and α is an imaginary quadratic integer. Let $\mathcal{O} = \mathbb{Z}[\alpha]$. The element α determines an

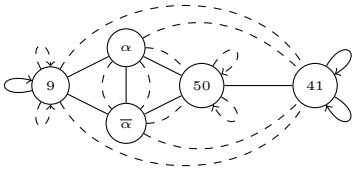

Figure 4: The supersingular $\{2,3\}$-isogeny graph over $\overline{\mathbb{F}}_{61}$, with vertices labelled by j-invariant in $\mathbb{F}_{61^2}$. Solid lines are 2-isogenies, dashed lines are 3-isogenies, and $\alpha, \overline{\alpha}$ denote conjugate j-invariants in $\mathbb{F}_{61^2} \setminus \mathbb{F}_{61}$. Loops may only be traversed in one direction, while all other edges are undirected and may be traversed in either direction.

embedding $\iota : \mathcal{O} \hookrightarrow \mathrm{End}(E)$, so p is inert in $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q}$ and $\mathcal{O} \in \mathcal{I}_p$. The orientation ι is determined by its image on the generator of $\mathcal{O}$, which is unique up to complex conjugation in $\mathcal{O}$. To see that ι is a primitive embedding (and thus $[(E,\iota)] \in \mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$), we show that there does not exist an endomorphism $\beta \in \mathrm{End}(E)$ such that $\alpha = [k] \circ \beta$ for some integer $k \neq \pm 1$. Since the norm of α is squarefree and $N([k] \circ \beta) = k^2 N(\beta)$, such a decomposition is indeed not possible. The map ρ is well-defined, and we will complete the proof by showing that ρ is a bijection.

To see that ρ is surjective, consider $((E,\iota), \{\pm\alpha\})$ where $\iota : \mathcal{O} \hookrightarrow \mathrm{End}(E)$ is a primitive embedding. The element $\alpha \in \mathrm{End}(E)$ necessarily has cyclic kernel, as it has squarefree degree, and so α can be factored into isogenies of prime-power degree, thus determining a principal isogeny cycle in $\mathrm{Cycles}(\{\ell_1, \ldots, \ell_r\})^{\mathrm{prin}} / \sim$.

To see that ρ is injective, note that an isogeny cycle determines the element α uniquely up to automorphism and $\mathrm{Aut}(E) = \{[\pm 1]\} = \mathcal{O}^\times$. The map ρ is thus a bijection, and by our previous count of the size of the codomain we get the formula in Equation (9). □

4.3. Examples

Code for these examples is given in `cycle_counts.ipynb`.

Example 31 (Principal $\{2,3\}$-isogeny cycles for $p = 61$)**.** For a small example, we wish to count the $\{2,3\}$-isogeny cycles in the supersingular $\{2,3\}$-isogeny graph over $\overline{\mathbb{F}}_{61}$, see Figure 4. In this case, *all* $\{2,3\}$-isogeny cycles are principal, as there is no possibility of backtracking.

A count in the graph finds ten distinct principal $\{2,3\}$-isogeny cycles:

$$(9,9), (9,9), (\alpha,\overline{\alpha}), (\overline{\alpha},\alpha), (\alpha,\overline{\alpha}), (\overline{\alpha},\alpha), (\alpha,50), (50,\alpha), (\overline{\alpha},50), (50,\overline{\alpha}).$$

Counting using ideals. First, note that $r_{\mathcal{O}}(6) = 0$ for any $\mathcal{O}$ of discriminant greater in absolute value than $4 \cdot 6$. The discriminants of orders in $\mathcal{I}_p$ which also satisfy $|\operatorname{disc}(\mathcal{O})| \leq 24$ are precisely:

$$-7, -8, -11, -23, -24.$$

These are discriminants of maximal orders, in number fields with class numbers $1, 1, 1, 3$, and 2, respectively. The elements of norm 6 can be found by factoring the ideals above 2 and 3 and searching for the generators of principal ideals of norm 6. The results are summarized in Table 1. Putting these values into the formula from Theorem 30:

$$|\mathrm{Cycles}(\{2,3\})^{\mathrm{prin}}/\sim| = \frac{1\cdot 4}{2} + \frac{3\cdot 4}{2} + \frac{2\cdot 2}{2} = 10.$$

$\mathrm{disc}(\mathcal{O})$	$h_{\mathcal{O}}$	Elements of norm 6
-7	1	-
-8	1	$\pm\left(\sqrt{-2}+2\right), \pm\left(\sqrt{-2}-2\right)$
-11	1	-
-23	3	$\pm\left(\frac{1+\sqrt{-23}}{2}\right), \pm\left(\frac{1-\sqrt{-23}}{2}\right)$
-24	2	$\pm\sqrt{-6}$

Table 1: Elements counted in the ideal-theoretic count of isogeny cycles in Example 31

Counting using Brandt matrices:. Let us now count the number of $\{2,3\}$-isogeny cycles, but using the Brandt matrix method. As in Section 4.1, we have that $|\operatorname{Cycles}_{61}(2,3)^{\mathrm{prin}}/\sim| = \operatorname{Tr}(B(2\cdot 3))$. This number can be computed directly using SageMath since we are working with small numbers.

Moreover, as discussed in Section 4.1, if we wanted to avoid computing this number directly, via Brandt matrices, we can use Theorems 5 and 6, together with the bound $H_p(D) \leq H(D)$ or the estimate $H_p(D) \approx \frac{1}{2}H(D)$. Indeed, we can easily compute (and this remains easy to do even when the degree, given by its prime factorisation, grows larger):

$$2\sum_{\substack{d\mid 6\\ d>0}} d - \sum_{\substack{d\mid 6\\ d>0}} \min\{d, 6/d\} = 18. \tag{11}$$

So, we thus obtain the bound $|\operatorname{Cycles}_{61}(2,3)^{\mathrm{prin}}/\sim| \leq 18$ and the estimate $|\operatorname{Cycles}_{61}(2,3)^{\mathrm{prin}}/\sim| \approx 9$, which is not far at all from the true answer 10.

Here we provide an example of a $\{\ell_1^{e_1}, \ldots, \ell_r^{e_r}\}$-isogeny cycle which does not arise from the theory of ideals discussed in Section 4.2.

Example 32. We consider again the setup of Example 31, namely we take $L = \{2,3\}, p = 61$. A subgraph of $\mathcal{G}(p,L)$ is shown in Figure 5. The 4 supersingular curves pictured; $E_{50}, E_{41}, E_{\alpha}, E_{\overline{\alpha}}$, have embeddings of the imaginary quadratic order $\mathbb{Z}[\sqrt{-11}]$, which has conductor 2. For the latter 3 curves, such an embedding will be primitive. In the case of E_{50} we find that this curve has a primitive embedding of the maximal order $\mathbb{Z}[\frac{1+\sqrt{-11}}{2}]$ contained in the imaginary quadratic field $\mathbb{Q}(\sqrt{-11})$. Note these orders have class numbers 3 and 1 respectively.

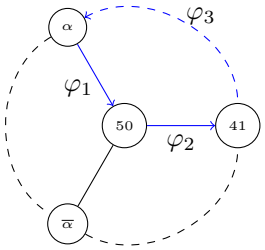

Figure 5: A subgraph of the supersingular $\{2,3\}$-isogeny graph over $\overline{\mathbb{F}}_{61}$ (see Figure 4).

Consider the endomorphism $\varphi_3 \circ \varphi_2 \circ \varphi_1$ of degree $2^2 \cdot 3$ depicted in blue. The isogeny φ_3 is horizontal with respect to the $\mathbb{Z}[\sqrt{-11}]$-embedding on E_{41} and E_α, and thus corresponds to the action of a prime ideal $\mathfrak{l}_3$ lying above $(3) \cdot \mathbb{Z}[\sqrt{-11}]$, namely $(3, 1+\sqrt{-11})$. On the other hand the 2-isogenies φ_1, φ_2 are ascending and descending respectively, and thus do not arise from the action of a prime ideal in this order.

Using degree $2^2 \cdot 3$ in our implementation of Theorem 30, yields a count of 0 principal isogeny cycles, when the correct count is 16. This is because our algorithm only counts cycles where steps are all horizontal in L-isogeny volcano rims, while some principal cycles such as $\varphi_3 \circ \varphi_2 \circ \varphi_1$ exist, which ascend and descend in the 2-isogeny volcano.

5. Graph cuts

In this section we study cuts of ℓ- and L-isogeny graphs with low edge expansion (see Definition 7). We no longer restrict to the case $p \equiv 1 \pmod{12}$. For a regular graph, when taking a random edge from a random vertex within the cut, a lower edge expansion implies a higher chance the edge is internal to the cut, and lower chance it leaves the cut. This is cryptographically relevant as most isogeny-based schemes use Ramanujan or rapid mixing properties of isogeny graphs. In particular, the distribution of endpoints of random walks approaches the discrete uniform distribution on the vertices quickly as the length of the walk increases. For low expansion cuts however, should they exist, the probabilities are imbalanced for vertices on opposite sides.

5.1. Edge expansion in small cuts

For convenience we now refer to a cut as a set of vertices $C \subset V$, where as a disjoint partition this is $V = C \sqcup (V \setminus C)$. We define a cut to be *(dis)connected* if the induced subgraph from the vertices in the cut is (dis)connected. As we are looking for cuts with relatively low edge expansion, we only care about connected cuts by the following result.

Lemma 33. *Let $G = (V, E)$ be a graph and C a disconnected cut such that $|C| \leq |V \setminus C|$. Then there exists a connected cut $T \subset C$ with $E(T, C \setminus T) = \varnothing$ and $h_G(T) \leq h_G(C)$.*

Proof. Let $C_1, C_2, \ldots, C_n$ be the vertex sets of the connected components of subgraph induced by C. Let E_i denote the set of edges in the subgraph induced by the vertices C_i. We have $\mathrm{Vol}(C) = \sum \mathrm{Vol}(C_i)$ and $|E(C, V \setminus C)| = \sum |E(C_i, V \setminus C_i)|$. Suppose for contradiction that $h_G(C_i) > h_G(C)$ for all i. Then

$$|E(C_i, V \setminus C_i)| = \mathrm{Vol}(C_i) \cdot h_G(C_i) > \mathrm{Vol}(C_i) \cdot h_G(C).$$

Summing over i gives $|E(C, V \setminus C)| > \mathrm{Vol}(C) \cdot h_G(C) = |E(C, V \setminus C|)$—this is a contradiction. Hence there must exist r and $T = C_r$ with $h_G(T) \leq h_G(C)$. □

We now consider small cuts of relatively small edge expansion in ℓ-isogeny graphs.

Lemma 34. *In a directed ℓ-isogeny graph G, take a connected cut C of n vertices with $n < k \cdot \log_\ell(p) + 1$ and define $R_{\ell,n} = \frac{n \cdot \ell + 1}{n}$. Suppose C contains no curves with additional automorphisms. Then $h_G(C) \leq R_{\ell,n}$ with equality if and only if the induced subgraph of C in G is a tree (with dual edges), and inequality implying C contains a curve with a cyclic endomorphism of degree at most p^k.*

Proof. Within G merge a maximal pairing of directed edges with their duals giving a graph with undirected edges and possibly loops which remain directed. It remains $(\ell + 1)$-regular, directed edges being self-dual loops or edges from curves with additional automorphisms. We ignore the latter, as no such curves are in C. Consider the induced subgraph G_C of C in G. Suppose G_C has no cycles (and so no loops). Then, it is by definition an undirected tree (since it is connected), and has $n - 1$ edges. As the graph is $(\ell + 1)$-regular, each vertex of C has $(\ell + 1)$ edges, so $E(C, V \setminus C) = n(\ell + 1) - (n - 1) = n\ell + 1$, giving $h_G(C) = R_{\ell,n}$. If G_C is not a tree, then it contains a (graph-theoretic) cycle, and so at least one additional edge giving $h_G(C) < R_{\ell,n}$. With n vertices the cycle has degree at most $\ell^{n-1} \leq \ell^{k \log_\ell(p)} = p^k$, and composes to a cyclic endomorphism by [2, Proposition 1]. □

One interpretation of *relatively* small edge expansion is compared to connected cuts of the same number of vertices n. Lemma 34 allows us to relate the distribution of such cuts to the previously studied distribution of p^k-valleys [24].

Remark 35. *By Lemma 34, small cuts with "relatively" small edge expansion must have edge expansion strictly less than $R_{\ell,n}$, as all these cuts have edge expansion less than or equal to $R_{\ell,n}$. Moreover, the lemma together with [25, Section 4] shows for a fixed constant $k < \frac{2}{3}$ and $n < k \cdot \log_\ell(p) + 1$, cuts of size n contain curves with endomorphisms of degree less that p^k. For sufficiently large p these curves are rare, appearing in small clusters which are far apart, referred to as p^k-valleys. For smaller k (and also n), the p^k-valleys become sparser [24].*

If we try the same argument for L-isogeny graphs, we run into an issue, due to the existence of graph-theoretic cycles which compose to scalar endomorphisms (see Figure 1). To maximize expansion, we seek to minimize cycles in the subgraphs induced by the cut. Trees have maximal expansion. In the ℓ-isogeny graph, consider non-backtracking paths in subtrees of the full graph. These paths correspond to cyclic isogenies which are not endomorphisms. Lemma 34 references the edge expansion of trees, and we need the appropriate analogue of a tree in the L-isogeny graph setting. In this setting, the correct structure is a connected induced subgraph of $\mathcal{G}(p, L)$ where all cycles compose to scalar endomorphisms, which we call a 'tree$^\blacklozenge$.'

Definition 36 (Minimum Uniform Edge Expansion $R_{L,n}$)**.** *Fix a set of primes primes $L = \{\ell_1, \ldots, \ell_r\}$ and $n > 0$. For a prime $p \notin L$ we define $R_{L,n}(p)$ as the minimum edge expansion over We then define $R_{L,n} = \lim_{p\to\infty} R_{L,n}(p)$.*

Lemma 37. *Fix L and n, as above. The limit $R_{L,n} = \lim_{p\to\infty} R_{L,n}(p)$ is well-defined, and is attained for sufficiently large p.*

Proof. The length of a cycle in an induced subgraph of $\mathcal{G}(p, L)$ on n vertices can be bounded by the fact there are exactly $n \cdot (\ell_i + 1)$ outgoing ℓ_i-isogenies, up to post-composition with automorphism, originating from vertices in the subgraph. Hence the degree of the endomorphism formed by composing a cycle can be upper bounded by $K = \prod_i \ell_i^{n\cdot(\ell_i+1)}$, which is independent of p. Hence for sufficiently large p we may assume $K \ll p$. By the results of [24], increasing p increases the distance between K-valleys until there is a large enough space in the graph for a curve E_p and depth-$(n + 1)$ neighbourhood of L-isogenies, disjoint from K-valleys. This means the neighbourhood is a tree$^\blacklozenge$, i.e. it contains no non-scalar cycles of degree less than or equal to K. The n-vertex cut defining $R_{L,n}(p)$ has an isomorphic cut (i.e., graph isomorphism of the induced subgraphs) within the neighbourhood of E_p. The neighbourhood of E_p is isomorphic as a graph to the neighbourhood arising from any larger value of p. Take a point far enough from the K-valleys so that the neighbourhood in the ℓ_i-isogeny graph looks like a tree, for each $i = 1, \ldots, n$. The cycles in the union of these ℓ_i-isogeny trees compose to scalar endomorphisms. Hence for large enough p, $R_{L,n} = R_{L,n}(p)$ is the minimum edge expansion of all n-vertex cuts in the neighbourhood of E_p. □

We now consider the analogue of Lemma 34 in L-isogeny graphs.

Lemma 38. *In a directed $L = \{\ell_1, \ldots, \ell_r\}$-isogeny graph, take a connected cut C of n vertices with $n < k \cdot \log_{\ell_r}(p) + 1$. Suppose C contains no curves with additional automorphisms. Then $h_G(C) < R_{L,n}$ implies C contains a curve with a cyclic endomorphism of degree at most p^k.*

Proof. Follows as in the proof of Lemma 34, replacing 'tree' with depth-n 'tree$^\blacklozenge$' from the definition of $R_{L,n}$. □

Remark 39. *For large p there are many depth-d L-isogeny neighbourhoods containing n vertices within $\mathcal{G}(p, L)$ without non-scalar cycles. An n-vertex cut*

C has "relatively small" edge expansion if $h_G(C) < R_{L,n}$. (For the L-isogeny graph, it is possible for $h_G(C) > R_{L,n}$, where as $h_G(C) \leq R_{\ell,n}$ by Lemma 34.) By Lemma 38 along with the result of [25, Section 4] implies that for $k < \frac{2}{3}$ and small enough n that any n-vertex cut of relatively small edge expansion, intersects a p^k-valley, which for larger p and smaller k are increasingly rare and sparse.

5.2. Clustering based on Fiedler's vector

We now discuss finding larger subsets of vertices which define cuts of relatively small edge expansion. We take an experimental approach on isogeny graphs $\mathcal{G}(p, L)$ with p as large as computationally feasible. Some preliminary attempts showed for p large enough, n-vertex sets with edge expansion less than $R_{L,n}$ in $\mathcal{G}(p, L)$ are rare. Searching through all n vertex subsets attempting to find one, being exponential time in p, is computationally infeasible. Instead of trying them all, techniques from spectral graph theory allow us to make a sensible guess.

While several algorithms exist using spectral methods to obtain low edge expansion cuts, we will use *Fiedler's algorithm*, which we explain here. For a graph $G = (V, E)$ with vertices denoted $v_1, \ldots, v_k$, it proceeds as follows:

1. Compute the eigenvector $\vec{x}$ of the Laplacian matrix corresponding to the second largest eigenvalue λ_2[4].
2. Order vertices v_i with respect to the quantity x_i, largest first, and relabel them $u_1, \ldots, u_k$ with respect to this ordering.
3. For $i = 1, \ldots, k$, let $C_i = \{u_1, \ldots, u_i\}$, called a *sweep cut*, and compute $h_G(C_i)$ and $h_G(V \setminus C_i)$.
4. Return the cut C_i such that $\max(h_G(C_i), h_G(V \setminus C_i))$ is minimised.

The algorithm is polynomial time in $|V| + |E|$, which for $G = \mathcal{G}(p, L)$ means polynomial time in p and $\sum \ell_i$. The eigenvector of the second largest eigenvalue λ_2—often called the *Fiedler vector*—is used as it is believed to provide a good measure of connectivity. Its use is attributable to Fiedler [26]. In his original work however, Fiedler suggested use of the *Fiedler cut* constructed as $C = \{v_i : x_i \geq 0\}$ for clustering. More generally, when using multiple eigenvectors, this is *hyperplane rounding*, which for a single vector amounts to picking a threshold to define the cut. A choice of 0 for the threshold is better suited to well-balanced graphs, while allowing it to vary can result in better clustering of unbalanced graphs. The concept of finding cuts to minimise edge expansion didn't arise until later, and it is now more common to use the minimal edge expansion to decide the threshold. Also, since we desire a partition into two 'inter-connected' clusters either side of the cut, we minimise $\max\{h_G(C_i), h_G(V \setminus C_i)\}$ rather than just $h_G(C_i)$, to ensure the complement also has low edge expansion. For random

[4]For regular graphs this is also the eigenvector of the second largest eigenvalue of the adjacency matrix.

graphs we expect the resulting cut size to be around half of the vertices. For a summary of spectral clustering techniques using Fiedler's vector see [27, 28].

Example 40. For the isogeny graph $\mathcal{G}(p, L)$ with $p = 419$ and $L = \{2, 3\}$, the second largest eigenvalue of the Laplacian is $\lambda_2 \approx 11.17$. The above algorithm finds sweep cut C_{18} has the minimal value $\max(h_G(C_{18}), h_G(V \setminus C_{18})) \approx 0.46$. The graph has a total of 36 vertices, so C_{18} contains exactly half. The code for this example in given in file `fiedler.ipynb`.

5.3. Other Algorithms for Computing Graph Cuts

To see if the cuts resulting from this algorithm have small edge expansion, when compared to other cuts of roughly the same size, we now give an alternative non-spectral method. To do this we replace the ordering of vertices in the previous algorithm with a different ordering. We want the vertices to form a subgraph that looks as 'complete' as possible. The most natural way to do this is to take the neighbourhood of a vertex. We present two variants of this idea.

Neighbour Expansion: Pick a starting vertex $[v_0]$ then add its neighbours $[v_0, v_1, v_2, v_3]$ then add the neighbours of v_1, then the neighbours of v_2 etc, until all vertices are ordered $[v_0, v_1, \ldots, v_n]$.

Greedy Neighbour Expansion: Pick a starting vertex $S_0 = [v_0]$ then define v_{i+1} to be the neighbour of vertices in S_i such that $\phi(S_i \cup \{v_{i+1}\})$ in minimised, then $S_{i+1} := S_i \cup \{v_{i+1}\}$. Here ϕ of a cut denotes the maximum of the edge expansion of the cut and the edge expansion of its complement.

Both of these algorithms are exponential in $\log(p)$ as they have to loop over all vertices in the graph. However the second is slower as it is quadratic in the number of vertices (as is Fiedler's algorithm), while the first is linear in the number of vertices.

The optimum solution, that is, the cut $S \subseteq V$ minimising the value of $\phi(S)$, is called the *Cheeger cut*. While computationally infeasible to compute, there are known upper and lower bounds,

$$\frac{\lambda_{-2}}{2} \leq \phi(S) \leq \sqrt{2\lambda_{-2}},$$

where λ_{-2} is the second smallest eigenvalue of the Laplacian. We also compute these bounds for reference.

Results are given in Figure 6 with code in file `fiedler_comparison.ipynb`. With the spectral ordering for reference, we try both algorithms three times, staring from different random vertices, and take the average of values ϕ found. From the results we make the following observations:

1. The spectral ordering performs badly, suggesting there is not one distinguishable low edge expansion cluster of size roughly $\frac{|V|}{2}$. Perhaps there are more, and the spectral ordering is merging such cuts, or perhaps better cuts have fewer vertices.

2. The neighbour ordering finds better cuts than the spectral ordering, suggesting neighbourhoods of vertices form smaller edge expansion clusters. The greedy neighbour ordering does even better, which motivates further study of the less-trivial structures that arise from this ordering.

3. As the size of the set L increases, the improvements reduce. The Cheeger lower bound also increases, and so there is less room for improvement. For the neighbour orderings this is likely as there are more edges, increasing the connectivity of any subset of vertices. This also suggests ℓ-isogeny graphs have smaller expansion clusters, and more of a bottleneck than L-isogeny graphs with larger L.

Figure 7 with code in `fiedler_viz.ipynb`, gives a visual comparison of these results. Each image represents an adjacency matrix where a non-zero entry is given by a dot. The different images correspond to vertices in the adjacency matrix ordered in different ways. We first order the vertices using the spectral ordering, neighbour ordering and greedy neighbour ordering. As the resulting cuts are around half the vertices, the edges within each cut are those in the top left quadrant of each image. The edges fully outside the cut are those in the bottom left quadrant. The edges between the cut and its complement lie in the off-diagonal quadrants. Hence a low edge expansion cut will have more vertices in the diagonal quadrants and less in the off-diagonal quadrants. To gain a perspective of the density within each quadrant, we give the same images with vertices within the cut randomly shuffled, and those in the complement also randomly shuffled.

For the spectral ordering we see this does no better than a random ordering of vertices, as all quadrants are equally full. The neighbour ordering sees some white space in the off diagonal quadrants, since for the first few curves we pick, their neighbours are guaranteed to also be early on in the ordering, and so there are no edges to vertices at end of the ordering. The greedy neighbour ordering, which performs best, essentially pushing all curves which are not close to the starting curves towards the end of the ordering.

p, L	Spectral Ordering	Neighbour Ordering	Greedy Ordering	Cheeger Lower Bound	Cheeger Upper Bound
419, {3}	0.597	0.319	0.225	0.071	0.533
419, {2,3}	0.452	0.289	0.286	0.096	0.619
419, {2,3,5,7,11}	0.488	0.491	0.413	0.151	0.776
5569, {3}	0.494	0.301	0.178	0.068	0.520
5569, {2,3}	0.484	0.322	0.195	0.092	0.608
5569, {2,3,5,7,11}	0.497	0.473	0.343	0.275	1.049
10007, {3}	0.489	0.287	0.170	0.065	0.508

Figure 6: Fiedler's algorithm with different vertex orderings.

6. Future work

There are several notable areas for further study.

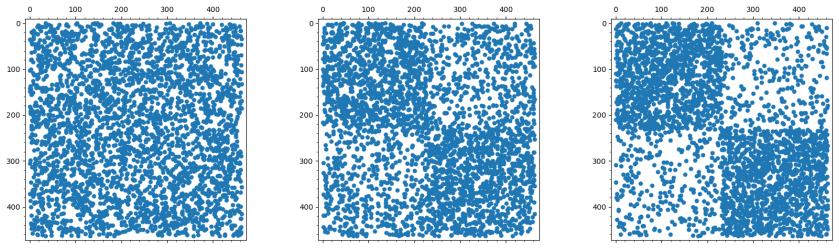

(a) Spectral ordering (b) Neighbour ordering (c) Greedy neighbour ordering

Figure 7: Adjacency matrix plot for $p = 5569$ and $L = \{2, 3\}$. Dots are edges. Vertices are ordered so that the first half of the vertices come from one side of the cut, and the second half of the vertices come from the other side.

6.1. Further principal cycle counting

Firstly, within Sections 3 and 4 we restricted to the case $p \equiv 1 \pmod{12}$. For other values of p, curves with additional automorphisms exist and so further consideration is required to determine if, and how, our principal cycle counts could be adapted to address this.

Furthermore Theorem 30, counting principal $(\ell_1^{e_1}\ell_2^{e_2}\cdots\ell_r^{e_r})$-isogeny cycles via ideals in class groups of imaginary quadratic orders, only works for products of distinct primes, with all exponents $e_i = 1$. Example 32 demonstrated how the approach fails for larger exponents. The issue arises for $e_i \geq 2$ and a starting curve with a primitive $\mathcal{O}$-embedding with $\ell_i \mid \mathrm{cond}(\mathcal{O})$, cycles may arise from ascending the ℓ_i-isogeny volcano then descending back down. These ascending and descending isogenies are not represented by the action of invertible ideals in $cl(\mathcal{O})$.

One method to address this is to use an approach based on the volcano 'rim-hopping' algorithm of [29]. Given the norm $\ell_1^{e_1}\ell_2^{e_2}\cdots\ell_r^{e_r}$ generator η of $\mathcal{O}$, one can obtain another generator from a translate $k+\eta$ with $k \in \mathbb{Z}$. Then factorising its norm $N(k+\eta) = q_1^{f_1}\cdots q_m^{f_m}$ look for principal products of ideals in $\mathcal{O}$ which consist of f_i ideals of norm q_i for each i. If k is chosen such that $q_i \nmid \mathrm{cond}(\mathcal{O})$, in terms of endomorphisms we obtain all degree $q_1^{f_1}\cdots q_m^{f_m}$ endomorphisms φ which give primitive $\mathcal{O}$-embeddings. Translating back, $\psi = \varphi - [k]$, we obtain all degree $\ell_1^{e_1}\ell_2^{e_2}\cdots\ell_r^{e_r}$ endomorphisms giving primitive $\mathcal{O}$-embeddings. From the product of norm q_i ideals we can easily check if φ is cyclic by ensuring ideal inverses do not appear in the product. However, further thought is needed to determine how to check, from the ideal product alone, which translated endomorphisms ψ will be cyclic. It is also possible more careful analysis of oriented isogeny volcanoes may result in an alternative approach which more closely matches Theorem 30.

6.2. Scalar cycle counting

Within Section 4, we restricted our counts to principal isogeny cycles, up to the equivalence of having the same canonical decomposition. Endomorphisms

arising from principal cycles have the advantage that all of their decompositions yield genuine (non-backtracking) cycles. However, counting principal cycles is not the same as counting graph-theoretic cycles. To bridge this gap, one would have to also consider non-principal cycles, as well as cycles not up to canonical equivalence. The task of counting non-principal cycles additionally requires determining how many prime decompositions of a given non-cyclic endomorphism have backtracking. In particular, if we counted the number of graph-theoretic cycles which compose to scalar endomorphisms, this would be a step towards understanding the L-isogeny cycles generally. This is also closely tied to computing the minimum uniform edge expansion $R_{L,n}$ defined in Section 5.1.

If $[n] = [\ell_1 \ell_2 \cdots \ell_r]$, the total number of walks corresponding to $[n]$ is $\frac{(2r)!}{2^r} \prod_{i=1}^{r} (\ell_i + 1)$. If $r = 1$, none of these are a graph-theoretic cycle (non-backtracking). If $r = 2$, precisely $2(\ell_1 + 1)(\ell_2 + 1)$ of these walks is a graph-theoretic cycle (without backtracking), namely the isogeny diamond and the same diamond traversed in the opposite direction. If $r > 2$, an inclusion-exclusion argument can be applied to find the number of graph-theoretic cycles.

6.3. Further graph clustering

From Section 5.3 the unexpectedly good performance of the greedy neighbour ordering prompts questions into the structure of this ordering, and what the consequences of this might be for sampling isogenies in cryptography. There are also alternative graph clustering techniques, such as flow-based clustering, where the results could give better indications of graph structure.

One may also wish to consider a more extreme interpretation of small clusters; those with a single vertex. Their connectivity with the rest of the graph can be represented by a *distance distribution*.

Definition 41 (Distance distribution)**.** *The distance distribution for a vertex $v \in V$ of a graph $G = (V, E)$ is defined as the set $(N_0(v), N_1(v), \dots)$, where $N_i(v)$ is equal to the number of vertices in the graph G which are distance i from the vertex v.*

For random graphs this typically resembles a normal distribution, and the same seems to hold for isogeny graphs, as can be seen in Figure 8. Examining outliers and skew of these distributions could be insightful.

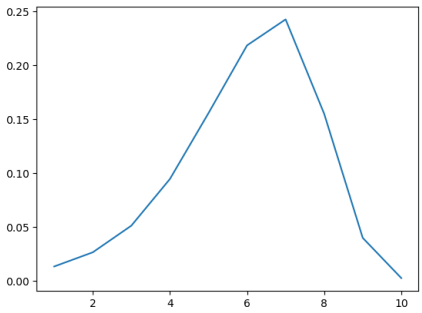

Figure 8: Distance distribution in 2-isogeny graph with $p = 2689$, from vertex with j-invariant 30. Distance is denoted by x-axis, and the y-axis gives the proportion of vertices of this distance.

References

[1] A. Pizer, An algorithm for computing modular forms on $\Gamma_0(N)$, J. Algebra 64 (2) (1980) 340–390. doi:10.1016/0021-8693(80)90151-9.
URL https://doi.org/10.1016/0021-8693(80)90151-9

[2] D. X. Charles, K. E. Lauter, E. Z. Goren, Cryptographic hash functions from expander graphs, J. Cryptology 22 (1) (2009) 93–113. doi:10.1007/s00145-007-9002-x.
URL https://doi.org/10.1007/s00145-007-9002-x

[3] C. Delfs, S. D. Galbraith, Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$, Des. Codes Cryptogr. 78 (2) (2016) 425–440. doi:10.1007/s10623-014-0010-1.
URL https://doi.org/10.1007/s10623-014-0010-1

[4] E. Bank, C. Camacho-Navarro, K. Eisenträger, T. Morrison, J. Park, Cycles in the supersingular ℓ-isogeny graph and corresponding endomorphisms, in: Research Directions in Number Theory: Women in Numbers IV, Springer, 2019, pp. 41–66.

[5] K. Eisenträger, S. Hallgren, C. Leonardi, T. Morrison, J. Park, Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs, in: ANTS XIV—Proceedings of the Fourteenth Algorithmic Number Theory Symposium, Vol. 4 of Open Book Ser., Math. Sci. Publ., Berkeley, CA, 2020, pp. 215–232. doi:10.2140/obs.2020.4.215.
URL https://doi.org/10.2140/obs.2020.4.215

[6] J. Fuselier, A. Iezzi, M. Kozek, T. Morrison, C. Namoijam, Computing supersingular endomorphism rings using inseparable endomorphisms (2023). arXiv:2306.03051.
URL https://arxiv.org/abs/2306.03051

[7] E. Orvis, Distribution of cycles in supersingular ℓ-isogeny graphs, available online: https://arxiv.org/abs/2403.14831 (2024). arXiv:2403.14831.
URL https://arxiv.org/abs/2403.14831

[8] W. Stein, et al., Sage Mathematics Software (Version 10.5), The Sage Development Team, http://www.sagemath.org (2025).

[9] Y. Kambe, A. Katayama, Y. Aikawa, Y. Ishihara, M. Yasuda, K. Yokoyama, Computing endomorphism rings of supersingular elliptic curves by finding cycles in concatenated supersingular isogeny graphs, Comment. Math. Univ. St. Pauli 72 (2024) 19–42.

[10] J. H. Silverman, Advanced topics in the arithmetic of elliptic curves, Vol. 151 of Graduate Texts in Mathematics, Springer-Verlag, New York, 1994. doi:10.1007/978-1-4612-0851-8.
URL https://doi.org/10.1007/978-1-4612-0851-8

[11] J. Voight, Quaternion Algebras, Vol. 288 of Graduate Texts in Mathematics, Springer, Cham, 2021. doi:10.1007/978-3-030-56694-4.
URL https://doi.org/10.1007/978-3-030-56694-4

[12] M. Deuring, Die Typen der Multiplikatorenringe elliptischer Funktionenkörper., Abh. Math. Sem. Hansischen Univ. 14 (1941) 197–272.

[13] J. H. Silverman, The arithmetic of elliptic curves, 2nd Edition, Vol. 106 of Graduate Texts in Mathematics, Springer, Dordrecht, 2009. doi:10.1007/978-0-387-09494-6.
URL https://doi.org/10.1007/978-0-387-09494-6

[14] S. Arpin, M. Chen, K. E. Lauter, R. Scheidler, K. E. Stange, H. T. N. Tran, Orientations and cycles in supersingular isogeny graphs, in: Research directions in number theory, Vol. 33 of Assoc. Women Math. Ser., Springer, Cham, 2024, pp. 25–86. doi:10.1007/978-3-031-51677-1_2.
URL https://doi.org/10.1007/978-3-031-51677-1_2

[15] W. Ghantous, Loops, multi-edges and collisions in supersingular isogeny graphs, Advances in Mathematics of Communications (2022).

[16] B. Gross, Heights and the special values of L-series, in: Conference Proceedings of the CMS, Vol. 7, 1987, pp. 115–186.

[17] W. Ghantous, S. Katsumata, F. Pintore, M. Veroni, Collisions in supersingular isogeny graphs and the SIDH-based identification protocol, Cryptology ePrint Archive, Paper 2021/1051 (2021).
URL https://eprint.iacr.org/2021/1051

[18] A. Hurwitz, Ueber relationen zwischen classenanzahlen binärer quadratischer formen von negativer determinante, Mathematische Annalen 25 (2) (1885) 157–196.

[19] A. Tucker, Applied Combinatorics, 5th Edition, John Wiley & Sons, Nashville, TN, 2007.

[20] E. H. Friend, Sorting on electronic computer systems, J. ACM 3 (3) (1956) 134–168. doi:10.1145/320831.320833.
URL https://doi.org/10.1145/320831.320833

[21] J. B. Rosser, L. Schoenfeld, Approximate formulas for some functions of prime numbers, Illinois Journal of Mathematics 6 (1) (1962) 64–94.

[22] L. Colò, D. Kohel, Orienting supersingular isogeny graphs, J. Math. Cryptol. 14 (1) (2020) 414–437. doi:10.1515/jmc-2019-0034.
URL https://doi.org/10.1515/jmc-2019-0034

[23] H. Onuki, On oriented supersingular elliptic curves, Finite Fields Appl. 69 (2021) Paper No. 101777, 18. doi:10.1016/j.ffa.2020.101777.
URL https://doi.org/10.1016/j.ffa.2020.101777

[24] J. Love, D. Boneh, Supersingular curves with small noninteger endomorphisms, in: ANTS XIV—Proceedings of the Fourteenth Algorithmic Number Theory Symposium, Vol. 4 of Open Book Ser., Math. Sci. Publ., Berkeley, CA, 2020, pp. 7–22. doi:10.2140/obs.2020.4.7.
URL https://doi.org/10.2140/obs.2020.4.7

[25] N. D. Elkies, The existence of infinitely many supersingular primes for every elliptic curve over $\mathbf{Q}$, Invent. Math. 89 (3) (1987) 561–567. doi:10.1007/BF01388985.
URL https://doi.org/10.1007/BF01388985

[26] M. Fiedler, Algebraic connectivity of graphs, Czechoslovak Math. J. 23 (2) (1973) 298–305.

[27] J. Wyss-Gallifent, Lecture notes on graph theory, https://www.math.umd.edu/~immortal/MATH246/lecturenotes/ch_all.pdf (2021).

[28] M. W. Mahoney, Lecture notes on spectral graph methods, CoRR abs/1608.04845 (2016). arXiv:1608.04845.
URL http://arxiv.org/abs/1608.04845

[29] S. Arpin, J. Clements, P. Dartois, J. K. Eriksen, P. Kutas, B. Wesolowski, Finding orientations of supersingular elliptic curves and quaternion orders, Des. Codes Cryptogr. 92 (11) (2024) 3447–3493. doi:10.1007/s10623-024-01435-5.
URL https://doi.org/10.1007/s10623-024-01435-5