ORIENTATIONS AND CYCLES IN SUPERSINGULAR ISOGENY GRAPHS

SARAH ARPIN, MINGJIE CHEN, KRISTIN E. LAUTER, RENATE SCHEIDLER, KATHERINE E. STANGE, HA T. N. TRAN

ABSTRACT. The paper concerns several theoretical aspects of oriented supersingular ℓ-isogeny volcanoes and their relationship to closed walks in the supersingular ℓ-isogeny graph. Our main result is a bijection between the rims of the union of all oriented supersingular ℓ-isogeny volcanoes over $\overline{\mathbb{F}}_p$ (up to conjugation of the orientations), and *isogeny cycles* (non-backtracking closed walks which are not powers of smaller walks) of the supersingular ℓ-isogeny graph over $\overline{\mathbb{F}}_p$. The exact proof and statement of this bijection are made more intricate by special behaviours arising from extra automorphisms and the ramification of p in certain quadratic orders. We use the bijection to count isogeny cycles of given length in the supersingular ℓ-isogeny graph exactly as a sum of class numbers of these orders, and also give an explicit upper bound by estimating the class numbers.

1. INTRODUCTION

Fix primes $\ell < p$. The supersingular ℓ-isogeny graph $\mathcal{G}_\ell$ is the directed graph whose vertices are the $\overline{\mathbb{F}}_p$-isomorphism classes of supersingular elliptic curves, and whose edges are isogenies, up to post-composition by an automorphism. The graph is $(\ell+1)$-regular, has around $p/12$ vertices, and consists of one connected component. However, it has few obvious symmetries: its structure more closely resembles that of a random regular graph in its statistical behaviour, including its graph spectrum. It is a Ramanujan graph.

The graph $\mathcal{G}_\ell$ was introduced into cryptography in 2006 in [15], where the hard problem of finding paths in the graph was proposed. In the cryptographic setting, p is always taken to be a large prime of cryptographic size and ℓ is a small prime like $\ell = 2$ or 3. A key exchange protocol called SIDH based on the path-finding problem in $\mathcal{G}_\ell$ was proposed in [21] and was considered for standardization in the NIST Post-Quantum Cryptography process (2017–2022). Although SIDH has been broken [11, 40, 47], so far there do not exist polynomial time algorithms (polynomial in $\log p$) to solve either the *path finding problem* (to find a path joining two given supersingular curves), or the *endomorphism ring problem* (to compute the endomorphism ring, either as a maximal order in a quaternion algebra or by giving a basis of endomorphisms) of a supersingular curve. These problems are closely related to one another [24, 57], and to other cryptographic protocols such as CSIDH [13] and SQISign [22].

The graph $\mathcal{G}_\ell$, like any superhero worth its salt, has an 'alter ego,' which is as the graph of maximal orders in a quaternion algebra $B_{p,\infty}$ ramified at p and ∞. For each elliptic curve vertex, the endomorphism ring of the curve is such a maximal order. However, even viewed in this way, the graph is still its apparently disordered self. To reveal some familiar structure with which to navigate the graph, one can focus, one at a time, on the quadratic rings embedded in the quaternion algebra.

The aforementioned structure recalls to mind the analogous graph for ordinary elliptic curves, whose endomorphism rings are merely quadratic orders. One can form an ordinary ℓ-isogeny graph just as one does a supersingular one: vertices are isomorphism classes of ordinary elliptic curves, and edges are ℓ-isogenies [26] [52]. In this case, the endomorphism rings of the curves in a connected component are all orders in the same quadratic field. Stratifying the graph according to the conductors of these orders, with larger orders at

FIGURE 1. An oriented 3-isogeny volcano. The rim is comprised of the five black vertices. These vertices are primitively oriented by a 3-fundamental order $\mathcal{O}$. The grey vertices at altitude 1 of the volcano are primitively oriented by the suborder of $\mathcal{O}$ of index 3. The five thick dotted edges along the rim indicate horizontal oriented 3-isogenies, and the remaining black edges are ascending/descending.

higher 'altitude,' we see at the top a cycle, or 'rim,' from which trees descend down the sides of the 'volcano' (see Figure 1). This simple, organized picture is in contrast to the complexity of the supersingular graph.

The storyline of the present paper is that the (finite) supersingular graph $\mathcal{G}_\ell$ can be understood as the result of superimposing infinitely many volcanoes, each obtained by focusing on a single quadratic field embedded inside the quaternion algebra $B_{p,\infty}$ and its intersection with the various maximal orders. The volcanoes in this case are not the volcanoes of the ordinary ℓ-isogeny graph, but rather the connected components of the *oriented supersingular ℓ-isogeny graph associated to a quadratic field K*, denoted by $\mathcal{G}_{K,\ell}$. This graph has recently been studied by Colò-Kohel [17] and Onuki [42] in the cryptographic context. The vertices of $\mathcal{G}_{K,\ell}$ are pairs (E, ι) consisting of an isomorphism class of supersingular curves, together with an embedding $\iota \colon K \to \mathrm{End}^0(E)$ of a quadratic field into the endomorphism algebra of the curve, called a *K-orientation*. Edges are again ℓ-isogenies. (For the precise definitions, see Section 2.)

Upon forgetting orientations, each of these volcanoes covers the supersingular ℓ-isogeny graph. In particular, each rim maps to a closed walk in the ℓ-isogeny graph.

We concern ourselves with a particular type of closed walk in $\mathcal{G}_\ell$: one with no backtracking and which is not a repeat of a smaller closed walk. We keep its direction (traversed 'forward' or 'backward'), but forget any choice of basepoint, and call this an *isogeny cycle* (Definition 3.1). The main result of this paper is that each isogeny cycle is obtained, by forgetting orientations, from exactly one oriented volcano rim (up to conjugation of the rim; see Section 3.1 for definitions).

Theorem 1 (See precise version in Theorem 3.2). *Fix distinct primes $\ell < p$. Let $r > 2$. The isogeny cycles of length r in $\mathcal{G}_\ell$ are in bijection with the rims of length r of the union of all oriented supersingular ℓ-isogeny volcanoes over $\overline{\mathbb{F}}_p$, up to conjugation of the orientations.*

The exact statement is discussed in Section 3.1, and a detailed example is given in Section 6. At the highest level, the bijection is given by a simple process. Given a rim, one can forget the orientation information to obtain an isogeny cycle in the supersingular graph. Conversely, given an isogeny cycle, one can compose the sequence of isogenies it comprises, to obtain an endomorphism. This endomorphism induces an orientation, placing the initial curve on a rim. The proof, however, spans around a third of the paper, in large part because of the existence of special cases and issues spawned by extra automorphisms and quadratic fields in which p ramifies. In particular, *the bijection is not canonical* due to these considerations. See also Remark 2.17, where we make some philosophical comments on the obstacles presented by extra automorphisms in the supersingular isogeny graph.

We use the bijection of Theorem 1 to count isogeny cycles of length r in $\mathcal{G}_\ell$. The essential ingredient is that the rims of K-oriented ℓ-isogeny volcanoes correspond to cycles of length r in the permutation on the oriented rim vertices arising from the action of an element of order r in the class group of an order of K.

Theorem 2 (See precise versions in Theorem 7.1, Corollary 7.3 and Corollary 7.7). *The number of isogeny cycles of length r in $\mathcal{G}_\ell$ is given by*

$$(1) \qquad \frac{1}{r} \sum_{\mathcal{O} \in \mathcal{I}_r} \epsilon_{\mathcal{O},\ell} h_{\mathcal{O}},$$

where $\mathcal{I}_r$ is a set of quadratic orders satisfying certain splitting properties for ℓ and p specified in Section 5, and $\epsilon_{\mathcal{O},r}$ is defined in Theorem 5.1, and satisfies $1 \le \epsilon_{\mathcal{O},r} \le 2$, with $\epsilon_{\mathcal{O},r} = 2$ whenever p is unramified in $\mathcal{O}$. The quantity (1) is bounded above by

$$\frac{2\pi e^{0.578} \log(4\ell)}{3} \, \ell^r (\log r + c) + O(\ell^{3r/4} \log r)$$

as $r \to \infty$, where the constant c is explicit, and the implied constant in the big O notation can be made explicit using Corollary 7.7. For $p \equiv 1 \pmod{12}$, the quantity (1) asymptotically approaches $\frac{\ell^r}{2r}$ as $r \to \infty$.

The asymptotic for the number of isogeny cycles is as expected for Ramanujan graphs (see Section 7.1) and we give a standard but self-contained proof for expander graphs. However, the class number formula and the explicit upper bound are consequences of the main bijection between rims and isogeny cycles. Figure 7 and Section 6 contain experimental data confirming our results.

Our strategy for counting isogeny cycles leads immediately to an algorithm to list all isogeny cycles of length r in $\mathcal{G}_\ell$. In particular, the discriminants of elements of $\mathcal{I}_r$ are bounded by ℓ^r. The j-invariants of the isogeny cycles are roots of the Hilbert class polynomials of the orders $\mathcal{O} \in \mathcal{I}_r$. The cycle in $\mathcal{G}_\ell$ can be determined by use of the modular polynomial Φ_ℓ. This generalizes a simple congruence condition check on p used in [15, Section 5.3.4] to rule out the existence of small cycles for the Charles-Goren-Lauter hash function. The exact value of $\epsilon_{\mathcal{O},r}$ in the ramified cases depends upon the theory of the ring class field of K, and is given in terms of the class number and genus number (Section 4).

In the course of the proof, we need to clarify a point about the number of orientations of a curve. When p does not split in $\mathcal{O}$, Onuki [42] considered the set $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ of pairs (E, ι) of a supersingular isogeny graph together with a primitive $\mathcal{O}$-orientation (meaning a K-orientation ι such that $\iota(K) \cap \mathrm{End}(E) = \mathcal{O}$), and the set $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ of pairs $(E, \iota) \in \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ obtained by reducing an elliptic curve over $\mathbb{C}$ with normalized complex multiplication by $\mathcal{O}$ (where normalized means that $[\alpha]^*\omega = \alpha\omega$ for the invariant differential ω; see Section 2.3.) Onuki showed that these two sets are not always equal; sometimes $\#\,\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} = 2\#\rho(\mathcal{E}\ell\ell(\mathcal{O}))$. This raises the natural question of when this case occurs. In this paper, we answer that it occurs if and only if p is inert in K. (In the ramified case, $\#\,\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} = \#\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ exactly.) We prove this using the ring class field and the action of conjugation (Theorem 4.4), but note that the result is also proved using Deuring lifting in [25, Lemma 3.2].

We also explicitly find the in- and out-degrees of the vertices of all possible K-oriented ℓ-isogeny graphs. The K-oriented graph is $(\ell+1)$-regular except at vertices corresponding to curves with extra K-automorphisms (i.e. automorphisms which preserve the orientation), where the out-degree may be slightly less than $\ell + 1$. This is made explicit in Remark 2.7 and Proposition 2.11. In connection with this, we provide a new direct proof of the volcano structure of the graph, and the count of ascending, descending and horizontal edges even in exceptional cases, in Proposition 2.15.

The motivation for this paper is prior work by the same authors, summarized in Section 8, in which the oriented volcano structure is used to navigate the supersingular ℓ-isogeny graph [5]. In particular, given an initial vertex in the supersingular ℓ-isogeny graph, and given an endomorphism of that curve, one asks to navigate to a target curve (typically $j = 1728$). The endomorphism provides an orientation of the initial curve, and hence a location in an oriented supersingular ℓ-isogeny volcano. Then, using methods of determining ascending and descending directions, one can navigate the oriented volcano. The target curve must be given an orientation on the same volcano (algorithms are provided in [5]). From both curves, one ascends to the rim and hopes to connect the paths. Having found a path on an oriented volcano, one takes the image in the original supersingular ℓ-isogeny graph.

The algorithms provided in [5] explicitly relied on heuristics about oriented volcanoes. They also implicitly hinted at the sort of bijection given in our first theorem above. In Section 9 of the present paper, we revisit the explicit heuristics of [5] and provide some partial results. In particular, we do the following.

(1) We give an explicit bound on the distance from uniform of the distribution of endpoints of non-backtracking random walks in $\mathcal{G}_\ell$ from an initial set of vertices when $p \equiv 1 \pmod{12}$ (Proposition 9.1).
(2) We estimate the expected distance of any fixed vertex to a set of j-invariants of a given size (Corollary 9.2).
(3) We fix an oriented volcano and bound the depth at which any j-invariant first occurs (Section 9.2).
(4) We fix a j-invariant and consider all the possible orientations by a given field, and the distribution of this collection of oriented curves on the various oriented volcanoes (Section 9.3). We are able to prove a result that differs slightly from the heuristic used in [5].

As an additional minor point of interest, Lemma 2.9 shows that the kernel of an endomorphism is fixed by an extra automorphism if and only if its field of definition lies in the field of definition of the automorphism. As a corollary, we can classify the behaviour of loops at $j = 0$ and $j = 1728$ in Corollary 2.13; this is a new method of proof, and slight generalization, of a known result [1, 36, 44].

The topic of closed walks[1] in the supersingular ℓ-isogeny graph was first studied in [15], where it was observed that closed walks give endomorphisms, and so the splitting behaviour of p in various extensions controlled the existence of small cycles. The endomorphisms created by closed walks were studied in the thesis of Kohel [33], and further in [7], with the goal of understanding when such cycles are independent and generate an endomorphism ring. These sources also discuss the computation of the trace of a cycle. The neighbourhoods of 0 and 1728 are particularly intricate, and their loops and cycles are discussed in [1, 35, 36, 44, 58]. Cycles in the SIKE graph which pass through the secret key are discussed in [43].

Further back in the literature, Gross [28] considered counting endomorphisms of a fixed degree by computing the traces of matrices derived from looking at certain modular forms of weight 2 for the group $\Gamma_0(N)$. Thus Gross counts closed walks of a different sort, and in particular allows backtracking (which we do not). As the degree increases, the proportion of the count due to backtracking increases as well. We provide a comparison at the end of our extended example in Section 6.

The study of K-orientations is in some sense as old as the study of optimal embeddings in quaternion algebras, and has naturally appeared in the study of isogeny graphs, for example in [37]. In the context of isogeny-based cryptography, the graph $\mathcal{G}_{K,\ell}$ was studied by Colò and Kohel in order to propose OSIDH, a new key exchange protocol [17] (see also [14]). Onuki considered the difference between $\mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$ and $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ [42]. In cryptography, the class group action on oriented curves has been used constructively [16], and the hardness of the group action problem has been considered both classically and quantumly [12] [56]; a new variation called the $\mathcal{O}$-uber isogeny problem appears in [20] (see also [14]). The volcano structure provides some algorithmic approaches to hard problems, which includes the recent work of the present authors [5] and related work [56]. Recent work also asks about the number of orientable curves [34]; see also [37]. Finally, specific orientations underlie the line of work in [23] and [20].

Outline of the paper. Section 2 contains the background on oriented and unoriented supersingular ℓ-isogeny graphs. In Section 3, we state and prove Theorem 1, and this section is the heart of the paper. Section 4 settles the question of when $\rho(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$ and discusses the action of conjugation on $\mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$. Section 5 specifies exactly the number of rims associated to a given quadratic order, in terms of class and genus numbers. Section 6 provides an extended example of the main bijection of Theorem 1. Section 7 counts the number of isogeny cycles in $\mathcal{G}_\ell$ asymptotically, as well as exactly in terms of class numbers, and gives an explicit upper bound. Section 8 gives an overview of our previous work and how it motivates the work in this paper, as well as the heuristics it depends upon. Section 9 partially addresses some of these heuristic questions by use of methods of expander graphs.

---

[1]often simply called cycles in the isogeny literature; we use the term *closed walk* to more accurately follow the graph theory literature

2. MATHEMATICAL PRELIMINARIES ON ORIENTED SUPERSINGULAR ℓ-ISOGENY GRAPHS AND THEIR VOLCANOES

2.1. The supersingular ℓ-isogeny graph $\mathcal{G}_\ell$. Let $\ell < p$ be distinct primes. We let $\mathcal{G}_\ell$ denote the supersingular ℓ-isogeny graph over $\overline{\mathbb{F}}_p$, defined as the directed graph whose vertices are isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$, and whose directed edges are isogenies of degree ℓ up to post-composition by an automorphism. The number $\#\mathcal{G}_\ell$ of supersingular j-invariants over $\overline{\mathbb{F}}_p$ (see [50]) is

$$(2) \qquad \#\mathcal{G}_\ell = \left\lfloor \frac{p}{12} \right\rfloor + \begin{cases} 0 & \text{if } p \equiv 1 \pmod{12}, \\ 1 & \text{if } p \equiv 5 \text{ or } 7 \pmod{12}, \\ 2 & \text{if } p \equiv 11 \pmod{12}. \end{cases}$$

Supersingular elliptic curves with j-invariants equal to 0 and 1728 have extra automorphisms, beyond the usual automorphisms $[\pm 1]$. Let $E_0 : y^2 = x^3 - 1$ and $E_{1728} : y^2 = x^3 - x$ with $j(E_0) = 0, j(E_{1728}) = 1728$. The curve $E_0/\overline{\mathbb{F}}_p$ is supersingular for any $p \equiv 2 \pmod 3$, and $E_{1728}/\overline{\mathbb{F}}_p$ is supersingular for any $p \equiv 3 \pmod 4$. For the sake of this discussion, we assume these curves are supersingular. We provide their explicit endomorphism rings below, referenced from [41, 51].

The automorphism group for E_{1728} is $\operatorname{Aut}(E_{1728}) = \{[\pm 1], [\pm i]\}$, where $[i](x, y) := (-x, iy)$ for $i \in \mathbb{F}_{p^2}$ with $i^2 = -1$. Write π_p for the p-power Frobenius endomorphism: $\pi_p(x, y) := (x^p, y^p)$. In $B_{p,\infty} := \operatorname{End}(E_{1728}) \otimes_{\mathbb{Z}} \mathbb{Q}$, we have:

$$(3) \qquad \operatorname{End}(E_{1728}) = \left\langle 1, [i], \frac{1 + \pi_p}{2}, \frac{[i] + [i]\pi_p}{2} \right\rangle.$$

The automorphism group for E_0 is $\operatorname{Aut}(E_0) = \{[\pm 1], [\pm \zeta_3], [\pm \zeta_3^2]\}$, where $[\zeta_3](x, y) := (\zeta_3 x, y)$ with ζ_3 a root of $T^2 + T + 1$ in $\mathbb{F}_{p^2}$. We let $[\sqrt{-3}] = 2[\zeta_3] + 1$. In $B_{p,\infty} := \operatorname{End}(E_0) \otimes_{\mathbb{Z}} \mathbb{Q}$, we have

$$(4) \qquad \operatorname{End}(E_0) = \left\langle 1, \frac{-1 + [\sqrt{-3}]}{2}, \pi_p, \frac{3 + [\sqrt{-3}] + 3\pi_p + [\sqrt{-3}]\pi_p}{6} \right\rangle.$$

Lemma 2.1. *Let $E \in \{E_0, E_{1728}\}$ and let $\varphi \in \operatorname{Aut}(E)$ such that $\varphi \neq \pm[1]$. If $\overline{\varphi}\theta = \theta\varphi$ for $\theta \in \operatorname{End}(E)$, then $\deg \theta \geq p$.*

Proof. Let $\eta = [i]$ or $\eta = [\sqrt{-3}]$ depending whether $E = E_{1728}$ or $E = E_0$, respectively. The condition $\overline{\varphi}\theta = \theta\varphi$ implies that θ must have the form $\theta = s\pi_p + t(\eta\pi_p)$ where $s, t \in \mathbb{Q}$.

For $E = E_{1728}$, using the $\mathbb{Z}$-basis for $\operatorname{End}(E_{1728})$ in (3), we have $\theta = m\pi_p + n([i]\pi_p)$ where $m, n \in \mathbb{Z}$. Thus $\deg \theta = p(m^2 + n^2) \geq p$.

For $E = E_0$, using the $\mathbb{Z}$-basis for $\operatorname{End}(E_0)$ in (4), we have $\theta = \frac{m}{2}\pi_p - \frac{n}{2}(\eta\pi_p)$ where $m, n \in \mathbb{Z}$ are of the same parity. Thus $\deg \theta = \frac{p}{4}(m^2 + 3n^2)$. Since $m^2 + 3n^2 \equiv 0 \pmod 4$, it follows that $\deg \theta \geq p$. $\qquad\square$

The graph $\mathcal{G}_\ell$ has regular out-degree $\ell + 1$ (the number of subgroups of $E[\ell]$ of size ℓ) [46]. It is very nearly an undirected graph, obtained by identifying each ℓ-isogeny with its dual. The exception is that since isogenies are identified under post-composition with an automorphism, at curves with extra automorphisms, this identification may fail (see Figure 3). Nevertheless, this issue occurs for edges adjacent to at most two vertices ($j = 0, 1728$), so one often considers it an $\ell + 1$ regular undirected graph for the purposes of spectral graph theory, without affecting the validity of the conclusions. Modulo this slight lie, or when $j = 0$ and $j = 1728$ are ordinary and hence absent from the graph (they are absent if and only if $p \equiv 1 \pmod{12}$), $\mathcal{G}_\ell$ is an $(\ell + 1)$-regular Ramanujan graph. In particular, if its adjacency matrix has eigenvalues $\ell + 1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$, then

$$(5) \qquad \max_{i > 1} |\lambda_i| \leq 2\sqrt{\ell}.$$

2.2. The oriented supersingular ℓ-isogeny graph $\mathcal{G}_{K,\ell}$. Let $\ell < p$ be distinct primes. Fix a supersingular elliptic curve $E/\overline{\mathbb{F}}_p$. Then $\operatorname{End}^0(E) \cong B_{p,\infty}$, where $B_{p,\infty}$ denotes the definite quaternion algebra ramified precisely at p and ∞ (unique up to isomorphism). Let K be an imaginary quadratic field in which p does not split. This condition on p gives the existence of an embedding of K into $B_{p,\infty}$. Let $\mathcal{O}$ denote an order in K and $h_\mathcal{O}$ the class number of this order. Let $\Delta_\mathcal{O}$ denote its discriminant. An order $\mathcal{O}$ of K is ℓ-*fundamental* if its conductor is relatively prime to ℓ.

Definition 2.2 (K-orientation)**.** A K-*orientation* of E is an embedding $\iota\colon K \to \mathrm{End}^0(E)$.

Definition 2.3 ($\mathcal{O}$-orientation)**.** A K-orientation $\iota\colon K \to \mathrm{End}^0(E)$ is an $\mathcal{O}$-*orientation* if $\iota(\mathcal{O}) \subseteq \mathrm{End}(E)$. Furthermore, we say that an $\mathcal{O}$-orientation is $\mathcal{O}$-*primitive* if $\iota(\mathcal{O}) = \mathrm{End}(E) \cap \iota(K)$.

Every K-orientation $\iota\colon K \to \mathrm{End}^0(E)$ is $\mathcal{O}$-primitive for a unique order $\mathcal{O}$ of K.

In the literature, for example [37], one will find the closely related term *optimal embedding*, in the context of a maximal order $\mathfrak{O}$ of a quaternion algebra A, for an embedding $\iota\colon K \to A$ such that $\iota(K) \cap \mathfrak{O} = \mathcal{O}$. An optimal embedding corresponds to a primitive orientation via the Deuring correspondence. For the theory of optimal embeddings, see [55, Chapter 30].

Definition 2.4 (K-oriented elliptic curve)**.** A K-*oriented elliptic curve* is a pair (E, ι), where E is an elliptic curve and ι is a K-orientation of E.

Isogenies of elliptic curves induce isogenies of oriented elliptic curves in the following sense: if $\phi\colon E \to E'$ is an isogeny and $\iota\colon K \to \mathrm{End}(E)$ is a K-orientation of E, then we define the unique induced orientation $\phi_*\iota$ on E' for any $\alpha \in K$ as follows:

$$(6) \qquad\qquad (\phi_*\iota)(\alpha) := \frac{1}{[\deg \phi]} \phi \circ \iota(\alpha) \circ \widehat{\phi}.$$

Definition 2.5 (K-oriented isogeny)**.** If $\phi_*\iota = \iota'$, then we say that $\phi\colon (E, \iota) \to (E', \iota')$ is an isogeny of K-oriented curves, also known as a K-*oriented isogeny*.

Two K-oriented curves $(E, \iota), (E', \iota')$ are K-*isomorphic* if there exists an isomorphism $\eta\colon E \to E'$ such that $\eta_*\iota = \iota'$. Two K-oriented isogenies $\phi\colon (E_1, \iota_1) \to (E_2, \iota_2)$ and $\phi'\colon (E_1', \iota_1') \to (E_2', \iota_2')$ are considered *equivalent* if there are K-oriented isomorphisms η and η' so that the following diagram commutes:

$$(7) \qquad\qquad
\begin{array}{ccc}
(E_1, \iota_1) & \xrightarrow{\phi} & (E_2, \iota_2) \\
\eta \downarrow & & \downarrow \eta' \\
(E_1', \iota_1') & \xrightarrow{\phi'} & (E_2', \iota_2')
\end{array}$$

We will frequently abuse notation by writing a representative (E, ι) or ϕ in place of its equivalence class.

An automorphism φ of E is a K-*automorphism* of (E, ι) if $\varphi_*\iota = \iota$. This occurs if and only if $\varphi \in \iota(K)$.

Definition 2.6 ($\mathcal{G}_{K,\ell}$)**.** The ℓ-*isogeny graph* $\mathcal{G}_{K,\ell}$ of K-oriented supersingular elliptic curves over $\overline{\mathbb{F}}_p$ is the directed graph whose vertices (E, ι) are K-isomorphism classes of K-oriented supersingular elliptic curves over $\overline{\mathbb{F}}_p$ and whose edges are equivalence classes of K-isogenies of degree ℓ between the given oriented curves.

Remark 2.7. The only vertices of $\mathcal{G}_{K,\ell}$ with automorphism group larger than $\{[\pm 1]\}$ are (E, ι) where either (i) E has j-invariant 1728 and $\iota(i) \in \mathrm{End}(E)$, or (ii) E has j-invariant 0 and $\iota(\zeta_3) \in \mathrm{End}(E)$.

Proposition 2.8. *The directed edges of $\mathcal{G}_{K,\ell}$ can be put into equivalence classes by making an isogeny equivalent to its dual isogeny. These classes are of size 1 or 2. The only self-dual edges are loops. By identifying isogenies with their duals, this graph can be drawn undirected.*

Proof. With reference to (7), if ϕ and ϕ' are equivalent, then their duals are equivalent. Therefore an equivalence class consists of one directed edge and its dual, or a single self-dual edge. $\square$

Although the number of ℓ-isogenies originating at a supersingular curve is always $\ell + 1$, some of these may be equivalent, so that in the graph $\mathcal{G}_{K,\ell}$, the out-degree may be strictly smaller. As usual, the culprits are the extra automorphisms, as in Remark 2.7. Before we state and prove the exact out-degree, we need a lemma.

Lemma 2.9. *Let E be a curve with extra automorphisms. Given a non-zero isogeny $\phi\colon E \to E'$, its kernel $\ker \phi$ is fixed by $\mathrm{Aut}(E) \setminus \{[\pm 1]\}$ if and only if $E \cong E'$. If we further assume that $\deg \phi < p$ then $\ker \phi$ is fixed by $\mathrm{Aut}(E) \setminus \{[\pm 1]\}$ if and only if $E \cong E'$ and ϕ (up to post-composition with an isomorphism $E' \to E$) is an element in $\mathbb{Q}(\mathrm{Aut}(E)) \cap \mathrm{End}(E)$.*

Proof. Let φ be a non-integer automorphism of E. First, suppose $\ker \phi = \varphi \ker \phi$. Then $\ker \phi = \ker \phi\varphi$, so there exists a unique $\eta \in \mathrm{Aut}(E')$ such that $\phi\varphi = \eta\phi$ by [50, III.4.11]. Then, working in $\mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q} \cong \mathrm{End}(E') \otimes_{\mathbb{Z}} \mathbb{Q}$, the element $\phi\varphi \frac{\hat{\phi}}{[\deg(\phi)]} = \eta \in \mathrm{Aut}(E') \subset \mathrm{End}(E')$ has the same minimal polynomial as $\varphi \in \mathrm{Aut}(E) \subset \mathrm{End}(E)$. The existence of non-integer elements in $\mathrm{Aut}(E)$ and $\mathrm{Aut}(E')$ with the same minimal polynomial gives $\mathrm{Aut}(E) \cong \mathrm{Aut}(E')$. Since $\mathrm{Aut}(E) \cong \mathrm{Aut}(E') \supsetneq \{[\pm 1]\}$, the automorphism group determines the isomorphism class of the curve and we have $E \cong E'$.

Next, suppose $E \cong E'$ via an isomorphism $\lambda \colon E' \to E$. In particular, $\mathrm{Aut}(E) \cong \mathrm{Aut}(E') \supsetneq \{[\pm 1]\}$. Then, $\lambda\phi \in \mathrm{End}(E)$ and $\ker(\lambda\phi) = \ker\phi$. By abuse of notation, we replace $\lambda\phi$ with ϕ and may henceforth assume $\phi \in \mathrm{End}(E)$ in this case. The endomorphism ϕ has the same minimal polynomial as $\varphi^{-1}\phi\varphi$, and so up to some automorphism $\eta \in \mathrm{Aut}(E)$, we have $\varphi^{-1}\phi\varphi = \eta\phi$, and thus $\ker\phi = \ker\phi\varphi$.

Now, assume $\deg\phi < p$ and suppose $E \cong E'$ and $\ker\phi$ is fixed by a generator φ of $\mathrm{Aut}(E)$. Then $\phi = \varphi^k\phi\varphi$ for some integer $0 \le k < |\mathrm{Aut}(E)|$. By cancellation, we know $\phi \neq \pm\phi\varphi$ since $\varphi \neq [\pm 1]$, hence $\overline{\varphi}\phi = \phi\varphi$ or else $\varphi\phi = \phi\varphi$. For $\varphi = [\zeta_6]$ (a primitive sixth root of unity), we have $k \in \{1, 2, 4, 5\}$. If $k \in \{1, 5\}$, then either $\overline{[\zeta_6]}\phi = \phi[\zeta_6]$ or $[\zeta_6]\phi = \phi[\zeta_6]$. Otherwise, if $k \in \{2, 4\}$, then observe that $\phi = [\zeta_6]^k\phi[\zeta_6] = [\zeta_6]^k([\zeta_6]^k\phi[\zeta_6])[\zeta_6] = [\zeta_6]^{2k}\phi[\zeta_6]^2 = [\zeta_3]^k\phi[\zeta_3]$, so $\overline{[\zeta_3]}\phi = \phi[\zeta_3]$ or $[\zeta_3]\phi = \phi[\zeta_3]$. Since $\deg\phi < p$, Lemma 2.1 rules out the case that $\overline{\varphi}\phi = \phi\varphi$. We are left to conclude that ϕ commutes with an element φ of $\mathrm{Aut}(E) \setminus \{[\pm 1]\}$, which implies that $\phi \in \mathbb{Q}(\varphi) \subseteq \mathrm{End}^0(E)$.

On the other hand, if $\ker\phi \neq \ker\phi\varphi$, then $\phi \neq \pm\varphi^k\phi\varphi$ for any k, so ϕ does not commute with φ, and hence $\phi \notin \mathbb{Q}(\mathrm{Aut}(E))$. $\square$

Definition 2.10 (Reduced Automorphism Group)**.** Let $\mathrm{Aut}'(E) := \mathrm{Aut}(E)/[\pm 1]$ denote the reduced automorphism group of the elliptic curve E.

Proposition 2.11. *Let (E, ι) be an oriented curve in $\mathcal{G}_{K,\ell}$ where ι is an $\mathcal{O}$-orientation. Then (E, ι) has out-degree $\ell + 1$, except at the oriented curves with extra K-automorphisms (those of Remark 2.7). In that exceptional case, the out-degree is $(\ell + 1 - s_\ell)/|\mathrm{Aut}'(E)| + s_\ell$ where s_ℓ is the number of $\phi \in \mathrm{End}(E) \cap \iota(K)$ of degree ℓ up to post-composition by $\mathrm{Aut}(E)$. In particular, the out-degree at (E, ι), when (E, ι) has extra K-automorphisms, is*

$$\begin{cases} (\ell - 1)/2 + 2 & \mathcal{O} = \mathbb{Z}[i] \text{ and } \ell \text{ splits in } \mathcal{O} \ (s_\ell = 2) \\ (\ell + 1)/2 & \mathcal{O} = \mathbb{Z}[i] \text{ and } \ell \text{ is inert in } \mathcal{O} \ (s_\ell = 0) \\ 2 & \mathcal{O} = \mathbb{Z}[i] \text{ and } \ell = 2 \ (s_\ell = 1) \\ 2(\ell - 1)/3 + 2 & \mathcal{O} = \mathbb{Z}[\zeta_3] \text{ and } \ell \text{ splits in } \mathcal{O} \ (s_\ell = 2) \\ 2(\ell + 1)/3 & \mathcal{O} = \mathbb{Z}[\zeta_3] \text{ and } \ell \text{ is inert in } \mathcal{O} \ (s_\ell = 0) \\ 3 & \mathcal{O} = \mathbb{Z}[\zeta_3] \text{ and } \ell = 3 \ (s_\ell = 1) \end{cases}$$

Proof. With reference to diagram (7), by taking $\eta \in \mathrm{Aut}(E)$, the two equivalent K-oriented isogenies ϕ and ϕ' from (E, ι) must satisfy $\iota' = \eta_*\iota = \iota$ and $\eta\ker\phi = \ker\phi'$. The first equality implies $\eta \in \iota(K)$. Hence (E, ι) has out-degree $(\ell + 1)$ if $\mathrm{Aut}(E) \cap \iota(K) = [\pm 1]$. If $\mathrm{Aut}(E) \cap \iota(K) \supsetneq [\pm 1]$, the out-degree may be less than $(\ell + 1)$, and we investigate this case below.

If $\mathrm{Aut}(E) \cap \iota(K)$ contains non-trivial automorphisms, then it contains all of them, and $\mathrm{End}(E) \cap \iota(K) = \iota(\mathcal{O}_K) \supseteq \mathrm{Aut}(E)$. The $\ell + 1$ outgoing ℓ-isogenies fall into equivalence classes according to whether their kernels are related by an extra automorphism. The kernel of an outgoing ℓ-isogeny ϕ from E is fixed by $\eta \in \mathrm{Aut}(E)$ if and only if $\phi \in \mathrm{End}(E) \cap \iota(K)$ (as in the proof of Lemma 2.9, we may assume that $\phi \in \mathrm{End}(E)$ by post-composing with an isomorphism; we have assumed throughout the paper that $\ell < p$). Let S_ℓ be the set of equivalence classes of elements in $\mathrm{End}(E) \cap \iota(K)$ of degree ℓ up to pre-composition by $\mathrm{Aut}(E)$, and put $s_\ell = \#_\ell$. Note that $\mathrm{End}(E) \cap \iota(K) = \iota(\mathcal{O}_K)$, so $s_\ell = 0, 1$ or 2, depending on whether ℓ is inert, ramified or split in $\mathcal{O}_K$. By the above, a degree ℓ-isogeny ϕ is in S_ℓ if and only if it is fixed by all $\eta \in \mathrm{Aut}(E)$. Otherwise its kernel has an orbit of size $|\mathrm{Aut}'(E)|$. Thus the number of out-isogenies is $(\ell + 1 - s_\ell)/|\mathrm{Aut}'(E)| + s_\ell$. $\square$

Remark 2.12. As an immediate corollary, we obtain a new method of proof and generalization of a known result about the number of loops at a vertex [1, 36, 44]. Specifically, the previous works concerned only large p (the last two sentences of the next corollary).

Corollary 2.13. *Let $j \in \{0, 1728\}$ correspond to a supersingular curve over $\overline{\mathbb{F}}_p$. Let $K = \mathbb{Q}(\zeta_3), \mathbb{Q}(i)$ for $j = 0$, 1728 respectively. Let n represent the number of loops at the vertex of j-invariant j, in $\mathcal{G}_\ell$. Then*

$$n = 1 + \left(\frac{\Delta_K}{\ell}\right) + k|\operatorname{Aut}'(E)|,$$

for some $k \in \mathbb{Z}$. If $4\ell < p$, then the number of loops at $j = 1728$ is exactly $1 + \left(\frac{\Delta_K}{\ell}\right)$. If $3\ell < p$, then the number of loops at $j = 0$ is exactly $1 + \left(\frac{\Delta_K}{\ell}\right)$.

Proof. Let $\mathcal{O}$ be $\mathbb{Z}[\zeta_3]$ or $\mathbb{Z}[i]$ when $j = 0$ or 1728, respectively. Whether ℓ splits or ramifies or is inert determines the number of elements of $\iota(\mathcal{O})$ of degree ℓ, up to automorphisms: two, one or zero, respectively. These elements are all loops on the vertex. All other loops fall into groups of size $|\operatorname{Aut}'(E)|$, as in the proof of Proposition 2.11. For the last two sentences of the statement, we revisit the endomorphism rings of the vertex (Section 2.1), to see that the smallest elements not in $\iota(\mathcal{O})$ have norm $p/4$ and $p/3$ respectively. $\square$

It is possible to compute k explicitly for any given value of ℓ in terms of p, by reference to the endomorphism rings as described in the proof.

The graph $\mathcal{G}_{K,\ell}$ consists of connected components, each with a volcano structure as pictured in Figure 1. The rim vertices of any given volcano are K-orientations of supersingular elliptic curves which are $\mathcal{O}$-primitive for some fixed ℓ-fundamental order $\mathcal{O}$ of K. The oriented supersingular elliptic curves which appear at depths below the rim are primitively oriented by orders whose index in $\mathcal{O}$ is given by increasing powers of ℓ: at depth one from the rim, the curves are primitively $(\mathbb{Z} + \ell\mathcal{O})$-oriented, and in general, at depth k, the curves are primitively $(\mathbb{Z} + \ell^k\mathcal{O})$-oriented. The edges of the volcano can thus be classified according to whether they increase, decrease or preserve the depth as we move from domain to codomain.

Definition 2.14 (Horizontal, descending and ascending isogenies)**.** Let $\phi : (E, \iota) \to (E', \iota')$ be a K-oriented degree ℓ isogeny. If ι is a primitive $\mathcal{O}$-orientation and ι' is a primitive $\mathcal{O}'$-orientation, then ϕ is one of the following three types of isogeny:

    (1) ϕ is *horizontal* if $\mathcal{O} = \mathcal{O}'$,
    (2) ϕ is *descending* if $\mathcal{O} \supsetneq \mathcal{O}'$,
    (3) ϕ is *ascending* if $\mathcal{O} \subsetneq \mathcal{O}'$.

The fact that each component has a volcano structure is the content of the following proposition; versions of this appear in various places in the literature [17, 42].

Proposition 2.15. *Let (E, ι) be such that ι is $\mathcal{O}$-primitive. If $\mathcal{O}$ is ℓ-fundamental, then (E, ι) has no ascending ℓ-isogeny and $\left(\frac{\Delta_\mathcal{O}}{\ell}\right) + 1$ horizontal ℓ-isogenies. Otherwise, (E, ι) has one ascending ℓ-isogeny and no horizontal isogenies. All other isogenies from (E, ι) are descending.*

Proof. This follows from [5, Proposition 4.8]. Write $\mathcal{O} = \mathbb{Z}[\theta]$ satisfying the hypotheses of that proposition. If $\mathcal{O}$ is ℓ-fundamental, then the norm of θ is not divisible by ℓ. Then θ acts to permute $E[\ell]$ and therefore has two non-zero eigenvalues. By [5, Proposition 4.8], there are two horizontal isogenies if and only if these eigenvalues are distinct and defined over $\mathbb{F}_\ell$, i.e. ℓ splits. If they are not distinct but defined over $\mathbb{F}_\ell$, there is one horizontal isogeny, i.e. ℓ ramifies. Otherwise there are no horizontal isogenies. If $\mathcal{O}$ is not ℓ-fundamental, then the norm of θ is divisible by ℓ and there is a zero eigenvalue (but only one, since ℓ does not divide θ). By [5, Proposition 4.8], this gives exactly one ascending isogeny. All other isogenies are descending. $\square$

Let $\mathcal{O}$ be an order in K. Every primitive $\mathcal{O}$-oriented curve (E, ι) is on a volcano in $\mathcal{G}_{K,\ell}$. If $\mathcal{O}$ is ℓ-fundamental, then (E, ι) is on the *rim* of this volcano. If ℓ divides the conductor of $\mathcal{O}$, then the power of ℓ dividing the discriminant of $\mathcal{O}$ is referred to as the *depth* of the curve on this volcano.

Proposition 2.16. *A connected component of $\mathcal{G}_{K,\ell}$, when identifying ℓ-isogenies with their duals, has at most one cycle, passing through the rim vertices.*

Proof. The fact that there is at most one ascending isogeny from any vertex below the rim implies that there are no cycles below the rim. Restricting to the rim, we have a 2-regular graph, hence a union of cycle graphs. But since the graph is connected by assumption, it consists of at most one cycle. $\square$

Therefore, we also use the term *rim* to refer to the undirected cycle graph that one obtains as a subgraph of $\mathcal{G}_{K,\ell}$ by restricting to the rim vertices and identifying ℓ-isogenies with their duals. We call a rim *directed* if we choose a direction of traversal (disallowing backtracking).

Remark 2.17. We pause briefly to justify the choice of equivalence on isogenies for $\mathcal{G}_{K,\ell}$ (which differs from the notion of equivalence in $\mathcal{G}_\ell$). In fact, when forming a graph whose vertices are isomorphism classes of any type of object, the graph is well-defined if we choose the edges to be maps up to pre- and post-composition by isomorphism, as in (7). In other words, it is reasonable to require of our graph definition that, under any two sets of representatives for the vertices, and any set of object-wise isomorphisms between those sets of representatives, the resulting graphs should be canonically isomorphic. This requirement, when unravelled, is essentially the requirement than the maps be taken up to pre- and post-composition. Historically, the supersingular isogeny graph is not defined in this robust way, instead identifying isogenies as equivalent merely up to post-composition. This allows for the computational convenience of identifying kernels with isogenies, and allows for an identification between non-backtracking walks and cyclic isogenies. However, it results in certain non-canonical behaviours. These behaviours can be observed only in proximity to the curves with extra automorphisms (having j-invariant $j = 0$ and $j = 1728$), particularly the phenomenon described in the second paragraph of Section 3.2. The historical definition of the supersingular isogeny graph resulted in substantial difficulties arising from extra automorphisms in proving our main Theorem 3.2.

2.3. **The sets $\mathrm{SS}_{\mathcal{O}}^{\mathbf{pr}}$ and $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ and the class group action.** Having stratified each volcano into depths where the orientations are all primitive with respect to a fixed order, it is natural to consider the complete set of curves with such orientations.

Definition 2.18 ($\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$)**.** Let $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ denote the set of primitively $\mathcal{O}$-oriented isomorphism classes of supersingular elliptic curves, up to K-isomorphism.

Here we recall the following conditions for the set $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ to be non-empty.

Proposition 2.19 ([42, Proposition 3.2])**.** *The set $\mathrm{SS}_{\mathcal{O}}^{pr}$ is not empty if and only if p does not split in K and does not divide the conductor of $\mathcal{O}$.*

If $\mathcal{O}$ is ℓ-fundamental, then $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ is a union of volcano rims in $\mathcal{G}_{K,\ell}$.

For the proof of our bijection in Section 3.3, we invoke the Deuring lifting theorem as outlined in [42, Section 3.2]. For this, we will need a number of definitions and lemmas about CM elliptic curves over number fields.

There is a number field L' (an extension of the ring class field L of $\mathcal{O}$) and a prime ideal $\mathfrak{p}$ above p in $\mathcal{O}_{L'}$ such that every elliptic curve with complex multiplication by $\mathcal{O}$ has a representative over L' with good reduction at $\mathfrak{p}$ [50, Section VII.5]. Fix such a choice.

Definition 2.20 ($\mathcal{E}\ell\ell(\mathcal{O})$)**.** Let $\mathcal{E}\ell\ell(\mathcal{O})$ denote the set of isomorphism classes of elliptic curves over L' with endomorphism ring isomorphic to $\mathcal{O}$ and good reduction at $\mathfrak{p}$.

By the theory of complex multiplication, $\#\mathcal{E}\ell\ell(\mathcal{O}) = h_{\mathcal{O}}$. Write $[\cdot]_E \colon \mathcal{O} \to \mathrm{End}(E)$ for the isomorphism which is normalized so that $([\alpha]_E)^*\omega_E = \alpha\omega_E$, where ω_E denotes the invariant differential of E ([42, Section 2.3], [49, Prop. II.1.1]). There is a map given by reduction modulo $\mathfrak{p}$:

$$\rho \colon \mathcal{E}\ell\ell(\mathcal{O}) \to \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}, \quad E \mapsto (\tilde{E}, \iota_E),$$

where $\iota_E \colon K \to \mathrm{End}^0(\tilde{E})$ is determined by its restriction $\iota_E \colon \mathcal{O} \to \mathrm{End}(\tilde{E})$ being the reduction modulo $\mathfrak{p}$ of the isomorphism $[\cdot]_E$ (in other words, $\iota_E(\alpha) = [\alpha]_E \pmod{\mathfrak{p}}$ for all $\alpha \in \mathcal{O}$). In particular, from the normalization, it holds that $\iota_E(\alpha)^*\omega_{\tilde{E}} = \alpha\omega_{\tilde{E}}$ by reduction modulo $\mathfrak{p}$.

The map ρ is injective on $\mathcal{E}\ell\ell(\mathcal{O})$, so that its image $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ is a subset of $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ of size $h_{\mathcal{O}}$. There is an action of the p-power Frobenius on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$, given by $\pi_p \cdot (E, \iota) \mapsto (E^{(p)}, \iota^{(p)})$, where $E^{(p)} = \pi_p(E)$ and $\iota^{(p)} = (\pi_p)_*\iota$.

Proposition 2.21 ([42, Proposition 3.3])**.** *For all $(E, \iota) \in \mathrm{SS}_{\mathcal{O}}^{pr}$, at least one of (E, ι), $(E^{(p)}, \iota^{(p)})$ belongs to $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$.*

As suggested by the above, surjectivity of ρ may indeed fail: $\#\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} \in \{h_{\mathcal{O}}, 2h_{\mathcal{O}}\}$. Theorem 4.4 describes when this failure occurs.

We define an action of ideals of $\mathcal{O}$ on oriented elliptic curves $(E, \iota) \in \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$.

Definition 2.22 (Ideal action on oriented elliptic curves)**.** Let $(E, \iota) \in \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ for some order $\mathcal{O}$ of K. Let $\mathfrak{a}$ be an integral ideal of $\mathcal{O}$ coprime to p. Define the intersection:

$$E[\iota(\mathfrak{a})] := \bigcap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha)).$$

This group defines an isogeny $\phi_{\mathfrak{a}}^{(E,\iota)} : E \to E/E[\iota(\mathfrak{a})]$. The action of $\mathfrak{a}$ on (E, ι) is defined as $\mathfrak{a} * (E, \iota) := (\phi_{\mathfrak{a}}^{(E,\iota)}(E), (\phi_{\mathfrak{a}}^{(E,\iota)})_* \iota)$.

When clear from context, we will drop the superscript and just write $\phi_{\mathfrak{a}}$. Note that to define the action of the ideal $\mathfrak{a}$, we require its norm to be coprime to p (see the proof of [42, Proposition 3.6]).

Theorem 2.23 ([42, Theorem 3.4])**.** *Assume that p does not split in K and p does not divide the conductor of $\mathcal{O}$. Then the ideal action of Definition 2.22 defines a free and transitive action of $\mathrm{Cl}(\mathcal{O})$ on $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$.*

This action is the 'same' as that of the CM theory, in the sense that $\rho(\sigma \cdot (E, \iota)) = \sigma \cdot (\rho((E, \iota)))$ for $\sigma \in \mathrm{Cl}(\mathcal{O})$ [42, Proof of Proposition 3.6].

Finally, we relate $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ to the volcanoes of $\mathcal{G}_{K,\ell}$.

Definition 2.24 ($\mathcal{O}$-cordillera)**.** A collection of volcanoes in $\mathcal{G}_{K,\ell}$. The collection of volcanoes whose rims are primitively oriented by $\mathcal{O}$ is called the $\mathcal{O}$-*cordillera*.

The set of vertices at the rim of the $\mathcal{O}$-cordillera is exactly the set $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$.

Next, we slightly strengthen results of Colò-Kohel [17] and Onuki [42] to give an action on oriented isogenies by a direct product of the class group with Frobenius.

Definition 2.25 (Action of Frobenius on oriented elliptic curves)**.** The two-element group $\langle \pi_p \rangle = \{1, \pi_p\}$ generated by the Frobenius automorphism π_p of $\mathbb{F}_{p^2}$ acts on $\mathcal{G}_{K,\ell}$ by

$$\pi_p \cdot (E, \iota) = (E^{(p)}, \iota^{(p)}), \quad \pi_p \cdot \varphi = \varphi^{(p)},$$

where $\iota^{(p)} := (\pi_p)_*(\iota)$.

For any isogeny φ, we have $\pi_p \circ \varphi(x, y) = \varphi^{(p)}(x^p, y^p) = \varphi^{(p)} \circ \pi_p(x, y)$. Hence, one has

$$\iota^{(p)}(\alpha) = (\pi_p)_*(\iota)(\alpha) = \frac{1}{p}\pi_p \circ \iota(\alpha) \circ \widehat{\pi_p} = \frac{1}{p}\iota(\alpha)^{(p)} \circ \pi_p \circ \widehat{\pi_p} = \iota(\alpha)^{(p)}.$$

Since $\varphi \mapsto \varphi^{(p)}$ gives an isomorphism $\mathrm{End}(E) \cong \mathrm{End}(E^{(p)})$, this yields an action on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ by $\langle \pi_p \rangle$. In fact, it is an action on the graph $\mathcal{G}_{K,\ell}$, i.e. it preserves adjacency.

Proposition 2.26. *The actions described in Definitions 2.22 and 2.25 above commute and hence give an action of $\mathrm{Cl}(\mathcal{O}) \times \langle \pi_p \rangle$ on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$. This action is transitive and its point stabilizers are either all trivial or all $\langle \pi_p \rangle$. In particular, $\#\,\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} \in \{h_{\mathcal{O}}, 2h_{\mathcal{O}}\}$.*

Proof. We have $\pi_p \cdot \varphi_{\mathfrak{a}} \cdot (E, \iota) = (\varphi_{\mathfrak{a}})^{(p)} \cdot \pi_p \cdot (E, \iota)$. To avoid confusion we momentarily use the more specific notation $\varphi_{\mathfrak{a}}^{(E,\iota)}$ to denote the isogeny $\varphi_{\mathfrak{a}}$ with domain (E, ι) and kernel $E[\iota(\mathfrak{a})]$. Then

$$\ker((\varphi_{\mathfrak{a}}^{(E,\iota)})^{(p)}) = \ker(\varphi_{\mathfrak{a}}^{(E,\iota)})^{(p)} = E[\iota(\mathfrak{a})]^{(p)} = \bigcap_{\theta \in \iota(\mathfrak{a})} \ker(\theta)^{(p)}$$

$$= \bigcap_{\theta \in \iota(\mathfrak{a})} \ker(\theta^{(p)}) = \bigcap_{\theta \in \iota^{(p)}(\mathfrak{a})} \ker(\theta) = E^{(p)}[\iota^{(p)}(\mathfrak{a})] = \ker((\varphi_{\mathfrak{a}}^{(E^{(p)},\iota^{(p)})})).$$

The calculation above implies that $(\varphi_{\mathfrak{a}}^{(E,\iota)})^{(p)} = \varphi_{\mathfrak{a}}^{(E^{(p)},\iota^{(p)})}$. Thus,

$$\pi_p \cdot \mathfrak{a} \cdot (E, \iota) = \mathfrak{a} \cdot \pi_p \cdot (E, \iota).$$

So the action of $\mathrm{Cl}(\mathcal{O}) \times \langle \pi_p \rangle$ on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ is well-defined.

The restriction of this action to $\mathrm{Cl}(\mathcal{O})$ acts freely and transitively on a subset of $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ which contains at least one of (E, ι) or $(E^{(p)}, \iota^{(p)})$ [42, Theorem 3.4], from which the rest of the statement follows. Transitivity implies that the stabilizers are all of the same size. $\qquad\square$

FIGURE 2. Here $p = 31, \ell = 2$ and $K = \mathbb{Q}(\sqrt{-47})$. This figure shows the set $\mathrm{SS}_{\mathcal{O}_K}$ which consists of two cycles. We denote by (j, ι_j) the vertex with j-invariant j and orientation ι_j; if the j-invariant appears twice, the second orientation is denoted by $\iota_{j,2}$. The notation $\bar{\iota}_j$ denotes the conjugate orientation. Conjugation (Section 4) maps (j, ι_j) to $(j, \bar{\iota}_j)$ and $(j, \iota_{j,2})$ to $(j, \bar{\iota}_{j,2})$. On the other hand, π_p maps $(1728, \iota_{1728})$ to $(1728, \bar{\iota}_{1728})$, $(2, \iota_2)$ to $(2, \bar{\iota}_{2,2})$, and $(4, \iota_4)$ to $(4, \bar{\iota}_{4,2})$.

Remark 2.27. One might expect to see the dihedral group, not a direct product. We will see the dihedral group coming from class field theory in Section 4, where we consider the action of conjugation on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$. But in our definitions here, Frobenius acts on both curve and orientation. In other words, $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ has an action by a direct product and by a dihedral group, both extensions of the same action of $\mathrm{Cl}(\mathcal{O})$ by $\mathbb{Z}/2\mathbb{Z}$, but these are not necessarily the same. For an example of the action of Frobenius and the action of conjugation being different, see Figure 2.

2.4. A lemma of algebraic number theory.

We make explicit a small lemma from algebraic number theory that will be used in Section 7.

Lemma 2.28. *Let K be an imaginary quadratic field and let $\mathcal{O}$ be an order in K with conductor f. Assume that ℓ is a prime that does not divide f. Then ℓ splits (ramifies, is inert, respectively) in $\mathcal{O}$ if and only if ℓ splits (ramifies, is inert, respectively) in the maximal order $\mathcal{O}_K$.*

Proof. By [18, Proposition 7.20], if $\mathfrak{a}$ is an $\mathcal{O}_K$-ideal above ℓ prime to f, then $\mathfrak{a}' = \mathfrak{a} \cap \mathcal{O}$ is an $\mathcal{O}$-ideal above ℓ with the same norm, and if $\mathfrak{a}'$ is an $\mathcal{O}$-ideal above ℓ prime to f, then $\mathfrak{a} = \mathfrak{a}'\mathcal{O}_K$ is an $\mathcal{O}_K$-ideal above ℓ with the same norm. Furthermore, these maps are inverse to one another, giving a bijection between ideals of $\mathcal{O}_K$ coprime to f, and invertible ideals of $\mathcal{O}$. Therefore the splitting, ramification and inertia properties of ℓ in $\mathcal{O}_K$ and in $\mathcal{O}$ are the same. $\square$

3. BIJECTION BETWEEN ORIENTED VOLCANO RIMS AND ISOGENY CYCLES IN THE SUPERSINGULAR ℓ-ISOGENY GRAPH

3.1. Statements of results.

In this section, we give a bijection between 'isogeny cycles' (closed walks satisfying a non-backtracking and non-repeating condition; see below) on the supersingular ℓ-isogeny graph and rims (up to conjugation) of the volcanoes in all oriented supersingular ℓ-isogeny graphs. We first give our definition of isogeny cycles.

Definition 3.1 (Isogeny cycle). An *isogeny cycle* is a closed walk, forgetting basepoint, in $\mathcal{G}_\ell$ containing no backtracking (no consecutive edges compose to multiplication-by-ℓ, Definition 3.5) which is not a power of another closed walk (i.e. not equal to another closed walk repeated more than once).

Isogeny cycles travelling the same vertices in opposite directions necessarily employ different directed edges and hence are considered different.

We need a few more items of terminology. The term *directed* refers to the notion of direction of traversal. Each rim corresponds to two directed rims, since there are two possible directions of traversal. The *conjugate* of a rim R is obtained by taking each vertex (E, ι) to $(E, \bar{\iota})$ and each edge $\phi \colon (E, \iota) \to (E', \iota')$ to $\phi \colon (E, \bar{\iota}) \to (E', \bar{\iota'})$. The conjugation of a rim is again a rim. Note that if $\phi_{\mathfrak{a}}^{(E,\iota)} \colon (E, \iota) \to (E', \iota')$, then $\phi_{\bar{\mathfrak{a}}}^{(E,\bar{\iota})} \colon (E, \bar{\iota}) \to (E', \bar{\iota'})$ since $E[\bar{\iota}(\bar{\mathfrak{a}})] = E[\iota(\mathfrak{a})]$. So in particular, the direction of a rim (E_0, E_1, etc.) is preserved by conjugation, but the ideal whose action traverses the rim in that direction is conjugated.

We now state our main theorem. For each integer $r > 2$, define two sets:

$$\mathcal{C}_r = \{ \text{ isogeny cycles in } \mathcal{G}_\ell \text{ of length } r \ \}.$$

$$\mathcal{R}_r = \{ \text{ directed rims } R \text{ of size } r \text{ in } \mathcal{G}_{K,\ell} \colon K \text{ is an imaginary quadratic field } \}.$$

Theorem 3.2. *Let $r > 2$. There is a bijection between $\mathcal{C}_r$ and $\mathcal{R}_r / \sim$, where $\sim$ denotes identifying any rim with its conjugate.*

Sketch of proof. Theorem 3.2 will be proved in Section 3.3. At the highest level, this bijection is easy to describe: given a rim $R \in \mathcal{R}_r$, forgetting orientations gives a cycle $C \in \mathcal{C}_r$. Conversely, given a cycle $C \in \mathcal{C}_r$ passing through a curve E, composing its component isogenies in order, one obtains an endomorphism $\theta \in \mathrm{End}(E)$, which induces an orientation ι_θ. Lifting E to (E, ι_θ), the isogenies of the cycle will give a rim in $\mathcal{G}_{K,\ell}$, where K is the quadratic field generated by the minimal polynomial of θ. These two maps are inverse.

However, the devil is in the details, especially the devil whose name is *extra automorphisms*. In particular, if $\mathcal{G}_\ell$ has no curves with extra automorphisms, the proof is much briefer. In what follows, we have endeavoured to clearly mark areas where we address extra automorphisms, so that at a first read, they can be skipped. In particular, if $\mathcal{G}_\ell$ has no extra automorphisms, the reader can entirely skip Section 3.4, and the word 'safe' (Definition 3.13) can be ignored. Section 3.2 should still be read, however, because its definitions and statements will be used in the proof of Theorem 3.2.

3.2. **Extra automorphisms and 'arbitrary assignment'.** Before we proceed, we must discuss a subtlety concerning extra automorphisms (see the related Remark 2.17). In particular, the composition of the isogenies along a walk in $\mathcal{G}_\ell$ is not well-defined, because the isogenies themselves are only defined up to post-composition by an automorphism. In the case of the automorphisms $[\pm 1]$, which commute with all the isogenies in the chain, this amounts to a single sign ambiguity. But extra automorphisms may wreak havoc: the endomorphism obtained by composing around a closed walk may not even have the same discriminant (and quadratic field) after a post-composition by an automorphism.

This would appear to lampoon any hope of associating endomorphisms nicely with closed walks. However, the remedy to this issue is the following observation. Consider two walks which differ by replacing a segment $j_1 \xrightarrow{\phi} 0 \xrightarrow{\psi} j_2$ with $j_1 \xrightarrow{[\zeta_3^k]\phi} 0 \xrightarrow{\psi[\zeta_3^{-k}]} j_2$, or $j_3 \xrightarrow{\phi} 1728 \xrightarrow{\psi} j_4$ with $j_3 \xrightarrow{[i]\phi} 1728 \xrightarrow{\psi[-i]} j_4$, where $[\zeta_3]$ and $[i]$ are extra automorphisms associated to primitive third and fourth roots of unity, respectively. In each case, the first arrow is the same directed edge of $\mathcal{G}_\ell$, but the second arrow differs as an edge of $\mathcal{G}_\ell$. The composition of the walk, however, is unchanged. In fact, there are typically three distinct walks $j_1 \to 0 \to j_2$ in $\mathcal{G}_\ell$, and three distinct resulting isogenies $j_1 \to j_2$, but there is no *canonical* identification between the elements of these two sets of size three. One way to set a non-canonical identification is to consider the incoming isogeny $j_1 \to 0$ to be a particular fixed (but arbitrary) representative of its equivalence class (recall that the equivalence class consists of all the isogenies with the same kernel, i.e. up to post-composition by an automorphism).

Definition 3.3 (Arbitrary assignment)**.** An *arbitrary assignment* is a choice of isogeny ϕ, up to sign, from every equivalence class of isogenies represented by a directed edge of $\mathcal{G}_\ell$.

There is no choice except when the codomain has extra automorphisms. **In what follows, we shall assume that such a fixed arbitrary assignment has been made for every arrow entering $j = 0$ or $j = 1728$.** This choice will affect the exact bijection we obtain in our main theorem, but not the fact that it is a bijection. Figure 3 is instructive as to the subtleties that may arise.

The following lemma is now immediate.

Lemma 3.4. *Assume we have made an arbitrary assignment for $\mathcal{G}_\ell$. Then any closed walk in $\mathcal{G}_\ell$ results in a unique endomorphism (up to sign) by composition of its component isogenies in order.*

FIGURE 3. This figure shows all the 2-isogenies between $j = 0$ and $j = 121$ for $p = 179$. Dual pairs are indicated by arrows matching in style. The isogenies entering $j = 0$ (upward in the picture) differ by post-composition by an automorphism: either $[\pm 1]$, $[\pm \zeta_3]$ or $[\pm \zeta_3^2]$. Therefore they are represented by a single directed edge in $\mathcal{G}_\ell$. The downward arrows from $j = 0$ to $j = 121$ are represented by distinct directed edges in $\mathcal{G}_\ell$. The two walks $\widehat{\varphi}[\zeta_3] \circ \varphi$ (up on dashed, down on solid) and $\widehat{\varphi} \circ [\zeta_3]\varphi$ (up on dotted, down on dashed) give the *same* endomorphism of $j = 121$, namely $\widehat{\varphi} \circ [\zeta_3] \circ \varphi$, despite being different walks in $\mathcal{G}_\ell$. In fact, the nine distinct closed walks of length 2 from 121 to 0 and back compose to only three distinct endomorphisms of $j = 121$, namely $\pm[2]$, $\pm\widehat{\varphi} \circ [\zeta_3] \circ \varphi$, or $\pm\widehat{\varphi} \circ [\zeta_3^2] \circ \varphi$. Finally, notice that *any* walk of length two from $j = 0$ to $j = 121$ and back is considered backtracking according to Definition 3.5. The same is not true for walks of length two from j=121 to $j = 0$ and back.

The endomorphism obtained by composing the isogenies of a cycle in order will simply be called the *composition* of the cycle. The fact that composition is well defined allows us to define a notion of backtracking without ambiguity around curves with extra automorphisms.

Definition 3.5 (Backtracking closed walk)**.** Given an arbitrary assignment on $\mathcal{G}_\ell$, two consecutive edges $\bullet \xrightarrow{\phi_0} \bullet \xrightarrow{\phi_1} \bullet$ are said to be *backtracking* if $\phi_1 \circ \phi_0 = [\ell]$, up to possible post-composition by an automorphism. A closed walk is considered to contain backtracking if any consecutive edges, including the last and first, are backtracking.

This definition of backtracking is the same definition as [7, Definition 4.3], except those authors leave out the composition of the last and first steps, which is more appropriate in their context. Note that this definition is not captured by the sequence of vertices alone. To wit, see Figure 3 for an example of two length-two walks, one of which is backtracking and one of which is not backtracking, but which are taken to each other by an isomorphism of the graph fixing vertices.

We also have a converse to Lemma 3.4.

Lemma 3.6. *Assume we have made an arbitrary assignment for $\mathcal{G}_\ell$. Consider an endomorphism $\theta \in \mathrm{End}(E)$ of ℓ-power degree. If θ is not divisible by $[\ell]$, then there is a unique closed walk in $\mathcal{G}_\ell$ with composition $\varphi\theta$, for some $\varphi \in \mathrm{Aut}(E)$. Furthermore, the walk has no backtracking, and φ is unique up to sign.*

Proof. Since $[\ell] \nmid \theta$, the kernel $\ker(\theta) \cap E[\ell]$ has size ℓ. Define the isogeny $\phi_1 : E \to E'$ corresponding to this subgroup (using the arbitrary assignment if applicable). Then any isogeny cycle which composes to θ must factor through this isogeny, hence this is the first step in the desired isogeny cycle: $\theta = \phi' \circ \phi_1$. Now take $\ker \phi' \cap E'[\ell]$ to find the next kernel, and the map ϕ_2 so that $\theta = \phi'' \circ \phi_2 \circ \phi_1$. Continuing, we obtain a closed walk in the ℓ-isogeny graph composing to an isogeny with the same kernel as θ. Which is to say, $\pm\theta$, up to post-composition by an automorphism of E. But by construction, this walk is unique. It has no backtracking, since backtracking would produce a composition divisible by $[\ell]$. The same closed walk is obtained if we begin with θ or a post-composition of θ by an automorphism on E. Thus, by Lemma 3.4, the automorphism φ of the statement is unique up to sign. $\square$

Finally, there is one more small lemma we will require, that shows that any closed walk composing to a power of an endomorphism is a power of a closed walk.

Lemma 3.7. *Suppose we have made an arbitrary assignment for $\mathcal{G}_\ell$ and let $\iota\colon K \to \mathrm{End}^0(E)$ be an $\mathcal{O}$-primitive orientation. Let $\alpha = u\beta^s$ for some $s > 1$, $\beta \in \mathcal{O}$, and some unit $u \in \mathcal{O}$. Assume that ℓ does not divide α, and that a closed walk C of $\mathcal{G}_\ell$ composes to $\pm\iota(\alpha)$. Then $C = C_\beta^s$ where C_β is a closed walk of $\mathcal{G}_\ell$ that composes to $\pm\iota(\beta)$.*

Proof. By Lemma 3.6, since ℓ does not divide β, there exists a unique closed walk C_β of ℓ-isogenies composing to $\varphi\iota(\beta)$ for some automorphism φ of E. If E has no extra automorphisms, then the proof is done.

If E has extra automorphisms, however, then we slightly modify C_β. Change the initial edge ϕ of C_β by replacing it with $\phi\varphi^{-1}$ (having the same domain and codomain). Call this modified closed walk C_β'. By the conventions of the arbitrary assignment (Definition 3.3), the composition around C_β' is $\varphi\iota(\beta)\varphi^{-1}$. The composition of $(C_\beta')^s$ is $\varphi\iota(\beta)^s\varphi^{-1} = \varphi\iota(u^{-1})\iota(\alpha)\varphi^{-1}$. But then we have closed walks C and $(C_\beta')^s$, composing to $\pm\iota(\alpha)$ and $\pm\varphi\iota(u^{-1})\iota(\alpha)\varphi^{-1}$, respectively. By the arbitrary assignment and Lemma 3.6, we can again modify the first edge of $(C_\beta')^s$ to obtain C' which composes to $\pm\varphi\iota(u^{-1})\iota(\alpha)$. But then $C = C'$ by Lemma 3.6, and so $\varphi = \pm\iota(u) \in \iota(\mathcal{O})$. Therefore $\varphi \in \iota(\mathcal{O})$, in which case $(C_\beta')^s$ composes to $\pm\iota(u^{-1})\iota(\alpha)$, and by Lemma 3.6, $C = (C_\beta')^s$. $\square$

3.3. **Proof of Theorem 3.2.** In order to prove Theorem 3.2, it is easier to first discuss rims and isogeny cycles with a basepoint (a marked point). Let E be a fixed vertex of $\mathcal{G}_\ell$. We define the following sets.

$$\mathcal{R}_{E,r} = \left\{ \begin{array}{c} \text{directed rims of size } r \text{ in } \mathcal{G}_{K,\ell} \\ \text{passing through } (E, \iota) \end{array} : \begin{array}{c} K \text{ is an imaginary quadratic field,} \\ \iota \text{ is a } K\text{-orientation on } E \end{array} \right\},$$

$$\mathcal{I}_{E,r} = \left\{ (\iota, \mathfrak{l}) : \begin{array}{c} \iota\colon K \to \mathrm{End}^0(E) \text{ a primitive } \mathcal{O}\text{-orientation,} \\ K \text{ is an imaginary quadratic field,} \\ \mathcal{O} \text{ is an } \ell\text{-fundamental order of } K, \\ (\ell) = \bar{\mathfrak{l}}\mathfrak{l} \text{ splits in } \mathcal{O}, \\ [\mathfrak{l}] \text{ has order } r \text{ in } \mathrm{Cl}(\mathcal{O}), \end{array} \right\}, \text{ and }$$

$$\mathcal{C}_{E,r} = \{\text{isogeny cycles of length } r \text{ in } \mathcal{G}_\ell \text{ starting and ending at } E \}.$$

Let $\sim$ denote conjugation on elements of $\mathcal{I}_{E,r}$, i.e. taking a pair $(\iota, \mathfrak{l})$ to $(\bar\iota, \bar{\mathfrak{l}})$. As described in Section 3.1, conjugation preserves the direction of the rim but conjugates the ideal whose action traverses the rim. We will continue to use the symbol $\sim$ to denote conjugation of a rim.

We assume a safe arbitrary assignment (Definition 3.13) has been made for $\mathcal{G}_\ell$. Depending upon this arbitrary assignment, we define the following maps:

(8)
$$\begin{array}{ccc} \mathcal{I}_{E,r}/\sim & \underset{\Psi'}{\overset{\Psi}{\rightleftarrows}} & \mathcal{R}_{E,r}/\sim \\ {\scriptstyle \Theta}\searrow & & \swarrow{\scriptstyle \Phi} \\ & \mathcal{C}_{E,r} & \end{array}$$

The goal is to show that the diagram commutes, where every map is a bijection. Then Theorem 3.2 will follow by 'forgetting' the basepoint.

Disregarding conjugation for the moment, we first define maps

$$\mathcal{I}_{E,r} \underset{\Psi'}{\overset{\Psi}{\rightleftarrows}} \mathcal{R}_{E,r} \ .$$

In Lemma 3.8, we show that these are inverse bijections and that they provide inverse bijections even after a quotient by the conjugation relation. In Lemma 3.9 we will define Θ and show it is a bijection. Lemma 3.10 will define Φ and show it is a bijection. Finally, in Theorem 3.11, we show that the triangle in (8) commutes.

Definition of Ψ'. Consider a directed rim $R \in \mathcal{R}_{E,r}$ passing through (E, ι). Denote by $\mathcal{O}$ the order such that ι is $\mathcal{O}$-primitive. Then $\mathcal{O}$ is ℓ-fundamental since (E, ι) is at the rim. Moreover, the $\mathcal{O}$-ideal (ℓ) splits and the $\mathcal{O}$-ideals above ℓ have order r in $\mathrm{Cl}(\mathcal{O})$ since (E, ι) belongs to a rim of size r. The direction along the rim corresponds to the action of one of the ideals above ℓ, which we denote $\mathfrak{l}$. Then Ψ' is defined as

$$\Psi'\colon \mathcal{R}_{E,r} \to \mathcal{I}_{E,r}, \quad R \mapsto (\iota, \mathfrak{l}).$$

Definition of Ψ. On the other hand, consider $(\iota, \mathfrak{l}) \in \mathcal{I}_{E,r}$. The orientation ι on E gives rise to a directed rim of size r passing through (E, ι) by the consecutive action of $\mathfrak{l}$. Let Ψ send $(\iota, \mathfrak{l})$ to this directed rim in $\mathcal{R}_{E,r}$, i.e.

$$\Psi \colon \mathcal{I}_{E,r} \to \mathcal{R}_{E,r}, \quad (\iota, \mathfrak{l}) \mapsto R = ((E, \iota) \to_{\phi_{\mathfrak{l}}} \cdots).$$

Lemma 3.8. *The maps Ψ and Ψ' defined above are inverse bijections. They also provide bijections*

$$\mathcal{I}_{E,r}/\!\sim \; \underset{\Psi'}{\overset{\Psi}{\rightleftarrows}} \; \mathcal{R}_{E,r}/\!\sim$$

where $\sim$ denotes equivalence between conjugates.

Proof. The fact that Ψ and Ψ' are inverse to each other follows from their definitions. It is also clear that conjugate rims correspond to conjugate orientations (while fixing E and K). $\qquad\square$

Definition of Θ. We now define a map Θ from $\mathcal{C}_{E,r}$ to $\mathcal{I}_{E,r}$. Let $C \in \mathcal{C}_{E,r}$. Starting at the vertex E, this cycle composes into an endomorphism $\pm\theta \in \mathrm{End}(E)$. This endomorphism $\pm\theta$ induces an orientation $\iota_\theta \colon K \to \mathrm{End}^0(E)$ on E (defined up to conjugation). Fixing a choice for ι_θ, let $\mathcal{O}$ be the order for which ι_θ is primitive. Then ℓ must split or ramify in $\mathcal{O}$ because $\theta \in \mathcal{O}$ has norm ℓ^r and is not divisible by ℓ (since $C \in \mathcal{C}_{E,r}$ has no backtracking). If $\phi_{\mathfrak{l}}^{(E, \iota_\theta)}$ is equal to the first isogeny of the cycle C in its indicated direction, for a prime $\mathfrak{l}$ above ℓ, then define $\Theta(C) := (\iota_\theta, \mathfrak{l})$.

Lemma 3.9. *The process described above gives a well-defined map $\Theta \colon \mathcal{C}_{E,r} \to \mathcal{I}_{E,r}/\!\sim$ (in particular, the ideal $\mathfrak{l}$ indicated in the definition above exists).*

Proof. Let $C_{E,r} \in \mathcal{C}_{E,r}$ and let $\theta \in \mathrm{End}(E)$ be the composition of the edges in $C_{E,r}$ (unique up to sign). Let α denote the preimage of θ in K via ι_θ. The orientation ι_θ is a primitive $\mathcal{O}$-orientation for exactly one order $\mathbb{Z}[\alpha] \subseteq \mathcal{O} \subseteq \mathbb{Q}(\alpha)$. If $\mathbb{Z}[\alpha]$ has conductor divisible by ℓ, then α has discriminant and norm divisible by ℓ, hence trace divisible by ℓ, so it is itself divisible by ℓ. But this is impossible because the isogeny cycle $C_{E,r}$ has no backtracking. Hence $\mathbb{Z}[\alpha]$, and therefore $\mathcal{O}$, are ℓ-fundamental orders.

Note that α has norm ℓ^r. We consider the factorization of (ℓ) in $\mathcal{O}$. As mentioned in the definition of Θ, ℓ splits or ramifies, so $(\ell) = \bar{\mathfrak{l}}\mathfrak{l}$. Then we can factor (α), using its norm: $(\alpha) = (\ell)^t \mathfrak{l}^s$ for some t, s (up to possibly replacing $\mathfrak{l}$ with its conjugate). However, $t = 0$ since ℓ does not divide α. Hence $s = r$, by comparing norms. That is, $(\alpha) = \mathfrak{l}^r$.

This shows that ℓ is represented by a class $[\mathfrak{l}]$ of order dividing r in $\mathrm{Cl}(\mathcal{O})$. We will show it is of order exactly r. If not, then $\mathfrak{l}^s = (\beta) \in \mathcal{O}$, so $\alpha = u\beta^s$ for some unit $u \in \mathcal{O}$. Then by Lemma 3.7, the cycle $C_{E,r}$ is a power of a smaller cycle, a contradiction to the definition of isogeny cycle.

Finally, since $(\alpha) = \mathfrak{l}^r$, the kernel $E[\mathfrak{l}]$ is contained in $\ker\theta$, and so $\phi_{\mathfrak{l}}$ defined on (E, ι_θ) is equal to the first isogeny of C in its direction of travel.

Note that our only arbitrary choice is that of ι_θ in place of its conjugate, but as we map into $\mathcal{I}_{E,r}/\!\sim$, the other choice results in the same image element, and the map is well-defined. $\qquad\square$

Definition of Φ. Next we consider the map Φ. Given a directed rim $R \in \mathcal{R}_{E,r}/\!\sim$, we get a walk in $\mathcal{G}_\ell$ by forgetting orientations. If there are no extra automorphisms on any curves of the rim, this is well-defined up to a sign on the isogenies. If there are extra automorphisms, we must use the process in Section 3.4 to safely forget orientations. The resulting cycle will be denoted $\Phi(R)$. Note that if R is conjugated, we obtain the same walk. This is clear when there are no extra automorphisms. When we use the process in Section 3.4, this follows from the observation that any valid set of representatives is again valid for the conjugate rim, in the proof of Proposition 3.14.

Lemma 3.10. *Let $R \in \mathcal{R}_{E,r}$. Then $\Phi(R) \in \mathcal{C}_{E,r}$.*

Proof. Let W be the walk obtained from R by forgetting the orientations. We need to show that W doesn't have backtracking and is not repeating smaller cycles. Let $(\iota, \mathfrak{l}) = \Psi'(R) \in \mathcal{I}_{E,r}/\!\sim$. First, we show that the action of $\mathfrak{l}$, which dictates the direction of R, cannot produce a backtracking path. Suppose the action of $\mathfrak{l}$ produces a backtracking path ϕ_1, ϕ_2. Then there is an endomorphism φ of E such that $\varphi[\ell] = \phi_2\phi_1$ has the same kernel as $\phi_{\mathfrak{l}^2}$, namely $E[\ell]$. This implies $\mathfrak{l}^2 = (\ell)$, so $r = 2$, a contradiction.

Now suppose that the walk W in $\mathcal{G}_\ell$ repeats a smaller closed walk $k > 1$ times. Let $\phi \in \mathrm{End}(E)$ be the composition of the smaller closed walk. Then $\phi = \phi_{\mathfrak{a}}$ for the ideal $\mathfrak{a} = \mathfrak{l}^{r/k}$, and $\phi^k = \phi_{\mathfrak{b}}$ where $\mathfrak{b} = \mathfrak{l}^r = (\beta)$

for some $\beta \in \mathcal{O}$. Hence as an element of $\mathrm{End}(E)$, $\phi^k = \iota(\beta)$, which implies that $\phi \in \iota(K)$. But then letting α be such that $\iota(\alpha) = \phi$, we have $\alpha^k = \beta$ and $\mathfrak{l}^{r/k} = (\alpha)$. This contradicts the rim being of size r. Thus, the walk W in $\mathcal{G}_\ell$ cannot repeat a smaller walk $k > 1$ times. $\square$

Theorem 3.11. *Given a safe arbitrary assignment for $\mathcal{G}_\ell$, we define maps Φ and $\Phi' := \Psi \circ \Theta$ as above. Then the sets $\mathcal{C}_{E,r}$ and $\mathcal{R}_{E,r}/\sim$ are in bijection under the inverse pair of maps Φ and Φ'.*

Proof. We make a safe arbitrary assignment (Definition 3.13), to define the maps above. We will show that the two maps are inverse to one another.

Suppose $C \in \mathcal{C}_{E,r}$ is taken to $[R] \in \mathcal{R}_{E,r}/\sim$ by Φ'. If θ is the composition of C, then R is the rim passing through (E, ι_θ) ($\overline{R}$ is the rim passing through $(E, \overline{\iota_\theta})$, respectively). In particular, $\mathfrak{l}^r = (\alpha)$ where $\iota_\theta(\alpha) = \theta$. Therefore, composing the isogenies around R results in the endomorphism $\pm\theta$. By Lemma 3.6 (and using Proposition 3.14 in the case of extra automorphisms), the rim R maps to C when orientations are forgotten (i.e. under Φ).

Next, suppose $[R] \in \mathcal{R}_{E,r}/\sim$ is taken to $(\iota_\theta, \mathfrak{l}) \in \mathcal{I}_{E,r}$ by Ψ'. Let $C = \Phi(R)$. Then C composes to θ (using Proposition 3.14 in the case of extra automorphisms). Thus C maps to $(\iota_\theta, \mathfrak{l}) \in \mathcal{I}_{E,r}$ under Θ. $\square$

The discussions so far relies on a choice of a basepoint E. We now prove the main theorem (Theorem 3.2) as a corollary to the basepointed version (Theorem 3.11).

Proof of Theorem 3.2. Make a safe arbitrary assignment for $\mathcal{G}_\ell$ and use it to define the maps of diagram (8) as described above.

There is a map $\mathcal{C}_{E,r} \to \mathcal{C}_r$ given by forgetting the basepoint of the isogeny cycle. There is also a map $\mathcal{R}_{E,r}/\sim \,\to \mathcal{R}_r/\sim$ given by forgetting the basepoint of the rim.

Given $[R] \in \mathcal{R}_r/\sim$, a basepoint can be chosen in r ways. For each such basepoint (E, ι), we obtain $R_{E,\iota} \in \mathcal{R}_{E,r}$, which we map to $C_{E,\iota} := \Phi(R_{E,\iota}) \in \mathcal{C}_{E,r}$ via Theorem 3.11. We must show that the resulting $C_{E,\iota}$ are all equal when basepoints are forgotten. But the map Φ to isogeny cycles is agnostic to the basepoint: this is clear for rims without extra automorphisms, as there are no choices to make, and is a consequence of Proposition 3.14 when there are extra automorphisms.

Similarly, a given $C \in \mathcal{C}_r$ corresponds to $C_i \in \mathcal{C}_{E_i,r}$ for E_i the i-th curve in the cycle. Each of these basepointed cycles gives a corresponding basepointed rim $R_i := \Phi'(C_i) \in \mathcal{R}_{E_i,r}$. We must show that the resulting R_i are all equal, when basepoints are forgotten. Define $(\iota_i, \mathfrak{l}_i) = \Theta(C_i) \in \mathcal{I}_{E,r}$. Define the composed endomorphism $\theta_i \in \mathrm{End}(E_i)$ for each E_i. Then by construction, $\iota_i = \iota_{\theta_i}$ and $\phi_i = \phi_{\mathfrak{l}_i}^{(E_i, \iota_i)}$. It suffices to show that all $(\iota_i, \mathfrak{l}_i)$ give rise to the same rim. The θ_i are related by $\theta_{i+1} = \phi_i \theta_i \phi_i^{-1}$. Thus, we obtain $\iota_{i+1} = (\phi_i)_* \iota_i$. In particular, the orientations ι_i are all associated to the same field K, and they are all primitive for the same order $\mathcal{O}$ of K, since ϕ_i is horizontal. Therefore (E_i, ι_i) and (E_{i+1}, ι_{i+1}) are elements of $\mathrm{SS}_\mathcal{O}^{\mathrm{pr}}$ connected by a horizontal isogeny ϕ_i of degree ℓ, i.e. they are on the same rim. As for direction, by construction $\phi_i = \phi_{\mathfrak{l}_i}^{E_i}$ for all i. Since ϕ_i all indicate the same direction, the $\mathfrak{l}_i$ all indicate the same direction. $\square$

3.4. **Mapping from $\mathcal{G}_{K,\ell}$ to $\mathcal{G}_\ell$.** We now deal with the one remaining item pushed aside during the proof: intricacies introduced by extra automorphisms when defining a map from rims to closed walks by forgetting orientations. One can hope to map any walk in $\mathcal{G}_{K,\ell}$ to a walk in $\mathcal{G}_\ell$ by forgetting orientations. However, isogenies in the oriented isogeny graph are taken up to pre- and post-composition by automorphisms (see (7)), so when mapping down to $\mathcal{G}_\ell$, there is a choice to be taken if the domain has extra automorphisms. Even more to the point, if we obtain a closed walk from a rim, by making the choice arbitrarily, we may end up with a closed walk whose composition does not lie in $\iota(K)$. (Recall that in order to obtain a bijection, we hope to recover the rim from the isogeny cycle using this composition.)

In this section, we sidestep this issue, with the help of a special type of arbitrary assignment, by showing that it is nevertheless possible to accomplish something more limited: to define a map from rims to isogeny cycles such that the resulting isogeny cycle has a composition in $\iota(K)$. We emphasize that this is a non-issue (and this entire section can be skipped) if $\mathcal{G}_\ell$ has no curves with extra automorphisms. To be somewhat more precise, consider a rim R in $\mathcal{G}_{K,\ell}$. If we choose representatives of the vertices and edges (which are, a priori, only equivalence classes), then we can forget orientations to obtain a closed walk in $\mathcal{G}_\ell$. However, the choice of representatives must be made with care. First, the edges must belong to the arbitrary assignment: if not, then when composing around the rim before and after forgetting orientations, we may differ by a

post-composition, and hence may produce an endomorphism from the wrong quadratic field (which we will see momentarily in Example 3.12 below). Second, the choice of representatives must be consistent: when we traverse the entire rim, beginning at representative (E, ι), we must return to representative (E, ι), and not merely an isomorphic oriented curve. In the proof that follows, we spend the vast majority of our energy in showing that such a choice of representative exists and is unique.

Example 3.12. An example demonstrates some of the problems the proof will need to overcome. Consider the example $p = 11$ and $\ell = 3$ where $\mathcal{G}_3$ has two loops at $j = 1728$. Recall the endomorphism ring of the curve with j-invariant $j = 1728$ from Section 2.1. There are eight endomorphisms of degree $\ell = 3$, namely

$$\frac{\pm 1 \pm \pi_p}{2} \quad \text{and} \quad \frac{\pm[i] \pm [i]\pi_p}{2}.$$

The first group have trace ± 1 and the second group have trace 0. The endomorphisms

$$\pm \frac{1 + \pi_p}{2}, \quad \pm \frac{[i] + [i]\pi_p}{2}$$

have the same kernel, which we denote κ. This corresponds to one of the loops, and these four endomorphisms differ by post-composition by an automorphism (namely $\pm 1, \pm[i]$). The endomorphisms

$$\pm \frac{1 - \pi_p}{2}, \quad \pm \frac{[i] - [i]\pi_p}{2}$$

all have the same kernel, namely the image of κ under $[i]$. This corresponds to the second loop, and these four endomorphisms also differ from one another only by post-composition by an automorphism.

If, in the arbitrary assignment, we choose $\frac{1+\pi_p}{2}$ for the first loop and $\frac{1-\pi_p}{2}$ for the second, then these endomorphisms generate orientations by $\mathbb{Z}\left[\frac{1+\sqrt{-11}}{2}\right]$. The endomorphisms $\pm([i] + [i]\pi_p)/2$ of trace zero generate orientations by $\mathbb{Z}[\sqrt{-3}]$. The graph $\mathcal{G}_{\mathbb{Q}(\sqrt{-3}),\ell}$ has one rim in $\mathrm{SS}_{\mathbb{Z}[\sqrt{-3}]}$, namely the loop on the oriented curve (E, ι) where E has j-invariant 1728 and $\iota(\sqrt{-3}) = [\pm\sqrt{-3}]$ (case (i) of Remark 2.7). Under this arbitrary assignment, there is no place to map this rim in such a way that the composition of the resulting loop will match the order $\mathbb{Z}[\sqrt{-3}]$ as desired.

This motivates modifying Definition 3.3 so that the two loops must compose to endomorphisms from different fields K. In particular, we must be careful to choose the arbitrary assignment 'in tandem' for multiple edges between curves with extra automorphisms. This issue arises in larger cycles also. For an explicit example, consider $p = 11$, and $\mathbb{Q}(\sqrt{-23})$, which has class group of size 3 generated by $\mathfrak{l}$ lying above $\ell = 2$; the rim for the maximal order is of size 3 with j-invariants $1728, 1728, 0$. We therefore make a special requirement on the arbitrary assignment which applies to this case.

Definition 3.13 (Safe arbitrary assignment). An arbitrary assignment is said to be *safe* if it obeys the following rule: whenever there is a collection of edges $E_1 \to E_2$ related by precomposition by an automorphism of E_1 besides $[\pm 1]$, we choose their representatives to be of the form $\{\pm\phi\varphi : \varphi \in \mathrm{Aut}(E_1)\}$ for a fixed choice ϕ.

The rule in Definition 3.13 is automatically satisfied at E_1 and E_2 by any arbitrary assignment unless both E_1 and E_2 have extra automorphisms, and there are multiple edges between them. It is also clear that a safe arbitrary assignment is always possible.

With regards to Example 3.12, note that $\frac{1+\pi_p}{2} = -[i]\frac{1-\pi_p}{2}[i]$, which violates Definition 3.13. This is the reason the example 'fails' the bijection.

Proposition 3.14. *Suppose a safe arbitrary assignment has been made for $\mathcal{G}_\ell$. There is a map from rims R of $\mathcal{G}_{K,\ell}$ to closed walks $C(R)$ in $\mathcal{G}_\ell$ with the following properties:*

(1) *The curves and isogenies of $C(R)$ are representatives of the isomorphism classes of curves and isogenies of R, in the same order; in particular, there is a map $R \to C(R)$ of directed graphs.*

(2) *The composition $\pm\theta$ of $C(R)$ at vertex E in $C(R)$ corresponding to (E, ι) in R, gives rise to the orientation $\pm\iota$ in the sense that $\theta = \iota(\alpha)$ for some $\alpha \in K$.*

Proof. **Defining the map.** It is possible to define such a map as follows. Consider a rim R of length r, associated to an ideal $\mathfrak{l}$ above ℓ in the associated ℓ-fundamental order $\mathcal{O}$, with basepoint (E_0, ι_0), written as

follows:

$$(E_0, \iota_0) \xrightarrow{\phi_0} (E_1, \iota_1) \xrightarrow{\phi_1} (E_2, \iota_2) \cdots (E_{r-1}, \iota_{r-1}) \xrightarrow{\phi_{r-1}} (E_0, \iota_0) .$$

Write $\phi := \phi_{r-1}\phi_{r-2}\cdots\phi_1\phi_0 \in \text{End}(E)$. The result ϕ depends on representative choices of the isomorphism classes of oriented curves (E_i, ι_i) and of the isogenies ϕ_i shown above. We seek to choose such representatives so that

(i) $\phi_*\iota_0 = \pm\iota_0$,
(ii) each ϕ_i is chosen from the safe arbitrary assignment.

Such a choice will be called a *valid set of representatives* during the course of this proof. With such a valid set of representatives, to define a map from rims to isogeny cycles, we then forget orientations to obtain an isogeny cycle in $\mathcal{G}_\ell$. This is now well-defined because of (ii) and will define the map referred to in the statement of the proposition. For any rim R, we will write $C(R)$ for the isogeny cycle which is its image.

Facts about valid sets of representatives. We will need a few facts about valid sets of representatives. First, we will show that if we have a valid set of representatives with respect to a basepoint (E_0, ι_0), then it is also a valid set of representatives for any other (E_i, ι_i) in the cycle. In other words, for any k, we have $(\phi_{r-k}\cdots\phi_{r-k+1})_*\iota_{r-k+1} = \pm\iota_{r-k+1}$ (where indices are taken modulo r). To see this for (E_1, ι_1), use $\iota_1 = (\phi_0)_*\iota_0$ to compute the left hand side, using the fact that $\phi_*\iota_0 = \pm\iota_0$; repeat for E_2, E_3 etc.

Let $\alpha \in \mathcal{O}$ be such that $\mathfrak{l}^r = (\alpha)$. Without causing confusion, we will write $\phi_\mathfrak{l}: (E_i, \iota_i) \to (E_{i+1}, \iota_{i+1})$ for each i (this isogeny is defined up to isomorphism, by the convention that its kernel is $E_i[\iota_i(\mathfrak{l})]$). We will now prove that any valid set of representatives also satisfies

(iii) $\ker \phi_i = \ker \phi_\mathfrak{l} = E_i[\iota_i(\mathfrak{l})]$,
(iv) $\phi = \pm\iota_0(\alpha)$.

To prove this, observe that $\ker \phi_{\mathfrak{l}^r} = E_i[\iota_0(\alpha)] = \ker \iota_0(\alpha)$. In particular, $\iota_0(\alpha)$ is K-isomorphic to $\phi_{\mathfrak{l}^r}$. Furthermore, $(\iota_0(\alpha))_*\iota_0 = \pm\iota_0$ because elements of $\iota_0(K)$ commute. Next, since $\phi_*\iota_0 = \pm\iota_0 = (\iota_0(\alpha))_*\iota_0$ (by (i)) it must be that $\phi = \pm\iota_0(\alpha)$, since this is the only pre-/post-composition ϕ of $\iota_0(\alpha)$ that continues to satisfy $\phi_*\iota_0 = \pm\iota_0$. This implies that $\ker \phi = \ker \phi_{\mathfrak{l}^r}$ and hence $\ker \phi_0 = \ker \phi_\mathfrak{l}$. Since the valid set of representatives is valid at each basepoint, we have $\ker \phi_i = \ker \phi_\mathfrak{l}$ by the same reasoning.

Claim: The map defined on a valid set of representatives satisfies requirements (1) **and** (2) **of Proposition 3.14.** By (iv), $\phi = \pm\iota_0(\alpha)$. The composition of the isogeny cycle $C(R)$ is $\pm\phi$, which demonstrates (2) at the vertex (E_0, ι_0). Then (2) holds at the other vertices in R by the fact that a valid set of representatives is valid at each of the vertices of R. The fact that the map satisfies (1) is by definition.

Proof outline. It remains to show that there is a set of representatives satisfying (i)–(ii) (i.e. it is a valid set of representatives), and that any such set leads to the same isogeny cycle (i.e. the map $R \to C(R)$ given above is well-defined). To do this, we will divide the possible rims into two cases: exceptional and non-exceptional. The exceptional case is the following: first, we assume every vertex of the rim has the same j-invariant with extra automorphisms, of which φ is a non-trivial ($\neq [\pm 1]$) element; then, defining ϕ as above for any basepoint, we assume that no matter which vertex we start with as basepoint, we have $\ker \phi = \ker \phi\varphi$.

The map $R \mapsto C(R)$ is well-defined. Let us show that the map $R \mapsto C(R)$ is well-defined in the non-exceptional cases. A priori, it may be that there are two choices of representatives satisfying (i)–(ii). Suppose so. Use primes (ι', ϕ_i' etc.) to denote the second valid set of representatives.

First note that if $(E_i, \iota_i') = (E_i, \pm\iota_i)$ at any vertex, then the two choices of representatives must be the same. To see this, note that the vertex representative $(E_i, \pm\iota_i)$ determines the kernel of ϕ_i (by (iii)), which in turn determines ϕ_i up to sign (by (ii)). Therefore $\phi_i' = \pm\phi_i$. But this implies that $(E_{i+1}, \iota_{i+1}') = (E_{i+1}, \pm\iota_{i+1})$. Proceeding in this way, the two valid sets of representatives are the same. Hence the proof is complete if any curve in the rim has no extra automorphisms (in which case there is no choice for ι_i).

Suppose that we have $\iota_0' = \varphi_*\iota_0$ for some non-trivial (not ± 1) automorphism $\varphi \in \text{End}(E_0)$. In particular, E_0 has extra automorphisms. At the same time, $\iota_1' = \eta_*\iota_1$ for some $\eta \in \text{Aut}(E_1)$ (possibly trivial this time). These relationships entail that $\phi_0' = \eta\phi_0\varphi^{-1}$. By (ii), ϕ_0 and ϕ_0' both belong to the safe arbitrary assignment, so $\eta = \pm 1$ and hence $\iota_1' = \pm\iota_1$; thus, we have reduced to the case in the previous paragraph.

Existence of a valid set of representatives in non-exceptional cases. We use the safeness of the safe arbitrary assignment (Definition 3.13). Let ϕ_0 be the map with kernel $E_0[\iota_0(\mathfrak{l})]$ and post-composition determined by the safe arbitrary assignment. Then let $\iota_1 = (\phi_0)_*\iota_0$. Continue in this manner around

the entire rim. When we close the rim (reaching (E_0, ι_0) again), we find that $\phi_* \iota_0 = \varphi \iota_0 \varphi^{-1}$ for some automorphism φ (since the rim closes, we must reach an oriented curve isomorphic to the initial oriented curve). If E_0 has no extra automorphisms, then the proof of this case is complete. Hence we are done if the cycle contains even one curve without extra automorphisms.

Let us therefore restrict to the case that E_0 has extra automorphisms and so do all other curves in the rim. If item (i) fails, then there is a nontrivial automorphism $\varphi \in \mathrm{Aut}(E_0)$ so that $\phi_* \iota_0 = \varphi_* \iota_0$. Then

$$(\phi \varphi^{-1})_* (\varphi_* \iota_0) = \phi_* \iota_0 = \varphi_* \iota_0.$$

Thus, replacing isomorphism class representatives (E_0, ι_0) with $(E_0, \varphi_* \iota_0)$ and ϕ_0 with $\phi_0 \varphi^{-1}$, we recover (i). Provided that ϕ_0 and $\phi_0 \varphi^{-1}$ have different kernels, which they do when $E_1 \not\cong E_0$ (Lemma 2.9), then using the conventions of the safe arbitrary assignment (Definition 3.13), we have the freedom to replace ϕ_0 with $\phi_0 \varphi^{-1}$ without violating item (ii).

Thus we are reduced to the case that all E_i are isomorphic (assumed equal without loss of generality), and ϕ_0 and $\phi_0 \varphi$ have the same kernel (for every choice of basepointing the rim). This is the exceptional case.

Existence of a valid set of representatives in the exceptional case. Suppose we are in the exceptional situation described in the proof outline. That is, all vertices are isomorphic curves with extra automorphisms, and $\ker \phi = \ker \phi \varphi$. Since $\phi \neq \pm \phi \varphi$ (as in the proof of Lemma 2.9), we have $\varphi^k \phi = \pm \phi \varphi$ for some integer k with $\varphi^k \neq \pm 1$. This means either $\phi \varphi = \varphi \phi$ or $\phi \varphi = \varphi \bar{\phi}$. Note that ϕ commutes with φ if and only if $\varphi \phi$ does, so commuting is independent of the safe arbitrary assignment.

The above is true for ϕ_0 as well as ϕ. But $\ell = N(\phi_0) < p$, which by Lemma 2.1 implies that ϕ and φ commute. The same argument applies to each ϕ_i for which $\ker \phi_i = \ker \phi_i \varphi$. But if there is a ϕ_i for which this is not true, we are out of the exceptional case. Hence every ϕ_i commutes with φ.

Therefore we may suppose ϕ commutes with φ. Then $\phi \in \mathbb{Q}(\varphi)$. Regardless of the safe arbitrary assignment, the commutativity implies $\phi_* \iota = \iota$ for the orientation $\iota(\zeta_3) = \varphi$ or $\iota(i) = \varphi$ for $j = 0$ and $j = 1728$, respectively. So suppose ϕ is the arbitrary assignment. Then if R is a rim with $\phi_\mathfrak{l} = \phi$ (up to post-composition by an automorphism), then R is associated to the field $K = \mathbb{Q}(\zeta_3)$ or $\mathbb{Q}(i)$ respectively, and ϕ provides a set of representatives satisfying (i)–(ii). Hence there is a valid set of representatives. $\square$

We observe the following immediate corollary to Theorem 3.2.

Corollary 3.15. *Let $(j_1, \ldots, j_n)$ be a finite sequence of j-invariants of supersingular elliptic curves in $\mathcal{G}_\ell$. Then the number of isogeny cycles in $\mathcal{G}_\ell$ on this ordered sequence of j-invariants is equal to the number of rims of $\bigcup_K \mathcal{G}_{K,\ell}$ on this ordered sequence of j-invariants.*

In fact, an alternate approach to proving Theorem 3.2 would be to observe that it is indeed equivalent to this corollary. One could then break into cases based on the existence of extra automorphisms. In most cases, there are no double edges, and this is a bijection of singleton sets.

3.5. **An undirected version of the main bijection.** The arbitrary assignment also allows us to set up a bijection between walks and their 'duals' (walking the same j-invariants backwards), despite the ambiguity at $j = 0$ and $j = 1728$. This means we can discuss the bijection between undirected rims and undirected isogeny cycles.

Lemma 3.16. *Assume we have made a safe arbitrary assignment for $\mathcal{G}_\ell$. Let $(j_0, j_1, \ldots, j_n)$ be a sequence of j-invariants for a closed walk in $\mathcal{G}_\ell$ (where j_i is adjacent to j_{i+1}, and j_n to j_0). Then the closed walks through this ordered list of j-invariants can be paired with the closed walks through the list $(j_n, j_{n-1}, \ldots, j_1, j_0)$ (i.e. going backwards) in such a way that the pair of endomorphisms obtained from each pair of walks (by composition) is a pair of dual endomorphisms.*

Proof. Each walk $\phi_0, \phi_1, \ldots, \phi_n$ with composition θ has a 'dual walk' obtained by replacing the individual isogenies with their duals and reversing the order: the walk $\widehat{\phi_n}, \widehat{\phi_{n-1}}, \ldots, \widehat{\phi_0}$. This results in a walk on the same vertices backwards. If the walk does not pass through $j = 0$ or $j = 1728$, this dual walk gives the dual endomorphism under composition (up to sign). The difficulty arises when choosing a dual for an outgoing edge $\phi_i \colon E_i \to E_{i+1}$ where E_i has j-invariant 0 or 1728, in which case the dual $\widehat{\phi_i} \colon E_{i+1} \to E_i$ may not be in the safe arbitrary assignment, and we are forced to use instead $\varphi \widehat{\phi_i} \colon E_{i+1} \to E_i$ for some $\varphi \in \mathrm{Aut}(E_i)$. In this case the endomorphism obtained by composition is no longer $\pm \widehat{\theta}$, because of the intervening φ between $\widehat{\phi_i}$

FIGURE 4. An example of an isogeny cycle that is a barbell (Definition 3.19). If we walk in the direction of the arrows from E_0, we obtain $\psi\widehat{\varphi_1}\cdots\widehat{\varphi_n}\psi'\varphi_n\cdots\varphi_1$. If we walk the cycle backwards (as defined in Lemma 3.16), we get $\widehat{\varphi_1}\cdots\widehat{\varphi_n}\widehat{\psi'}\varphi_n\cdots\varphi_1\widehat{\psi}$. If we ignore the basepoint and have $\psi=\widehat{\psi}$ and $\psi'=\widehat{\psi'}$, then these are the same cyclic sequence of arrows, in which case this is a barbell.

and $\widehat{\phi_{i-1}}$. However, this can be adjusted by choosing an appropriate replacement for $\widehat{\phi_{i-1}}\colon E_i \to E_{i-1}$ from among the available isogenies in the safe arbitrary assignment from E_i to E_{i-1}, which differ by precomposition by an automorphism. In other words, can obtain the composition $\widehat{\phi_{i-1}}\widehat{\phi_i}$ as $\widehat{\phi_{i-1}}\varphi^{-1}\varphi\widehat{\phi_i}$. In this way, we can obtain a dual path whose composition according to the safe arbitrary assignment is $\pm\widehat{\theta}$. Since this process can be performed in both directions (from clockwise paths to counterclockwise, and vice versa), inversely (by Lemma 3.6), the statement is proved. $\qquad\square$

The preceding lemma allows us to define, for any closed walk, its *backward walk*.

Lemma 3.17. *If $\overline{R}$ is the conjugate of R, then $\Phi(\overline{R})$ and $\Phi(R)$ are backward isogeny cycles of each other.*

Proof. The rim $\overline{R}$ consists of the set of edges and vertices traversed backward, where an isogeny is replaced with its dual. From the proof of Lemma 3.16, this results in the backward isogeny cycle of $\Phi(R)$. $\qquad\square$

Corollary 3.18. *There is a bijection between undirected rims of length r of $\mathcal{G}_{K,\ell}$ for all K, and undirected isogeny cycles of length r in $\mathcal{G}_\ell$.*

Definition 3.19. A closed walk in $\mathcal{G}_\ell$ is called a *barbell* if traversing it in the opposite direction as in Lemma 3.16 results in the same walk, up to a choice of basepoint.

An example is shown in Figure 4, justifying the name of 'barbell.'

Lemma 3.20. *Barbells necessarily involve at least two loops and have an even number r of vertices and edges. The total number of barbells having r edges is at most $\#\mathcal{G}_\ell(\ell+1)\ell^{(r-2)/2}$.*

Proof. Any barbell traverses a palindromic sequence of vertices of the following form:

$$j_0, j_1, \ldots, j_{n-1}, j_n, j_n, j_{n-1}, \ldots, j_1, j_0.$$

In particular, the 'turnaround' points must be loops, because otherwise they would involve backtracking. This implies that the number of edges and vertices is even. The count follows by counting the number of walks of length $(r-2)/2$ between two loops, and bounding the number of loops by the number of edges of the graph. $\qquad\square$

The above bound on the number of barbells is far from tight. There are congruence conditions modulo p depending on ℓ which can rule out loops in $\mathcal{G}_\ell$ and hence rule out the existence of barbells. (These conditions can be worked out by considering the specialization of the modular polynomial $\Phi_\ell(x,x) \bmod p$ and choosing p such that the roots are not supersingular.)

Theorem 3.21. *Under the bijection of Theorem 3.2, the self-conjugate rims in $\mathcal{G}_{K,\ell}$ are in bijection with the barbells in $\mathcal{G}_\ell$.*

Proof. If R is a self-conjugate rim, then conjugation returns the same rim traversed in the opposite direction. Hence $\Phi(R)$ must be a barbell. Conversely, if R is a rim such that $\Phi(R)$ is a barbell, then the rim R' directed in the opposite direction maps to $\Phi(R)$, the cycle $\Phi(R)$ traversed in the opposite direction. Hence $\Phi(R) = \Phi(R')$ and therefore R is self-conjugate. $\qquad\square$

4. Supersingular curves oriented by $\mathcal{O}$

This section answers the question of when $\mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}} = \rho(\mathcal{E}\ell\ell(\mathcal{O}))$. Its purpose is to prepare for the next section, where we discuss the fibre sizes of the map taking a rim to its corresponding quadratic order. In particular, we wish to understand the action of conjugation on $\mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$.

Let p be inert or ramified in K, and let $\mathcal{O}$ be an order of K. In this section, we study the set $\mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$ in greater detail, using class field theory. In particular, Theorem 4.4 characterizes when $\rho(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}$, which holds precisely when p ramifies in K. This also characterizes the scenario when conjugation takes $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ back to itself.

Let $K = \mathbb{Q}(\beta)$ be an imaginary quadratic field of discriminant Δ_K with a distinguished root β of a fixed minimal polynomial $f_K(x)$ generating $K/\mathbb{Q}$. Let $\mathcal{O}$ be an order of K. Let L/K be the ring class field of $\mathcal{O}$. By the general theory of class fields, the field L has Galois group $\mathrm{Gal}(L/\mathbb{Q}) \cong \mathrm{Cl}(\mathcal{O}) \rtimes \langle \tau \rangle$ where $\mathrm{Cl}(\mathcal{O}) \cong \mathrm{Gal}(L/K)$, and τ is the non-trivial automorphism of $K/\mathbb{Q}$, and where $\sigma\tau = \tau\sigma^{-1}$ for any $\sigma \in \mathrm{Cl}(\mathcal{O})$ [18, Lemma 9.3]. Let $H_{\mathcal{O}}(x)$ be the Hilbert class polynomial (sometimes called the ring class polynomial) of the order $\mathcal{O}$. Define

$$\mathcal{L}_K := \{j \in L : H_{\mathcal{O}}(j) = 0\}, \quad \mathcal{L}_{\mathbb{Q}} := \{(j, \alpha) \in L \times K : H_{\mathcal{O}}(j) = 0 = f_K(\alpha)\}.$$

The Galois groups $\mathrm{Gal}(L/K)$ and $\mathrm{Gal}(L/\mathbb{Q})$ act freely and transitively on $\mathcal{L}_K$ and $\mathcal{L}_{\mathbb{Q}}$, respectively, via $j \mapsto j^{\sigma}$ and $(j, \alpha) \mapsto (j, \alpha)^{\sigma} = (j^{\sigma}, \alpha^{\sigma})$, respectively. We have $\#\mathcal{L}_K = h_{\mathcal{O}} = \#\mathrm{Gal}(L/K)$ and $\#\mathcal{L}_{\mathbb{Q}} = 2h_{\mathcal{O}} = \#\mathrm{Gal}(L/\mathbb{Q})$.

It is well known that $\mathcal{L}_K$ is in bijection with $\mathcal{E}\ell\ell(\mathcal{O})$, the isomorphism classes of elliptic curves over $\mathbb{C}$ with complex multiplication by $\mathcal{O}$, by the map $j \mapsto E(j)$ taking a j-invariant to its isomorphism class of curves. (We will abuse notation by denoting an element of $\mathcal{E}\ell\ell(\mathcal{O})$ as a curve E without equivalence class brackets $[E]$.) This endows $\mathcal{E}\ell\ell(\mathcal{O})$ with an action of $\mathrm{Gal}(L/K)$.

We now give an analogous result for $\mathcal{L}_{\mathbb{Q}}$. Given $E \in \mathcal{E}\ell\ell(\mathcal{O})$, there are two isomorphisms $\iota \colon \mathcal{O} \to \mathrm{End}_{\mathbb{C}}(E)$ that differ by conjugation. Denote these by ι_E and $\bar{\iota}_E$, where ι_E is normalized so that $\iota_E(\alpha)^*\omega_E = \alpha\omega_E$ for the invariant differential ω_E of E and any $\alpha \in \mathcal{O}$. Let $\mathcal{E}\ell\ell^*(\mathcal{O})$ denote the set of cardinality $2h_{\mathcal{O}}$ of pairs consisting of a curve from $\mathcal{E}\ell\ell(\mathcal{O})$ together with one of the two orientations:

$$\mathcal{E}\ell\ell^*(\mathcal{O}) := \{(E, \iota) \colon E \in \mathcal{E}\ell\ell(\mathcal{O}), \ \iota \colon \mathcal{O} \xrightarrow{\sim} \mathrm{End}(E)\}.$$

Proposition 4.1. *There is a bijection of $\mathcal{L}_{\mathbb{Q}}$ with $\mathcal{E}\ell\ell^*(\mathcal{O})$ via*

$$(j, \delta) \mapsto \left(E = E(j), \begin{cases} \iota_E & \text{if } \delta = \beta \\ \bar{\iota}_E & \text{if } \delta = \bar{\beta} \end{cases} \right).$$

This endows $\mathcal{E}\ell\ell^*(\mathcal{O})$ with an action of $\mathrm{Gal}(L/\mathbb{Q})$.

The bijection is evidently non-canonical; we could make the opposite choice for ι_E.

Choose a finite extension L'/L such that there is a prime $\mathfrak{p}$ of L' above p for which all curves in $\mathcal{E}\ell\ell(\mathcal{O})$ have good reduction at $\mathfrak{p}$. Let $\rho_{\mathfrak{p}}$ denote reduction modulo $\mathfrak{p}$, and put $\tilde{E} = \rho_{\mathfrak{p}}(E)$ for brevity. Reduction provides a map

$$\rho \colon \mathcal{E}\ell\ell^*(\mathcal{O}) \to \mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}} \ \text{ given by } (E, \iota_E) \mapsto (\tilde{E}, \iota_E \bmod \mathfrak{p}), \text{ and } (E, \bar{\iota}_E) \mapsto (\tilde{E}, \bar{\iota}_E \bmod \mathfrak{p}).$$

There is a natural embedding of $\mathcal{E}\ell\ell(\mathcal{O})$ in $\mathcal{E}\ell\ell^*(\mathcal{O})$ given by $E \mapsto (E, \iota_E)$ (i.e. choosing the normalized isomorphism), and to ease notation, we will consider $\mathcal{E}\ell\ell(\mathcal{O})$ a subset of $\mathcal{E}\ell\ell^*(\mathcal{O})$. Consider the following diagram.

$$
\begin{array}{ccc}
\mathrm{Cl}(\mathcal{O}) \cong \mathrm{Gal}(L/K) & \longhookrightarrow & \mathrm{Gal}(L/\mathbb{Q}) \cong \mathrm{Cl}(\mathcal{O}) \rtimes \langle \tau \rangle \\
\curvearrowright \downarrow & & \curvearrowright \downarrow \\
\mathcal{L}_K & & \mathcal{L}_{\mathbb{Q}} \\
\downarrow \wr & & \downarrow \wr \\
\mathcal{E}\ell\ell(\mathcal{O}) & \longhookrightarrow & \mathcal{E}\ell\ell^*(\mathcal{O}) \\
\downarrow \rho_{\mathfrak{p}} & & \downarrow \rho_{\mathfrak{p}} \\
\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) & \longhookrightarrow & \mathrm{SS}^{\mathrm{pr}}_{\mathcal{O}}
\end{array}
$$

The actions of the Galois groups shown are free and transitive. The inclusion map $\mathcal{E}\ell\ell(\mathcal{O}) \hookrightarrow \mathcal{E}\ell\ell^*(\mathcal{O})$ is equivariant with respect to the action of $\mathrm{Gal}(L/K) \subseteq \mathrm{Gal}(L/\mathbb{Q})$.

In the vertical arrows of the lower half of the diagram, we see reduction modulo $\mathfrak{p}$. The square in this portion commutes. The descending arrow on the left is a bijection.

First, we consider the easier unramified case, and show that $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) \neq \mathrm{SS}_{\mathcal{O}}^{pr}$. The map $(E, \iota) \mapsto (E, \bar{\iota})$ on $\mathrm{SS}_{\mathcal{O}}^{pr}$ will be called *conjugation*.

Proposition 4.2. *Let p be inert in K. Then the two isomorphism classes of $\mathcal{O}$-oriented curves represented by (E, ι) and $(E, \bar{\iota})$ cannot both belong to $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$, so $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) \neq \mathrm{SS}_{\mathcal{O}}^{pr}$, in which case $2\#\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = \# \mathrm{SS}_{\mathcal{O}}^{pr}$.*

Proof. Since p is unramified, there exists $\alpha_0 \in \mathcal{O}$ such that $\iota(\alpha_0 - \overline{\alpha_0})$ is separable. Suppose $E_1, E_2 \in \mathcal{E}\ell\ell(\mathcal{O})$. Then there are isomorphisms $[\cdot]_{E_i} : \mathcal{O} \to \mathrm{End}(E_i)$, where we use the normalization of Section 2.3, namely we choose the isomorphisms so that $[\alpha]_{E_i}^* \omega_{E_i} = \alpha \omega_{E_i}$ for all $\alpha \in \mathcal{O}$. With notation as in Section 2.3, suppose that the reduction $(\tilde{E}_1, \iota_{E_1})$ of $(E_1, [\cdot]_{E_1})$ modulo $\mathfrak{p}$ is in the isomorphism class of (E, ι) and the reduction $(\tilde{E}_2, \iota_{E_2})$ of $(E_2, [\cdot]_{E_2})$ is in the isomorphism class of $(E, \bar{\iota})$.

Let us first assume for simplicity that $(\tilde{E}_1, \iota_{E_1}) = (E, \iota)$ and $(\tilde{E}_2, \iota_{E_2}) = (E, \bar{\iota})$, in which case the proof is more transparent. Then, as discussed in Section 2.3, the normalization is preserved by reduction, which is to say,
$$\iota(\alpha)^* \omega_E = \alpha \omega_E = \bar{\iota}(\alpha)^* \omega_E$$
for all $\alpha \in \mathcal{O}$. But then
$$\iota(\alpha)^* \omega_E = \bar{\iota}(\alpha)^* \omega_E = \iota(\overline{\alpha})^* \omega_E,$$
for all $\alpha \in \mathcal{O}$, which implies $\iota(\alpha - \overline{\alpha})^* \omega_E = 0$ for all α, a contradiction for $\alpha = \alpha_0$ since $\iota(\alpha_0 - \overline{\alpha_0})$ is separable.

In general, however, the map ρ is only defined on isomorphism classes, prompting a certain amount of bookkeeping. In particular, assume that $(\tilde{E}_1, \iota_{E_1}) = (\phi_1 E, \phi_1 \iota \phi_1^{-1})$, and $(\tilde{E}_2, \iota_{E_2}) = (\phi_2 E, \phi_2 \bar{\iota} \phi_2^{-1})$ where $\phi_i : E \to \phi_i E$, $i = 1, 2$, are isomorphisms. Then for all $\alpha \in \mathcal{O}$,
$$\alpha \omega_E = \alpha \phi_2^* \omega_{\phi_2 E} = \phi_2^* \alpha \omega_{\phi_2 E} = \phi_2^* (\phi_2 \bar{\iota}(\alpha) \phi_2^{-1})^* \omega_{\phi_2 E} = \bar{\iota}(\alpha)^* \phi_2^* \omega_{\phi_2 E} = \bar{\iota}(\alpha)^* \omega_E = \iota(\overline{\alpha})^* \omega_E$$
and by a similar computation, $\iota(\alpha)^* \omega_E = \alpha \omega_E$. The rest of the proof is as before. $\qquad \square$

Next we consider the ramified case, where we show that $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}_{\mathcal{O}}^{pr}$.

Proposition 4.3. *Suppose p is ramified in K. Then the action of $\mathrm{Gal}(L/\mathbb{Q})$ on $\mathcal{L}_{\mathbb{Q}}$ induces an action on $\mathrm{SS}_{\mathcal{O}}^{pr}$. Under this action, τ acts by conjugation, i.e. $\tau \cdot (E, \iota) = (E, \bar{\iota})$.*

Proof. Define an action on $\mathrm{SS}_{\mathcal{O}}^{pr}$ by $\mathrm{Cl}(\mathcal{O})$ as in Definition 2.22, and by $\langle \tau \rangle$ as $\tau \cdot (E, \iota) = (E, \bar{\iota})$. Let $\mathfrak{P} = \mathfrak{p} \cap L$ be the restriction of $\mathfrak{p}$ to L. Since p is ramified in K, τ is in the inertia group of $\mathfrak{P}$, which implies that $j^\tau \equiv j$ (mod $\mathfrak{P}$) and hence also modulo $\mathfrak{p}$, where $j = j(E)$. Thus
$$\rho_{\mathfrak{p}}(\tau \cdot (E_j, \iota)) = \rho_{\mathfrak{p}}((E_{j^\tau}, \bar{\iota})) = (E_{j^\tau \bmod \mathfrak{p}}, \bar{\iota} \bmod \mathfrak{p}) = (E_{j \bmod \mathfrak{p}}, \bar{\iota} \bmod \mathfrak{p})$$
$$= \tau \cdot (E_{j \bmod \mathfrak{p}}, \iota \bmod \mathfrak{p}) = \tau \cdot \rho_{\mathfrak{p}}((E_j, \iota)).$$

The map ρ is also known to satisfy $\rho_{\mathfrak{p}}(\sigma \cdot (E, \iota)) = \sigma \cdot (\rho_{\mathfrak{p}}((E, \iota)))$ for $\sigma \in \mathrm{Cl}(\mathcal{O})$ [42, Proof of Proposition 3.6]. Therefore, the actions described on $\mathrm{SS}_{\mathcal{O}}^{pr}$ give a well-defined action of $\mathrm{Gal}(L/\mathbb{Q})$ with respect to which ρ is equivariant. $\qquad \square$

Theorem 4.4. *The prime p ramifies in K if and only if $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}_{\mathcal{O}}^{pr}$ if and only if conjugation takes $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$ to itself.*

Proof. If $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}_{\mathcal{O}}^{pr}$, then p ramifies by Proposition 4.2. Conversely, if p ramifies in K, then there is some $\theta \in \iota(\mathcal{O}) \subseteq \mathrm{End}(E)$ such that $\theta^2 = -np$ for some integer n coprime to p. The endomorphism θ is a horizontal isogeny with respect to $\mathcal{O}$ (as an element of $\iota(\mathcal{O})$). Furthermore, θ must factor as
$$E \xrightarrow{\pi_p} E^{(p)} \xrightarrow{\phi} E,$$
where ϕ is of degree n and separable. But since $\theta \in \iota(\mathcal{O})$, $\theta_* \iota = \iota$ and we have the following chain of K-oriented isogenies:
$$(E, \iota) \xrightarrow{\pi_p} (E^{(p)}, \iota^{(p)}) \xrightarrow{\phi} (E, \iota).$$
The maps shown above are horizontal isogenies (π_p is horizontal and so is θ, hence ϕ is also). We wish to show that ϕ represents the action of an element of $\mathrm{Cl}(\mathcal{O})$, up to multiplication by an integer. We can

factor n into prime powers and prove it for each one, so without loss of generality, we assume that n is a prime power q^k, for $q \neq p$.

The map ϕ is given by a path in the oriented q-isogeny graph determined by successive kernels, and as it is horizontal, this path must have as many descending as ascending steps. Using the volcano structure of the graph, all the ascending and descending edges can be paired up into consecutive pairs $\phi_2 \circ \phi_1$ where ϕ_1 descends and ϕ_2 ascends. Furthermore, $\phi_2 = \varphi_1 \widehat{\phi_1} \varphi_2$ for some automorphisms φ_1, φ_2. Then $\varphi_1^{-1} \phi_2 \phi_1 = \widehat{\phi_1} \varphi_2 \phi_1$ has trace and norm divisible by q, hence it is divisible by $[q]$. So $\phi_2 \phi_1 = \varphi[q]$ for some automorphism φ. Since we only care about ϕ up to multiplication by an integer, we can replace $\phi_2 \phi_1$ with φ, which is horizontal. Repeating this process leaves a purely horizontal path, which must therefore represent an action of an element of $\mathrm{Cl}(\mathcal{O})$ (since $q \neq p$, $\mathrm{Cl}(\mathcal{O})$ acts on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$, so any horizontal q-isogeny is the action of an element of the class group). It follows that $(E^{(p)}, \iota^{(p)}) \in \mathrm{Cl}(\mathcal{O}) \cdot (E, \iota)$, which in turn implies that $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ by [42, Propositions 3.3 and 3.4].

Finally, we show that $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ if and only if conjugation takes $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$ to itself. Since p is ramified, by Proposition 4.3, the action of $\mathrm{Gal}(L/\mathbb{Q})$ is defined on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ and τ acts by conjugation, fixing the j-invariant $j(E)$ modulo $\mathfrak{p}$. First, note that if (E, ι) and $(E, \bar{\iota})$ are related by an element of the class group, then the entire orbit $\mathrm{Cl}(\mathcal{O}) \cdot (E, \iota)$ consists of elements related to their conjugates by the class group, and therefore, conjugation takes $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$ into itself. Therefore conjugation takes $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$ to itself, or else maps it to $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} \backslash \rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$. We know that $\#\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O})) = h_{\mathcal{O}}$, and $\#\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} \in \{h_{\mathcal{O}}, 2h_{\mathcal{O}}\}$ (Section 2.3). But $\#\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} = 2h_{\mathcal{O}}$ if and only if $\rho_{\mathfrak{p}} \colon \mathcal{E}\ell\ell^*(\mathcal{O}) \to \mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ is an isomorphism if and only if $\mathrm{Gal}(L/\mathbb{Q})$ acts freely and transitively on $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ (calling upon Proposition 4.3 again). But $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$ forms one $\mathrm{Cl}(\mathcal{O})$-orbit, so the only way the action can be free and transitive is if conjugation takes $\rho_{\mathfrak{p}}(\mathcal{E}\ell\ell(\mathcal{O}))$ outside itself. $\square$

5. ISOGENY CYCLES AND THEIR ASSOCIATED ORDERS

The following theorem describes a way to relate isogeny cycles in $\mathcal{C}_r$ to imaginary quadratic orders. Define

$$
\mathcal{I}_r = \left\{ \text{imaginary quadratic orders } \mathcal{O} : \begin{array}{c} p \text{ does not split in the field containing } \mathcal{O} \\ p \text{ does not divide the conductor of } \mathcal{O} \\ \mathcal{O} \text{ is an } \ell\text{-fundamental order,} \\ (\ell) = \bar{\mathfrak{l}}\mathfrak{l} \text{ splits in } \mathcal{O}, \\ \text{and } [\mathfrak{l}] \text{ has order } r \text{ in } \mathrm{Cl}(\mathcal{O}). \end{array} \right\}.
$$

These are exactly the orders for which $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$ is non-empty and the permutation on the vertices arising from the action of a prime ideal in $\mathcal{O}$ above ℓ has a decomposition consisting of cycles of length r only. The following theorem describes how often each order in $\mathcal{I}_r$ is obtained from an isogeny cycle in $\mathcal{G}_\ell$ via the bijection of Theorem 3.2. If p is inert, one expects to obtain $2h_{\mathcal{O}}/r$ directed isogeny cycles of size r which give rise to an order $\mathcal{O}$, since $\#\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}} = 2h_{\mathcal{O}}$ and the rims are identified in pairs by conjugation, and given a direction. However, in the case that p ramifies in the field K containing $\mathcal{O}$, the exact count may differ from this.

Let $\mathcal{O}$ be a quadratic order with fraction field K. Let L/K be the ring class field of $\mathcal{O}$. Let $g_{\mathcal{O}}$ be the number of genera in $\mathrm{Cl}(\mathcal{O})$. Let $\sigma_\ell := \mathrm{Frob}_\ell \in \mathrm{Gal}(L/K)$ have order r. Define $C_\ell := \langle \sigma_\ell \rangle$ (so $\#C_\ell = r$).

Theorem 5.1. *Let $r > 2$. The map $\mathcal{R}_r \to \mathcal{I}_r$ which takes a rim to the order $\mathcal{O}$ with respect to which all its vertices are primitively $\mathcal{O}$-oriented as in the bijection of Theorem 3.2 has fibres of size $2h_{\mathcal{O}}/r$ or $4h_{\mathcal{O}}/r$, depending on whether or not p ramifies in the field K containing $\mathcal{O}$, respectively. The map factors through $\mathcal{R}_r/\sim$, so the induced map $\mathcal{C}_r \to \mathcal{I}_r$ has a fibre of size $\epsilon_{\mathcal{O},\ell} h_{\mathcal{O}}/r$ over $\mathcal{O}$ for some $1 \leq \epsilon_{\mathcal{O},\ell} \leq 2$. Furthermore,*

(1) Suppose p is inert K. Then $\epsilon_{\mathcal{O},\ell} = 2$.
(2) Suppose p ramifies in K. Then there are two possibilities:
  (a) There are no self-conjugate rims in $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$. Then $\epsilon_{\mathcal{O},\ell} = 1$.
  (b) There are self-conjugate rims in $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$, in which case r is even, and

$$
\epsilon_{\mathcal{O},\ell} := 1 + \frac{rg_{\mathcal{O}}}{2h_{\mathcal{O}}}.
$$

  In this case, $1 < \epsilon_{\mathcal{O},\ell} \leq 2$.

The following commutative diagram illustrates the theorem:

(9)
$$
\mathcal{C}_r \xrightarrow[\text{Thm. 3.2}]{\cong} \mathcal{R}_r/\sim \xleftarrow[\substack{\text{fibres} \\ \text{of size} \\ \text{1 or 2}}]{} \mathcal{R}_r
$$

with $\mathcal{R}_r$ mapping down (fibres of size $2h_{\mathcal{O}}/r$ or $4h_{\mathcal{O}}/r$) to $\mathcal{I}_r$.

Proof. Define a map $\mathcal{R}_r \to \mathcal{I}_r$ taking a rim to the associated order as described in the statement of the theorem. The size of the fibre above $\mathcal{O}$ is $2\#\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}/r$ (the factor of 2 arises because rims are directed), which is $2h_{\mathcal{O}}/r$ or $4h_{\mathcal{O}}/r$ in the ramified and inert cases, respectively, by Theorem 4.4 (the conditions on p in the definition of $\mathcal{I}_r$ imply $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$ is nonempty by Proposition 2.19). Since a conjugate rim is associated to the same order as the original, the map factors through $\mathcal{R}_r/\sim$. The fibres of $\mathcal{R}_r \to \mathcal{R}_r/\sim$ are of size 1 or 2 (depending upon whether the rim in question was self-conjugate).

In the inert case, $\#\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}} = 2h_{\mathcal{O}}$ by Theorem 4.4. The fibres of $\mathcal{R}_r \to \mathcal{R}_r/\sim$ are uniformly of size 2 for rims in $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$ (since conjugation swaps $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ and its complement in $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$ by Theorem 4.4). The map $\mathcal{R}_r \to \mathcal{I}_r$ has a fibre of size $4h_{\mathcal{O}}/r$ over $\mathcal{O}$. Hence the map $\mathcal{C}_r \to \mathcal{I}_r$ has a fibre of size $2h_{\mathcal{O}}/r$ above $\mathcal{O}$.

Now we consider the ramified case. Call upon the conventions and set-up of Section 4. Let $G = \operatorname{Gal}(L/\mathbb{Q})$. We consider the stabilizers $\operatorname{Stab}_G((E,\iota))$ of elements of $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$. All stabilizers are of size 2 since $\#G = 2h_{\mathcal{O}} = 2\#\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$ by Theorem 4.4, and no stabilizers are contained in $\operatorname{Cl}(\mathcal{O})$ (whose action is free). Therefore, for some (E,ι), we have $\operatorname{Stab}_G((E,\iota)) = \{id, \sigma\tau\}$, for some $\sigma \in \operatorname{Cl}(\mathcal{O}) \cong \operatorname{Gal}(L/K)$. The set of stabilizers is the orbit of $\operatorname{Stab}_G((E,\iota))$ under conjugation by elements of $G \cong \operatorname{Cl}(\mathcal{O}) \rtimes \langle\tau\rangle$. For $\eta \in \operatorname{Cl}(\mathcal{O})$, we have

$$
\eta\operatorname{Stab}_G((E,\iota))\eta^{-1} = \{id, \eta\sigma\tau\eta^{-1}\} = \{id, \eta^2\sigma\tau\}.
$$

Also $\tau\operatorname{Stab}_G((E,\iota))\tau^{-1} = \{id, \tau\sigma\} = \{id, \sigma^{-2}\sigma\tau\}$. Thus, the nontrivial elements of the stabilizers, taken together, form a set $\operatorname{Cl}(\mathcal{O})^2\sigma\tau$. Set the notation $G_\sigma := \operatorname{Cl}(\mathcal{O})^2\sigma$. Then $\sigma \in \operatorname{Cl}(\mathcal{O}) \setminus \operatorname{Cl}(\mathcal{O})^2$, as τ does not stabilize any pairs (E,ι). Hence G_σ is a non-principal genus of $\operatorname{Cl}(\mathcal{O})$.

An element (E,ι) is in the same ℓ-isogeny rim as its conjugate $(E,\bar{\iota})$ if and only if $\sigma_\ell^k\tau \in \operatorname{Stab}_G((E,\iota))$ for some integer k if and only if $C_\ell\tau \cap \operatorname{Stab}_G((E,\iota)) = C_\ell\tau \cap G_\sigma\tau \neq \emptyset$. If $C_\ell \cap G_\sigma = \emptyset$, then there are no stabilizers intersecting $C_\ell\tau$, so there are no self-conjugate rims. Since $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ is taken to itself by conjugation, in this case all rims pair up in conjugate pairs. Therefore the fibres are of size $h_{\mathcal{O}}/r$.

Suppose there are self-conjugate rims, in which case $\sigma_\ell^k \in G_\sigma$ for some integer k. Then k must be odd, as otherwise $\sigma_\ell^k \in G_\sigma \cap \operatorname{Cl}(\mathcal{O})^2 = \emptyset$. It follows that $\sigma_\ell \in G_\sigma$. A similar argument forces r to be even, as otherwise $\sigma_\ell = \sigma_\ell^{r+1} \in G_\sigma \cap \operatorname{Cl}(\mathcal{O})^2 = \emptyset$.

There are $\#G_\sigma = \#\operatorname{Cl}(\mathcal{O})^2 = h_{\mathcal{O}}/g_{\mathcal{O}}$ distinct stabilizers, of which $\#G_\sigma \cap C_\ell$ belong to elements in the same rim as their conjugates. By the same logic as before, using the fact that r is even, the intersection of G_σ with C_ℓ consists of the odd powers of σ_ℓ, and is therefore of size $\#C_\ell/2 = r/2$. Since all rims are the same size, this means the proportion of rims which are self-conjugate is

$$
c := \frac{r/2}{\#\operatorname{Cl}(\mathcal{O})^2} = \frac{rg_{\mathcal{O}}}{2h_{\mathcal{O}}}.
$$

This represents the proportion of rims (equivalently, of vertices) of $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$ that are self-conjugate. This proportion is the same when we restrict to $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$.

The proportion $1 - c$ of non-self-conjugate rims is double-counted when we wish to count ℓ-isogeny cycles in $\mathcal{G}_\ell$, so we have

$$
\epsilon_{\mathcal{O},\ell} = 2c + (1 - c) = 1 + c.
$$

For the final inequality, note that in the ramified case, the fibres of $\mathcal{R}_r \to \mathcal{I}_r$ are of size $2h_{\mathcal{O}}/r$, and reference diagram (9). $\qquad\square$

Example 5.2. The following example illustrates the most complicated type of fibre. Consider $\mathcal{G}_\ell$ with $p = 241$, $\ell = 11$, and the quadratic order $\mathcal{O}$ of discriminant $-4p$, namely $\mathcal{O} = \mathbb{Z}[\sqrt{-p}]$. Since p ramifies, $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}} = \rho(\mathcal{E}\ell\ell(\mathcal{O}))$ by Theorem 4.4. Let $r = 4$. The class group of $\mathcal{O}$ is cyclic of order 12. Therefore, there are $12/4 = 3$ rims of size 4 in $\operatorname{SS}_{\mathcal{O}}^{\mathrm{pr}}$ under the action of a prime ideal $\mathfrak{l}$ in $\mathcal{O}$ above ℓ. That means 6 directed rims. Using an isomorphism $(\operatorname{Cl}(\mathcal{O}), \cdot) \cong (\mathbb{Z}/12\mathbb{Z}, +)$ which takes σ_ℓ of order 4 to the element 3, we have $\operatorname{Cl}(\mathcal{O})^2 = \{0, 2, 4, 6, 8, 10\}$ and $C_\ell = \{0, 3, 6, 9\}$. So $1/3$ of the elements of $\operatorname{Cl}(\mathcal{O})^2$ are in C_ℓ. This implies that

FIGURE 5. An illustration of Example 5.2 with $p = 241$, $\ell = 11$ and $\mathcal{O} = \mathbb{Z}[\sqrt{-p}]$, for which $h_{\mathcal{O}} = 12$. On the left, the rims of $\mathcal{G}_{\mathbb{Q}(\sqrt{-p}),11}$ drawn in such a way that conjugation is a mirror symmetry through the dotted line. On the right, the subgraph of $\mathbb{F}_p$-vertices of $\mathcal{G}_{11}$.

$1/3$ of the total directed rims are self-conjugate: two of the six (one of the three undirected rims). Using the formula of Theorem 5.1, with $g_{\mathcal{O}} = 2$, we obtain

$$\epsilon_{\mathcal{O},\ell} = 1 + \frac{4 \cdot 2}{2 \cdot 12} = \frac{4}{3},$$

which gives a fibre size of $\frac{4}{3} \cdot \frac{12}{4} = 4$ for the map $\mathcal{C}_4 \to \mathcal{I}_4$ above $\mathcal{O}$. In other words, there should be four directed isogeny cycles of size 4 (equivalently, two undirected isogeny cycles) in $\mathcal{G}_{11}$ which give rise to $\mathcal{O}$.

The set $\rho(\mathcal{E}\ell\ell(\mathcal{O}))$ has three undirected length 4 isogeny cycles: $(64, 93, 216, 240)$ (twice, conjugate to one another) and $(8, 8, 28, 28)$ (self-conjugate). The subgraph of $\mathcal{G}_{11}$ of $\mathbb{F}_p$-points is shown in Figure 5, where we observe that there is only one 4-cycle through $(964, 93, 216, 240)$, and that $(8, 8, 28, 28)$ collapses to a 'barbell shaped' isogeny cycle. The commutative diagram (9), if labelled with the (undirected) elements above $\mathbb{Z}[\sqrt{-p}]$, looks something like the following (please compare to Figure 5):

6. AN EXTENDED EXAMPLE

In this section, we give an example that demonstrates our bijection (Theorem 3.2). We choose $p = 179$, $\ell = 2$, and we consider $r = 3$, 4, 5 and 6. Computing the set $\mathcal{R}_r$ involves first finding all ℓ-fundamental imaginary quadratic orders $\mathcal{O}$ such that $\mathcal{O}$ embeds into $B_{p,\infty}$, the ideal $(\ell) = \mathfrak{l}\bar{\mathfrak{l}}$ splits in $\mathcal{O}$ and $[\mathfrak{l}]$ has order r in $\mathrm{Cl}(\mathcal{O})$. A complete list of such orders $\mathcal{O}$ and prime ideals $\mathfrak{l}, \bar{\mathfrak{l}}$ can be calculated as follows:

  (1) list all representations of ℓ^r by norm forms of sufficiently small discriminant;
  (2) each representation gives rise to an imaginary quadratic element α such that $N(\alpha) = \ell^r$ (up to sign and conjugation);
  (3) check the splitting behaviour of p in $\mathbb{Q}(\alpha)$, discarding α if p splits in $\mathbb{Q}(\alpha)$;

(4) consider all ℓ-fundamental orders in $\mathbb{Q}(\alpha)$ that contain α and check whether the ideal class $[\mathfrak{l}]$ has order r in $\mathrm{Cl}(\mathcal{O})$, obtaining a list L_r of desirable orders $\mathcal{O}$.

As an example, consider representing $2^3 = 8$. We need only check discriminants Δ with $|\Delta| \leq 4 \cdot 8 = 32$. We find two representations: $1^2 + 31 \cdot 1^2 = 4 \cdot 8$ and $3^2 + 23 \cdot 1^2 = 4 \cdot 8$. These gives rise to two possible α up to conjugation and sign:

$$\frac{\pm 1 \pm \sqrt{-31}}{2}, \frac{\pm 3 \pm \sqrt{-23}}{2}.$$

However, 179 splits in $\mathbb{Q}(\sqrt{-23})$ and not in $\mathbb{Q}(\sqrt{-31})$. Hence we discard the second α. We find that $\mathfrak{l} = (2, \frac{1+\sqrt{-31}}{2})$ is not principal and $\mathfrak{l}^3 = (\alpha)$. Hence $L_3 = \left\{ \mathbb{Z}\left[\frac{1+\sqrt{-31}}{2}\right] \right\}$.

Each element α corresponds to an orientation ι_α (up to conjugation), which gives a rim (up to conjugation and direction), as in the map Θ of Lemma 3.9. Such a list therefore gives the number of undirected rims: $\#L_r = \frac{1}{2}\#(\mathcal{R}_r/\sim)$.

Theorem 3.2 concerns directed isogeny cycles and rims, but in this example we will ignore the issue of direction, which is allowed by Corollary 3.18. Therefore, we use L_r as a proxy for $\mathcal{R}_r/\sim$ for computational purposes. We compare L_r with $\mathcal{C}_r$, verifying the bijection theorem.

For each $r \in \{3, 4, 5, 6\}$, we compute the set $\mathcal{C}_r$ on the supersingular 2-isogeny graph by reference to Figure 5 showing $\mathcal{G}_2$. We will actually consider $\mathcal{C}'_r$, the set of isogeny cycles of $\mathcal{C}_r$ but forgetting direction, allowed by Corollary 3.18. Then we determine the associated endomorphisms by composing around the cycles as in the map Θ; we do this by computing the trace of the composed endomorphism θ using the formulas $\theta + \widehat{\theta}$ or $1 + \deg(\theta) - \deg(1 - \theta)$ in SageMath [53].

For example, there is only one isogeny cycle of length 3 (in the sense of Theorem 3.2, but ignoring direction) in $\mathcal{C}'_3$ for $\mathcal{G}_2$, namely $(j_3, \overline{j_3}, 171)$. Composing around the triangle, one finds an endomorphism of trace ± 1. Since the endomorphism has norm 8, it has minimal polynomial $x^2 \pm x + 8$, and hence is of the form $\frac{\pm 1 \pm \sqrt{-31}}{2}$. This generates an order with class number 3 (the class numbers were computed with SageMath).

In Table 1, we give the lists $\mathcal{C}'_r$, and L_r, and put them in correspondence. This represents the first row of the table which puts $\mathcal{C}'_3 = \{(j_3, \overline{j_3}, 171)\}$ and $L_3 = \left\{ \mathbb{Z}\left[\frac{1+\sqrt{-31}}{2}\right] \right\}$ (as a proxy for $\mathcal{R}_3/\sim$ modulo direction) in bijection.

The table also verifies Corollary 7.3. In our case, we have observed no fields ramified at p, so $\epsilon_{\mathcal{O},\ell} = 2$. We compute:

$$2 = 2\#\mathcal{C}'_3 = \#\mathcal{C}_3 = \frac{2}{3}h\left(\mathbb{Z}\left[\frac{1+\sqrt{-31}}{2}\right]\right) = \frac{2}{3} \cdot 3 = 2.$$

$$2 = 2\#\mathcal{C}'_4 = \#\mathcal{C}_4 = \frac{2}{4}h\left(\mathbb{Z}\left[\frac{1+\sqrt{-39}}{2}\right]\right) = \frac{2}{4} \cdot 4 = 2.$$

$$2 = 2\#\mathcal{C}'_5 = \#\mathcal{C}_5 = \frac{2}{5}h\left(\mathbb{Z}\left[\frac{1+\sqrt{-47}}{2}\right]\right) = \frac{2}{5} \cdot 5 = 2.$$

$$14 = 2\#\mathcal{C}'_6 = \#\mathcal{C}_6 = \frac{2}{6}\left(h\left(\mathbb{Z}\left[\frac{1+\sqrt{-87}}{2}\right]\right) + h\left(\mathbb{Z}\left[\frac{1+\sqrt{-231}}{2}\right]\right) + \right.$$
$$\left. h\left(\mathbb{Z}\left[\frac{1+\sqrt{-247}}{2}\right]\right) + h\left(\mathbb{Z}\left[\frac{1+\sqrt{-255}}{2}\right]\right) + h\left(\mathbb{Z}\left[3\frac{1+\sqrt{-15}}{2}\right]\right) \right)$$
$$= \frac{2}{6} \cdot (6 + 12 + 6 + 12 + 6) = 14.$$

This verifies Corollary 7.3 in this example.

The cycle $(0, 121, 112, 35, 112, 121)$. The case of the j-invariant cycle $(0, 121, 112, 35, 112, 121)$ in Figure 6 (E) is particularly interesting. This sequence of j-invariants forms a 6-cycle in 12 possible ways, as follows. There are 4 ways to travel from 112 to 35 and back. There are 3 ways to travel from 121 to 0 and back, producing three distinct endomorphisms of $j = 121$ (but not canonically), each composing to one of ± 2, $\pm\widehat{\varphi} \circ \zeta_3 \circ \varphi$, or $\pm\widehat{\varphi} \circ \zeta_3^2 \circ \varphi$, as described in Figure 3. Thus, one endomorphism is backtracking and the other two are duals of each other. The isogenies from 112 to 35 each have a dual: hence two of the 2-cycles from 112 to 35 and back are backtracking (obtained by pairing an isogeny with its dual). There

isogeny cycle	figure	length	endomorphism	order	class number
$(j_3, \overline{j_3}, 171)$	Fig. A	3	$\frac{\pm 1 \pm \sqrt{-31}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-31}}{2}\right]$	3
$(61, j_1, 140, \overline{j_1})$	Fig. A	4	$\frac{\pm 5 \pm \sqrt{-39}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-39}}{2}\right]$	4
$(22, \overline{j_2}, \overline{j_3}, j_3, j_2)$	Fig. A	5	$\frac{\pm 9 \pm \sqrt{-47}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-47}}{2}\right]$	5
$(22, \overline{j_2}, \overline{j_1}, 140, j_1, j_2)$	Fig. B	6	$\frac{\pm 13 \pm \sqrt{-87}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-87}}{2}\right]$	6
$(140, j_1, j_2, j_3, 171, 120)$ $(140, \overline{j_1}, \overline{j_2}, \overline{j_3}, 171, 120)$	Fig. C Fig. D	6	$\frac{\pm 5 \pm \sqrt{-231}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-231}}{2}\right]$	12
$(0, 121, 112, 35, 112, 121)^*$	Fig. E	6	$\frac{\pm 3 \pm \sqrt{-247}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-247}}{2}\right]$	6
$(22, j_2, j_3, 171, \overline{j_3}, \overline{j_2})$ $(0, 121, 112, 35, 112, 121)^*$	Fig. F Fig. E	6	$\frac{\pm 1 \pm \sqrt{-255}}{2}$	$\mathbb{Z}\left[\frac{1+\sqrt{-255}}{2}\right]$	12
$(61, j_1, j_2, 22, \overline{j_2}, \overline{j_1})$	Fig. G	6	$\frac{\pm 11 \pm 3\sqrt{-15}}{2}$	$\mathbb{Z}\left[3\left(\frac{1+\sqrt{-15}}{2}\right)\right]$	6

TABLE 1. Cycles of lengths three through six, with the associated endomorphisms to which the cycles compose. The two cycles labelled with an asterisk $*$ are not fully specified by giving their j-invariants. However, as discussed in Section 3.2, the bijection between these two cycles and the two associated α is not canonical, but if we choose a safe arbitrary assignment, we can make a non-canonical bijection. So there is no reason to specify which cycle is which in the table here.

are therefore two non-backtracking 2-cycles to choose from among the four ways to travel from 112 to 35 and back. These two non-backtracking 2-cycles are actually duals of each other, and are obtained by simply changing the direction of traversal. If the traces of the paths that include backtracking are computed, we obtain the associated endomorphisms ± 8, $4(\pm 1 \pm \sqrt{-3})$, and $2(\pm 1 \pm \sqrt{-15})$, all of which are divisible by 2. This leaves exactly four non-backtracking 6-cycles on the j-invariants $(0, 121, 112, 35, 112, 121)$, coming in pairs obtained by changing the direction of traversal. These two pairs give the two endomorphisms listed in Table 1.

The cycle $(61, j_1, j_2, 22, \overline{j_2}, \overline{j_1})$**.** Another interesting cycle is $(61, j_1, j_2, 22, \overline{j_2}, \overline{j_1})$ in Figure 6 (G). This cycle corresponds to an element generating the non-maximal order $\mathcal{O}_3 := \mathbb{Z}\left[3\left(\frac{1+\sqrt{-15}}{2}\right)\right]$. That element also exists in the order $\mathcal{O}_1 := \mathbb{Z}\left[\frac{1+\sqrt{-15}}{2}\right]$. How do we know which order the volcano rim corresponds to? The easiest answer is that only the class number of $\mathcal{O}_3$ is divisible by 6. However, even without that observation (sometimes several class numbers may be divisible by the correct integer r), we can observe that the orientation ι given by the endomorphism on any of the j-invariants on the cycle is $\mathcal{O}_3$-primitive, not $\mathcal{O}_1$-primitive. This can be seen because if $\text{End}(E)$ contained $\frac{1+\sqrt{-15}}{2}$, which has norm 16, then this element would be realized as a 4-cycle from E, which it is not. (The curves $j_2, \overline{j_2}$, and 22 have no 4-cycle. The curves 61, j_1 and $\overline{j_1}$ have the 4-cycle associated to $\frac{\pm 5 \pm \sqrt{-39}}{2}$.) The moral of this observation is that each cycle can only arise from a single quadratic order, as dictated by Theorem 3.2.

Repeating j-invariants. A final remark is the observation that repeating j-invariants in a cycle can only happen when the discriminant is large with respect to the prime. This is a consequence of a result by Kaneko which gives the necessary condition $|\text{disc}(\mathcal{O})| \geq p$ [30, Theorem 2']. In our case, we do get a cycle with repeating j-invariants, but the associated discriminants are quite large at -247 and -255, both bigger in absolute value than the prime $p = 179$.

(A)

(B)   (C)   (D)

(E)   (F)   (G)

FIGURE 6. Figure (A) is the supersingular 2-isogeny graph $\mathcal{G}_2$ over $\mathbb{F}_{p^2}$, where $p = 179$. Here $j_1 = 64i + 5$, $j_2 = 99i + 107$, $j_3 = 5i + 109$, where i denotes a square root of -1 in $\mathbb{F}_{p^2}$. Figures (B) – (G) are the isogeny cycles of length 6 as listed in Table 1. There are two distinct isogeny cycles included in (E), both sharing the same sequence of vertices.

Comparison with Gross [28] Working with the same example, Gross counts cycles of length $\log_\ell(m)$ in the supersingular ℓ-isogeny graph $\mathcal{G}_\ell$ via the formula provided in [28, Proposition 1.9]. In the first row of Table 1, we count one isogeny cycle of length three in the 2-isogeny graph, namely $(j_3, \overline{j_3}, 171)$ arising from the order of discriminant -31. Gross counts nine via the formula [28, Proposition 1.9]. In particular, six of this count come from the order of discriminant -31. The three-cycle $(j_3, \overline{j_3}, 171)$, which we count once, is counted six times by Gross because it is a valid three-cycle starting at any of $j_3, 171, \overline{j_3}$, in either direction. Two more of Gross's count come the order of discriminant -16. The cycle corresponding to the order of discriminant -16 is necessarily $(j_{22}, j_{1728}, j_{1728})$. This cycle is only without backtracking when we start at

22, so Gross counts it once in each direction. We do not count this cycle, as it contains backtracking as soon as we forget the basepoint. The last one of Gross's count comes from the order of discriminant -4, as j_{1728} has an endomorphism of degree-8 corresponding to composing the loop at j_{1728} three times, with appropriate automorphisms in between to avoid backtracking.

7. COUNTING ISOGENY CYCLES IN $\mathcal{G}_\ell$

Theorem 5.1 makes possible explicit formulas for the number of isogeny cycles of length r in the ℓ-isogeny graph $\mathcal{G}_\ell$. In this section, we do the following:

(1) Theorem 7.1 uses spectral graph theory to give an asymptotic for the expected number of length r isogeny cycles, and this estimate is compared to empirical data in Figure 7. The use of spectral graph theory requires us to restrict to $p \equiv 1 \pmod{12}$, but for any large enough p we still expect the estimates of this theorem to hold.

(2) Corollary 7.3 gives an explicit formula for the number of length r isogeny cycles in terms of class numbers of certain quadratic orders, and Theorem 7.4 reformulates this as a sum over these class numbers with simpler conditions on which quadratic orders to include.

(3) Theorem 7.4 enables an explicit upper bound on the length r isogeny cycle count, given in Corollary 7.7. Unfortunately a lower bound would seem to be connected to an open problem in number theory (Remark 7.8).

7.1. **Asymptotic number of isogeny cycles.** It is a standard fact that for $(\ell + 1)$-regular Ramanujan graphs, the number of basepointed non-backtracking closed walks of length r (where the non-backtracking condition does not rule out entering the final vertex along the initial edge of the walk) is

$$(\ell+1)\ell^{r-1} + \#\mathcal{G}_\ell\, O\left(r(\ell+1)^{r/2}\right),$$

where the implied constant does not depend on $\#\mathcal{G}_\ell$. This follows, for example, from [19, Theorem 1.4.6]; some details are provided at [8]. If we disallow entering the final vertex along the initial edge and then forget basepoints, then one would heuristically expect this to become

(10) $$\ell^r/2r + \#\mathcal{G}_\ell\, O\left((\ell+1)^{r/2}\right).$$

For small r, estimating the number of graph-theoretic r-cycles is similar (recall that a cycle allows no repeated edges or vertices); this is because random d-regular graphs on n vertices have on average $(d-1)^r/2r$ such r-cycles [9, Theorem 2]. In our case, this would give $\ell^r/2r$. The number of r-cycles in some Ramanujan graphs, including LPS graphs, but not supersingular graphs, has been analysed, and these graphs appear to behave as random graphs [54].

In what follows, we give a brief proof of an asymptotic version of (10).

Theorem 7.1. *Let $p \equiv 1 \pmod{12}$. Let $\mathcal{G}_\ell$ be the supersingular ℓ-isogeny graph over $\overline{\mathbb{F}}_p$ and let $\#\mathcal{G}_\ell$ denote its number of vertices. The number of non-backtracking closed walks of length r in $\mathcal{G}_\ell$, taken up to order of traversal and starting point, asymptotically approaches $\ell^r/2r$ as $r \to \infty$. Thus, the number of isogeny cycles in $\mathcal{G}_\ell$ is also asymptotically $\ell^r/2r$ as $r \to \infty$.*

Before giving a formal proof, we give a heuristic description. Choose one of $\#\mathcal{G}_\ell$ initial vertices j_0. Label its $\ell + 1$ neighbours $j_1, j_2, \ldots, j_{\ell+1}$. Choose one of these neighbours for the first step, say j_1 (relabelling without loss of generality). Consider the ℓ^{r-1} non-backtracking length $r - 1$ walks that could follow the edge $j_0 \to j_1$. Such a walk provides a non-backtracking closed walk of length r through $j_0 \to j_1$ if and only if the walk terminates at j_0 via one of $j_2, \ldots, j_{\ell+1}$. The probability of this termination vertex j_0 should be expected to be $1/\#\mathcal{G}_\ell$. However, only ℓ of the $\ell + 1$ incoming edges reach j_0 via one of $j_2, \ldots, j_{\ell+1}$. Finally, we obtain a typical closed non-backtracking walk of length r exactly $2r$ times in this way, since there are r choices of initial vertex and 2 directions. Hence the number of closed non-backtracking walks of length r should be obtained by combining the steps in the outline above:

$$\#\mathcal{G}_\ell(\ell+1)\ell^{r-1}(1/\#\mathcal{G}_\ell)(\ell/(\ell+1))/2r = \ell^r/2r.$$

Proof of Theorem 7.1. The natural object for studying non-backtracking walks is a Markov chain on the *non-backtracking matrix* B, which is the adjacency matrix of an auxiliary graph $\mathcal{G}'$ whose vertices are the directed edges of $\mathcal{G}_\ell$, and whose edges indicate adjacency between edges of $\mathcal{G}_\ell$ that respects the directedness

FIGURE 7. A semilog plot of cycle counts for $\mathcal{G}_\ell$ for $p = 3361$ and $\ell = 2$ (green squares, lower graph), and for $p = 3229$ and $\ell = 3$ (red circles, upper graph). The black curves upon which the data points appear to sit are $\ell^x/(2x)$. The blue lines above the data in each case indicate the dominant term of Corollary 7.7, indicating our proved upper bound. On the left, the full data. On the right, a close-up of the behaviour for small cycle sizes. The graphs were generated in part using code from [4].

and does not backtrack. In other words, the matrix is a scaling of the transition probability matrix of a non-backtracking random walk on the directed edges of $\mathcal{G}_\ell$. In particular, the entries of B^r give the number of non-backtracking walks of length $r+1$ that start and end on certain edges. This is studied in [31], where it is shown that this Markov chain converges to a uniform stationary distribution (the mixing property). There are $\#\mathcal{G}_\ell(\ell+1)$ directed edges in $\mathcal{G}_\ell$. Choose one of these as the initial directed edge of a walk, originating at a vertex which we may call j_0. From that point, take a random non-backtracking walk of length $r-1$ (there are ℓ^{r-1} of these). Then we have created a closed non-backtracking walk if and only if we end on one of the ℓ edges that are directed into j_0 without backtracking the initial edge. The mixing property implies that the probability of ending on of these ℓ incoming edges is asymptotically $\ell/(\#\mathcal{G}_\ell(\ell+1))$. Therefore, asymptotically, we obtain

$$\#\mathcal{G}_\ell(\ell+1)\ell^{r-1}\left(\frac{\ell}{\#\mathcal{G}_\ell(\ell+1)}\right) = \ell^r$$

total non-backtracking closed walks with a direction and a basepoint. Most of these can be directed in two ways and a basepoint can be chosen in r ways. The exceptions, namely the barbells (Definition 3.19) are those closed walks that involve the same edges when the direction is reversed. These exceptions are rare enough (Lemma 3.20) that they do not affect the asymptotic. Finally, isogeny cycles are those non-backtracking closed walks which are not powers of smaller non-backtracking closed walks; this requirement clearly also does not affect the asymptotic.   □

Remark 7.2. In Theorem 7.1, we restrict to $p \equiv 1 \pmod{12}$ because the graph theory involved requires an $(\ell+1)$-regular graph. We have remarked at length the difficulties which arise when vertices $j = 0, 1728$ appear in $\mathcal{G}_\ell$. However, these issues are very localized – the $(\ell+1)$-regularity of the graph may only be affected at vertices adjacent to $j = 0, 1728$. For $p \gg \ell$, this is a very small number of potential deviations from regularity, and we do not expect the estimates provided here to be greatly affected.

By reference to [31] and [2], the mixing rate of the Markov chain could be used to analyse the rate of convergence of the asymptotic and make this theorem more precise.

The proof of Theorem 7.1 lends itself to computation: using powers of the non-backtracking matrix, we can compute the number of isogeny cycles in $\mathcal{G}_\ell$, at least up to the issue of barbells. Figure 7 shows some data collected in SageMath for the number of non-backtracking closed walks of length r on some small supersingular ℓ-isogeny graphs chosen to avoid loops (so $\mathcal{G}_\ell$ is $(\ell+1)$-regular without any barbells), and its agreement with the asymptotic.

7.2. **Counting isogeny cycles with class numbers.** A corollary to Theorems 3.2 and 5.1 counts cycles exactly as a sum of class numbers of certain quadratic orders.

Corollary 7.3. *Let $r > 2$ be fixed and consider primes $\ell \neq p$. Let $\mathcal{I}_r$ be as in Theorem 5.1. Let $c_r = |\mathcal{C}_r|$ denote the number of directed isogeny cycles of length r. Then*

$$c_r = \frac{1}{r} \sum_{\mathcal{O} \in \mathcal{I}_r} \epsilon_{\mathcal{O},\ell} h_{\mathcal{O}},$$

where $\epsilon_{\mathcal{O},\ell}$ is as defined in Theorem 5.1. In particular, $1 \leq \epsilon_{\mathcal{O},\ell} \leq 2$ and $\epsilon_{\mathcal{O},\ell} = 2$ if p is inert in the quadratic field containing $\mathcal{O}$. If $\ell^r < p$, then $\epsilon_{\mathcal{O},\ell} = 2$ for all $\mathcal{O} \in \mathcal{I}_r$.

Proof. This follows from Theorems 3.2 and 5.1. The cycles $\mathcal{C}_r$ of length r in $\mathcal{G}_\ell$ are in bijection with $\mathcal{R}_r/\sim$. Each $[R]$ then maps to some order $\mathcal{O} \in \mathcal{I}_r$. Summing the elements in $\mathcal{I}_r$ weighted by the fibre sizes counts $\#\mathcal{C}_r$. The fibre sizes are given in Theorem 5.1. $\qquad\square$

On the graph side, it is clear that this count is finite. On the side of ideal classes, finiteness of $\mathcal{I}_r$ follows from the observation that if $\mathrm{ord}([\mathfrak{l}]) = r$, then ℓ^r is represented by the norm form of the quadratic field. For sufficiently large discriminant, this representation is impossible except as $\ell^r = (\ell^{r/2})^2$ if r is even; but in this case, ℓ is ramified. This observation can be turned into a counting argument based on quadratic residues.

Theorem 7.4. *Assume $p > 2$. For any positive integer N, define*

$$\mathcal{Q}(N) = \left\{ 0 < x < 2\ell^{N/2} : \begin{array}{c} x \not\equiv 0 \pmod{\ell}, \\ x^2 - 4\ell^N \text{ is not a quadratic residue modulo } p, \\ \text{and has valuation } \leq 1 \text{ at } p \end{array} \right\},$$

and

$$Q_N = \sum_{x \in \mathcal{Q}(N)} \sum_{f^2 \mid (x^2 - 4\ell^N)} \epsilon_{\mathcal{O},\ell}\, h\left(\frac{x^2 - 4\ell^N}{f^2}\right),$$

where $h(D)$ denotes the class number of the order of discriminant D. Let c_N, $N \geq 3$, be as in Corollary 7.3. Then there are values for c_1 and c_2 such that

$$\sum_{r \mid N} r c_r = Q_N, \quad c_N = \frac{1}{N} \sum_{r \mid N} \mu(r) Q_{N/r}.$$

Proof. Let $r \geq 3$. Recall from Theorems 3.2 and 5.1 that we have maps

$$\mathcal{C}_r \longrightarrow \mathcal{R}_r/\sim \longrightarrow \mathcal{I}_r.$$

The first map is a bijection, and the second map has fibre sizes $\epsilon_{\mathcal{O},\ell} h_{\mathcal{O}}/r$ above $\mathcal{O}$, as described in Theorem 5.1. Using the fibre sizes, we obtain

$$c_N = \frac{1}{N} \sum_{\mathcal{O} \in \mathcal{I}_N} \epsilon_{\mathcal{O},\ell} h_{\mathcal{O}}.$$

Then

$$\sum_{r \mid N} r c_r = \sum_{\mathcal{O} \in \bigcup_{r \mid N} \mathcal{I}_r} \epsilon_{\mathcal{O},\ell} h_{\mathcal{O}}.$$

By definition, $\mathcal{O} \in \bigcup_{r \mid N} \mathcal{I}_r$ if and only if there exists an imaginary quadratic integer $\alpha \in \mathcal{O}$, not divisible by ℓ, such that $N(\alpha) = \ell^r$, p does not split in $\mathbb{Q}(\alpha)$, and ℓ is invertible and splits in $\mathcal{O}$. Since these properties of α are invariant up to conjugation and sign, we can write $\alpha = \frac{x + \sqrt{\Delta}}{2}$, where Δ is the discriminant of the order $\mathbb{Z}[\alpha]$, and $x \geq 0$. The condition that α has norm ℓ^N is equivalent to $x^2 - \Delta = 4\ell^N$. The fact that α is imaginary is equivalent to $0 > \Delta = x^2 - 4\ell^N$, or equivalently, $x < 2\ell^{N/2}$. Next we show that the prime ℓ is invertible and splits in $\mathbb{Z}[\alpha]$ if and only if $x \not\equiv 0 \pmod{\ell}$. To that end, invertibility is equivalent to the condition that ℓ does not divide the conductor of $\mathbb{Z}[\alpha]$, which is equivalent to $x \not\equiv 0 \pmod{\ell}$ (since $x^2 - \Delta = 4\ell^N$), and if these equivalent conditions hold, then $1 = \left(\frac{\Delta}{\ell}\right) = \left(\frac{x^2 - 4\ell^N}{\ell}\right)$, which holds if and only if ℓ splits in the maximal order containing $\mathcal{O}$ (by Lemma 2.28). The prime p does not split in this maximal order if and only if $x^2 - 4\ell^N$ is not a quadratic residue modulo p. Finally, α is not divisible by ℓ if and only if ℓ does not divide Δ. To conclude then, $\mathcal{O} \in \bigcup_{r \mid N} \mathcal{I}_r$ if and only if $\Delta_{\mathcal{O}} = x^2 - 4\ell^N$ for $x \in \mathcal{Q}(N)$ of the

theorem statement. Note that each $\mathcal{O}$ corresponds to one $\Delta_{\mathcal{O}}$, which in turn corresponds to one value of x. Therefore,

$$\sum_{r|N} rc_r = \sum_{\mathcal{O} \in \bigcup_{r|N} \mathcal{I}_r} \epsilon_{\mathcal{O},\ell} h_{\mathcal{O}} = \sum_{x \in \mathcal{Q}(N)} \sum_{f^2 | x^2 - 4\ell^N} \epsilon_{\mathcal{O},\ell} h\left(\frac{x^2 - 4\ell^N}{f^2}\right) = Q_N.$$

Calling upon Möbius inversion, we have proved the theorem.   □

Remark 7.5. The key observation we make here – that small cycles arise from a combination of splitting behaviour between p and ℓ – is used in [15, Section 5.3.4] to eliminate small cycles by use of congruence conditions on p.

The advantage of Theorem 7.4 is that in contrast to the sum of class numbers in Corollary 7.3, we obtain a sum of class numbers that has no conditions on the class groups. The condition on the order of $[\mathfrak{l}]$ present in Corollary 7.3 is removed when summing over $x \in \mathcal{Q}(N)$, since $\mathcal{Q}_N$ is entirely given in terms of 'elementary' conditions, namely Legendre symbols and valuations at p and ℓ.

Example 7.6. We illustrate the theorem and its proof with the example of Section 6, where $\ell = 2$ and $p = 179$. We have, from Table 1, the directed isogeny cycle counts

$$c_3 = 2, \ c_4 = 2, \ c_5 = 2, \ c_6 = 14.$$

Using Sage to compute $\mathcal{Q}(N)$, we obtain, showing the data as tuples $(x, x^2 - 4\ell^n, h_{x^2-4\ell^n})$,

$$\mathcal{Q}(1) = \emptyset, \mathcal{Q}(2) = \{(1, -15, 2)\}, \mathcal{Q}(3) = \{(1, -31, 3)\}, \mathcal{Q}(4) = \{(5, -39, 4), (7, -15, 2)\}, \mathcal{Q}(5) = \{(9, -47, 5)\},$$

$$\mathcal{Q}(6) = \{(1, -255, 12), (3, -247, 6), (5, -231, 12), (11, -135, 6), (13, -87, 6), (15, -31, 3)\}.$$

Observe that the last tuple $(15, -31, 3)$ in $\mathcal{Q}(6)$ corresponds to $\alpha = \frac{\pm 15 \pm \sqrt{-31}}{2} = \pm\left(\frac{1 \pm \sqrt{-31}}{2}\right)^2$. The element $\beta = \frac{\pm 1 \pm \sqrt{-31}}{2}$ corresponds to the unique 3-cycle in Table 1. Note that $\mathbb{Z}[\alpha] = \mathbb{Z}[\beta] = \mathcal{O}_{\mathbb{Q}(\sqrt{-31})}$. Accordingly, α corresponds to the 3-cycle repeated twice.

The fourth tuple $(11, -135, 6)$ in $\mathcal{Q}(6)$ corresponds to $\alpha = \frac{\pm 11 \pm 3\sqrt{-15}}{2} = \left(\frac{\pm 1 \pm \sqrt{-15}}{2}\right)^3$. The element $\frac{\pm 1 \pm \sqrt{-15}}{2}$ corresponds to the 2-cycle $(112, 35)$. In this case, $\mathbb{Z}[\beta] = \mathcal{O}_{\mathbb{Q}(\sqrt{-15})}$, but $\mathbb{Z}[\alpha]$ is a suborder of $\mathbb{Z}[\beta]$ of conductor 3. Thus, α corresponds to two different types of 6-cycles: the 2-cycle repeated 3 times (therefore not an isogeny cycle); and also an isogeny cycle of size 6, as in Table 1.

Therefore,

$$Q_1 = 0,$$

$$Q_2 = 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-15}}{2}\right]\right) = 4,$$

$$Q_3 = 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-31}}{2}\right]\right) = 6,$$

$$Q_4 = 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-15}}{2}\right]\right) + 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-39}}{2}\right]\right) = 4 + 8 = 12,$$

$$Q_5 = 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-47}}{2}\right]\right) = 10,$$

$$Q_6 = 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-255}}{2}\right]\right) + 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-247}}{2}\right]\right) + 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-231}}{2}\right]\right) +$$

$$2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-15}}{2}\right]\right) + 2h\left(\mathbb{Z}\left[3\frac{1 + \sqrt{-15}}{2}\right]\right) + 2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-87}}{2}\right]\right) +$$

$$2h\left(\mathbb{Z}\left[\frac{1 + \sqrt{-31}}{2}\right]\right) = 24 + 12 + 24 + 4 + 12 + 12 + 6 = 94.$$

Recall that for the Möbius function, $\mu(1) = 1, \mu(2) = -1, \mu(3) = -1, \mu(4) = 0, \mu(5) = -1, \mu(6) = 1$. So the sum $S(N) := \frac{1}{N} \sum_{r|N} \mu(r) Q_{N/r}$ takes on the following values, which we can verify to match c_N for

$3 \leq N \leq 6$:

$$S(1) = Q_1 = 0, \quad S(2) = \frac{1}{2}(Q_2 - Q_1) = 2, \quad S(3) = \frac{1}{3}(Q_3 - Q_1) = 2 = c_3,$$

$$S(4) = \frac{1}{4}(Q_4 - Q_2) = 2 = c_4, \quad S(5) = \frac{1}{5}(Q_5 - Q_1) = 2 = c_5,$$

$$S(6) = \frac{1}{6}(Q_6 - Q_3 - Q_2 + Q_1) = 14 = c_6.$$

We briefly confirm that the conclusion of Theorem 7.4 is in heuristic agreement with Theorem 7.1. Heuristically, we expect $\#\mathcal{Q}(N) \approx \ell^{N/2}$. Estimating the inner sum in Q_N with a Hurwitz class number heuristically of size $\sqrt{4\ell^N - x^2}$ (ignoring log factors), we have

$$Q_N \approx \sum_{x=1}^{\lfloor 2\ell^{N/2} \rfloor} \sqrt{4\ell^N - x^2} \approx \int_1^{2\ell^{N/2}} \sqrt{4\ell^N - x^2} \, dx \approx \ell^N,$$

from which follows $c_N \approx \ell^N/N$ as expected. In fact, a rigorous upper bound of magnitude $\ell^N \log N$ on c_N is given in Corollary 7.7, whose proof is based on this rough idea.

Corollary 7.7. *With the notation of Theorem 7.4,*

$$Q_N < B_N := \frac{2}{3}(e^\gamma \log\log(2\ell^{N/2}) + 7/3)\log(4\ell^N)(\pi\ell^N + 2\ell^{3N/4})$$

and

$$c_N < \frac{B_N}{N} + \left(e^\gamma \log\log N + 7/3 - \frac{1}{N}\right) B_{N/2}$$

$$< \frac{2\pi e^\gamma \log(4\ell)}{3} \ell^N (\log N + \log\log(2\sqrt{\ell}) + 7/3) + O(\ell^{3N/4} \log N) \text{ as } N \to \infty,$$

where $\gamma = 0.577\ldots$ is the Euler-Mascheroni constant.

Proof. For brevity, put $X = \lfloor 2\ell^{N/2} \rfloor$. For our bound, we simply ignore all the conditions on x imposed by membership in the set $\mathcal{Q}_N$ of Theorem 7.4 and use $\epsilon_{\mathcal{O},\ell} \leq 2$. That is, we upper-bound Q_N by $Q'_N \geq Q_N$, where

$$Q'_N = 2 \sum_{x=1}^{X} \sum_{f^2 | x^2 - 4\ell^N} h\left(\frac{x^2 - 4\ell^N}{f^2}\right).$$

Fix an integer x with $0 \leq x \leq X$ and write $4\ell^N - x^2 = y^2 d$ for positive integers d, y with d squarefree. For any imaginary quadratic order $\mathcal{O}$ of discriminant Δ, we have $h_{\mathcal{O}} < (\sqrt{|\Delta|} \log |\Delta|)/3$ by [39, Equation (8.11)], so

$$\sum_{f^2 | (x^2 - 4\ell^N)} h\left(\frac{x^2 - 4\ell^N}{f^2}\right) < \frac{1}{3} \sum_{f^2 | (x^2 - 4\ell^N)} \frac{\sqrt{4\ell^N - x^2}}{f} \log\left(\frac{4\ell^N - x^2}{f}\right)$$

$$< \frac{1}{3} \log(4\ell^N) \sum_{f^2 | y^2 d} \frac{\sqrt{y^2 d}}{f}$$

$$= \frac{1}{3} \log(4\ell^N) \sqrt{y^2 d} \sum_{f | y} \frac{1}{f}$$

$$= \frac{1}{3} \sqrt{4\ell^N - x^2} \log(4\ell^N) \frac{\sigma(y)}{y},$$

where $\sigma(\cdot)$ denotes the sum of divisors function. It is known (see [3, Exercise 3.9 (a)] for example) that $\sigma(n)\varphi(n) \leq n^2$ for all $n \geq 1$, where $\varphi(\cdot)$ is Euler's totient function. By [48, Theorem 15], we have

$$\frac{n}{\varphi(n)} < e^\gamma \log\log n + \frac{2.50637}{\log\log n}$$

for all $n \geq 3$. If $y \geq 19$, then

$$\frac{2.50637}{\log\log y} \leq \frac{2.50637}{\log\log(19)} < \frac{7}{3},$$

so $y = \sqrt{(4\ell^N - x^2)/d} < 2\ell^{N/2}$ yields

$$\frac{\sigma(y)}{y} \leq \frac{y}{\varphi(y)} < e^\gamma \log\log y + \frac{2.50637}{\log\log(y)} < e^\gamma \log\log(2\ell^{N/2}) + \frac{7}{3}.$$

Simple numerical verification shows that $\sigma(y)/y \leq 7/3$ for $1 \leq y \leq 18$ (the largest value is $\sigma(12)/12 = 7/3$), so this upper bound holds for all $y > 0$. It follows that

$$\sum_{f^2 \mid x^2 - 4\ell^N} h\left(\frac{x^2 - 4\ell^N}{f^2}\right) < \frac{1}{3}(e^\gamma \log\log(2\ell^{N/2}) + 7/3)\log(4\ell^N)\sqrt{4\ell^N - x^2}.$$

It remains to bound $\sum_{x=1}^X \sqrt{4\ell^N - x^2}$. Since $f(x) = 4\ell^N - x^2$ is monotonically decreasing, we have

$$\sum_{x=1}^X \sqrt{4\ell^N - x^2} \leq \int_0^X \sqrt{4\ell^N - x^2}\, dx = \frac{1}{2}X\sqrt{4\ell^N - X^2} + 2\ell^N \arcsin\left(\frac{X}{2\ell^{N/2}}\right).$$

Now $2\arcsin(X/2\ell^{N/2}) \leq 2\arcsin(1) = \pi$. Furthermore, $0 \leq 2\ell^{N/2} - X < 1$ and $0 < 2\ell^{N/2} + X < 4\ell^{N/2}$, so

$$\frac{1}{2}X\sqrt{4\ell^N - X^2} = \frac{1}{2}X\sqrt{(2\ell^{N/2} - X)(2\ell^{N/2} + X)} < \frac{1}{2}2\ell^{N/2}\sqrt{4\ell^{N/2}} = 2\ell^{3N/4}.$$

It follows that $Q'_N < B_N$ where

$$B_N = \frac{2}{3}\left(e^\gamma \log\log(2\ell^{N/2}) + 7/3\right)\log(4\ell^N)(\pi\ell^N + 2\ell^{3N/4}).$$

Next, we bound c_N. The quantity B_N is monotonically increasing as a function of N, so $Q_{N/r} < B_{N/r} \leq B_{N/2}$ for all divisors r of N with $r \geq 2$. Thus,

$$c_N = \frac{1}{N}\left(Q_N + \sum_{1 < r \mid N} \mu(r)Q_{N/r}\right) \leq \frac{1}{N}\left(Q_N + \sum_{1 < r \mid N} Q_{N/r}\right) < \frac{1}{N}\left(B_N + (\sigma(N) - 1)B_{N/2}\right).$$

Since $N \geq 3$, we can bound $\sigma(N)$ as before, i.e. $\sigma(N) < N(e^\gamma \log\log N + 7/3)$, to obtain

$$c_N < \frac{1}{N}\left(B_N + (N(e^\gamma \log\log N + 7/3) - 1)B_{N/2}\right) = \frac{B_N}{N} + \left(e^\gamma \log\log N + 7/3 - \frac{1}{N}\right)B_{N/2}.$$

The asymptotic bound on c_N is obtained by bounding the quantity $2\ell^{N/2}$ in the double log by $(2\sqrt{\ell})^N$. $\qquad\square$

Remark 7.8. The proof of Corollary 7.7 simply ignores the restriction that p should not split and ℓ should split in any of the orders under consideration. If we wish to prove a lower bound, these restrictions must be addressed. There is a widely believed heuristic that a Legendre symbol $\left(\frac{f(x)}{p}\right)$ should take values ± 1 with equal probability, independent of the size of x. However, known bounds on the smallest x for which a polynomial $f(x)$ is a non-residue modulo p are comparatively weak. There is a bound following from work of Weil; for example, Burgess mentioned this in 1967, observing there is a non-residue $x < H$ for $H \gg \sqrt{p}\log p$ [10]. Improvements have only been made in general for polynomials in at least two variables (in contrast to the special case of $f(x) = x$, where this problem is known as the least non-residue problem); for a discussion and literature review, see [45].

8. Path finding with oriented isogeny volcanoes

In this section, we review the previous work of the present authors [5] which provides inspiration for the subject of this paper.

As discussed in the introduction, finding a path between two curves in a supersingular ℓ-isogeny graph compromises the security of a variety of cryptographic systems. The bijection of Theorem 3.2 serves as inspiration for producing such a path, by highlighting the many volcano structures present in the graph, hinting that one might navigate the graph by reference to these oriented volcanoes. In fact, we saw that every cycle in the supersingular ℓ-isogeny graph is a rim of a volcano in some oriented isogeny graph. In previous work [5], the authors presented explicit classical and quantum algorithms for finding a path from any elliptic curve E to an initial curve E_{init}. Here, E_{init} is generally taken to be the supersingular elliptic curve of j-invariant 1728, but other choices for E_{init} are also possible. The algorithms use the knowledge of one endomorphism θ on E, or equivalently, an explicit orientation of E by some imaginary quadratic field K,

to navigate the oriented ℓ-isogeny graph $\mathcal{G}_{K,\ell}$. Their complexity is subject to certain plausible heuristics, further detailed below, and is in part governed by the time it takes to evaluate θ on points on E. The reader is encouraged to consult [56] for a different approach to solving the endomorphism ring problem using the data of an orientation.

The runtime of the classical algorithm additionally depends subexponentially on the degree d of θ and linearly on a certain class number that can be significantly smaller than the class number of the quadratic order of discriminant Δ defined by θ. This runtime can be substantially improved in many cases, leading in particular to certain new families of endomorphisms on *every* supersingular elliptic curve, whose exposure gives rise to a classical polynomial-time algorithm for finding an ℓ-isogeny path from E to E_{init}. While this fact may on first glance seem alarming, the proof is non-constructive and thus does not provide an efficient strategy for finding such an offending endomorphism. So this result does not pose a threat to the security of isogeny based cryptosystems.

The quantum path finding algorithm relies on two different quantum subroutines: solving an oriented variant of the VECTORIZATION problem, which asks to find a class group element acting to take one oriented curve to another (see also [16, 56]), and a new problem called PRIMITIVEORIENTATION, which asks, for a given endomorphism, to determine the quadratic order for which the orientation it induces is primitive. The resulting algorithm finds a smooth ℓ-isogeny from E to E_{init} such that the smoothness bound and the runtime are subexponential in $\log |\Delta|$.

For completeness, we reproduce the key results of [5] here and give a brief road map of the path finding algorithms. Let $E/\overline{\mathbb{F}}_p$ be a supersingular elliptic curve and θ an endomorphism on E of degree d and discriminant Δ not divisible by p. Denote by $T_\theta(k, p)$ the time it takes to evaluate θ on points on E defined over $\mathbb{F}_{p^k}$. If θ is given by rational maps, then $T_\theta(k, p)$ is a polynomial in d, k and $\log p$. But if θ is given as the product of endomorphisms of smaller degree, then the dependence of $T_\theta(k, p)$ on d is much more favourable. All the algorithms of [5] assume in practice one of these two representations, which are typical in applications involving elliptic curve endomorphisms.

For any discriminant Δ of some quadratic order, the ℓ-*fundamental part* of Δ is the unique discriminant Δ' such that $\Delta = \ell^{2m}\Delta'$, with $v_\ell(\Delta') \leq 1$ when ℓ or Δ are odd and $v_\ell(\Delta'/4) \leq 1$ when $\ell = 2$ and Δ is even. For brevity, write $L_x(1/2) = \exp(O(\sqrt{\log x \log \log x}))$.

Theorem 8.1 ([5, Theorem 11.2]). *Let $E/\overline{\mathbb{F}}_p$ be a supersingular elliptic curve and θ an endomorphism on E of degree d and discriminant Δ coprime to p. Let $\ell \neq p$ be a prime (assumed to be constant in the runtime estimate), Δ' the ℓ-fundamental part of Δ, and $h_{\Delta'}$ the class number of the quadratic order of discriminant Δ'. If $L_d(1/2)$ is bounded below by a polynomial in $\log p$ and $|\Delta'| \leq p^2$, then there is a classical algorithm that finds an ℓ-isogeny path of length $O(\log p + h_{\Delta'})$ from E to the curve E_{init} of j-invariant $j = 1728$ in heuristic runtime $T_\theta(L_d(1/2), p) + h_{\Delta'}L_d(1/2)\operatorname{poly}(\log p)$.*

Theorem 8.2 ([5, Theorem 11.3]). *Let $E/\overline{\mathbb{F}}_p$ be a supersingular elliptic curve and θ an endomorphism on E of discriminant Δ coprime to p such that $T_\theta(k, p)$ is bounded below by a polynomial in $\log(p^k)$. If $|\Delta| \leq p^2$, then there is a quantum algorithm that finds an $L_{|\Delta|}(1/2)$-smooth isogeny of norm $O(\sqrt{|\Delta|})$ from E to E_{init} in heuristic runtime $T_\theta(O(\log^2 d), p)L_{|\Delta|}(1/2)$.*

The basic strategy of the algorithm of Theorem 8.1 ([5, Algorithm 8.1]) is as follows. Use the orientation induced by θ as an "orienteering tool" to climb from E to the rim of the oriented ℓ-isogeny volcano containing E. Generate this entire rim via the action of the class group of the associated ℓ-fundamental order $\mathcal{O}$ on the rim curves and store all their j-invariants in a list L. Orient the curve E_{init} (whose endomorphism ring is known) by the same orientation as E and use it to climb its volcano. Hoping to reach the same rim as the path originating at E, i.e. a vertex in L, construct the desired path from E to E_{init} from the two ascending paths and one of the two rim segments connecting them.

Throughout the computation, it is imperative to keep the size of the endomorphisms and the discriminants under consideration reasonable. To that end, endomorphisms are maintained in factored form as a product of endomorphisms of bounded prime power degree, and a technique akin to sieving produces such a representation for which the trace of the endomorphism is also of manageable size. Navigating the oriented ℓ-isogeny graph requires effective subroutines for detecting whether an edge is ascending/descending/horizontal, traversing an edge and carrying along the orientation, dividing an endomorphism by ℓ (to traverse an ascending edge), evaluating an endomorphism on ℓ-torsion points (required for the class group action), and factoring an endomorphism into a product of endomorphims of prime power degree while

controlling the size of its reduced trace. Algorithms for performing all these computational tasks, accompanied by proofs of correctness and a complete complexity analysis, can be found in [5]. Code is available at [6] and was used for a proof-of-concept example to generate an explicit 2-isogeny path of length 5 from $E : y^2 = x^3 + (7i + 86)x + (45i + 174)$ to $E_{\text{init}} : y^2 = x^3 - x$ over $\mathbb{F}_{179^2}$, where $i^2 = -1$.

A crucial ingredient in the path finding technique is orienting E_{init} appropriately. Given an ℓ-fundamental order $\mathcal{O} \subset K$ of discriminant Δ, this requires finding an endomorphism $\theta \in \text{End}(E_{\text{init}})$ that generates an order $\mathcal{O}' \subseteq \mathcal{O}$. The relative index $[\mathcal{O} : \mathcal{O}']$ is a power of ℓ, say ℓ^r for some $r \geq 0$. For reasons of efficiency, r should be small, situating $\mathcal{O}'$ as close to a rim as possible. At the same time, to guarantee the existence of a path from the starting curve E to E_{init} in $\mathcal{G}_{K,\ell}$, the orientation induced by θ must place E_{init} in the same connected component (i.e. volcano) of $\mathcal{G}_{K,\ell}$ as the oriented starting curve. A natural strategy is thus to loop over $r = 0, 1, 2, \ldots$ until the desired endomorphism θ and order $\mathcal{O}'$ are found. This raises the key question of how large r needs to be to ensure that a suitable orientation is encountered, at least with high probability. The answer is governed by how vertices corresponding to a fixed curve are spread around the cordillera. We will return to this issue in a moment.

A number of the algorithms outlined above depend on heuristics. Quoting from the introductory text of [5]:

> We rely on a number of heuristic assumptions: (i) The Generalized Riemann Hypothesis. (ii) Powersmoothness in a quadratic sequence or form is as for random integers (a powersmooth analogue of the heuristic assumption underlying the quadratic sieve; see [5, Heuristics 5.10 and 9.3]). (iii) The orientations of a fixed j-invariant are distributed reasonably across all suitable volcanoes ([5, Heuristic 3.7]). (iv) This distribution is independent of a certain integer factorization ([5, Heuristic 6.7]). (v) The aforementioned integer factorization is prime with the same probability as a random integer ([5, Heuristic 6.4]; this heuristic is similar to those used in [23] and [32]).

Among these, item (i) is standard, and items (ii) and (v) are purely number theoretical questions. This leaves items (iii) and (iv) as the only heuristics which concern oriented isogeny graphs. Unfortunately, item (iv) seems out of reach at the moment, so we will focus on item (iii) which indeed addresses our earlier question. It asks us to consider, for a fixed j-invariant, how its various orientations appear in a cordillera of volcanoes: does it prefer certain volcanoes over others? At the same time, it is natural to consider the 'opposite' question: for a fixed volcano, how do the various j-invariants lie upon it? Do some j-invariants sit at greater depth than others? This can be rephrased as a question about the covering map from an oriented volcano to the usual supersingular ℓ-isogeny graph $\mathcal{G}_\ell$. In Section 9, we address some of these questions.

9. Random walks and consequences for oriented isogeny volcanoes

In this section, we use a standard random walk result for Ramanujan graphs to address some questions about the appearance of j-invariants on oriented ℓ-isogeny volcanoes.

9.1. **Random walks in an ℓ-isogeny graph.** In this entire section, as a matter of convenience, **we assume that $p \equiv 1 \pmod{12}$ to avoid issues of extra automorphisms**. In particular, strictly speaking, Proposition 9.1 depends upon results about regular undirected Ramanujan graphs. Of course, the result can be expected to hold for all $\mathcal{G}_\ell$, regardless of p, which are at worst very nearly undirected regular graphs, having some perturbations around the two vertices $j = 0$ and $j = 1728$ if these belong to $\mathcal{G}_\ell$.

We refer to the length (number of ℓ-isogenies) of a shortest path between two curves in the ℓ-isogeny graph $\mathcal{G}_\ell$ as the *distance* between the two curves . The diameter of $\mathcal{G}_\ell$ (the maximum distance between pairs of points) is known to be bounded by $2 \log p$ [46, Theorem 1]; see also [38]. One way to obtain such results is using the Ramanujan property and standard methods for expander graphs [29]; our version and proof here closely follow [27, Theorem 1].

Proposition 9.1. *Let $p \equiv 1 \pmod{12}$. Let ℓ be a prime and $\mathcal{G}_\ell$ be the supersingular ℓ-isogeny graph over $\overline{\mathbb{F}}_p$, with $\#\mathcal{G}_\ell$ vertices. Let $\mathcal{S}$ represent a subset of $\#\mathcal{S}$ initial vertices. Let j_1 represent a target vertex. Let j represent the vertex reached by a random walk of t steps from a random vertex in $\mathcal{S}$. We denote by Pr_t the probability of an event over all such random walks of length t. Then we have*

$$\left| \text{Pr}_t(j = j_1) - \frac{1}{\#\mathcal{G}_\ell} \right| \leq \frac{1}{\#\mathcal{S}} \left(\frac{2\sqrt{\ell}}{\ell + 1} \right)^t.$$

Proof. Choose a basis of a vector space of dimension $\#\mathcal{G}_\ell$ to correspond to the vertices of $\mathcal{G}_\ell$. Let A_ℓ represent the $\#\mathcal{G}_\ell \times \#\mathcal{G}_\ell$ adjacency matrix of $\mathcal{G}_\ell$. Let $\mathbf{v}_0$ represent an initial vector whose entry for vertex v is $1/\#\mathcal{S}$ if $v \in \mathcal{S}$ and 0 otherwise. This is the probability vector for the initial vertex of the random walk. Let λ_1 be the largest eigenvalue of A_ℓ, and λ_2 the second largest. Recursively define $\mathbf{v}_t = \frac{1}{\lambda_1} A_\ell \mathbf{v}_{t-1}$, so $\mathbf{v}_t$ is the probability vector for the t-th vertex of the random walk.

Let $\mathbf{u}$ have all entries $1/\#\mathcal{G}_\ell$, so $\mathbf{u}$ is the probability vector for the uniform distribution over the vertices of $\mathcal{G}_\ell$. Then $\mathbf{u}$ is an eigenvector for λ_1, which has an eigenspace of dimension 1. Choose an orthogonal basis of eigenvectors $\mathbf{e}_i$ for $\mathbb{R}^{\#\mathcal{G}_\ell}$ such that $\mathbf{e}_1 = \mathbf{u}$. Then since $\mathbf{v}_{t-1}$ is a probability distribution, $\langle \mathbf{v}_{t-1}, \mathbf{1} \rangle = \langle \mathbf{u}, \mathbf{1} \rangle = 1$. Writing $\mathbf{v}_{t-1} = \sum_i \alpha_i \mathbf{e}_i$, this implies that $\alpha_1 = 1$. Hence $\mathbf{v}_{t-1} - \mathbf{u} = \sum_{i=2}^{\#\mathcal{G}_\ell} \alpha_i \mathbf{e}_i$. Therefore,

$$A_\ell(\mathbf{v}_{t-1} - \mathbf{u}) = \sum_{i=2}^{\#\mathcal{G}_\ell} \lambda_i \alpha_i \mathbf{e}_i.$$

Consequently, we have

$$\|\mathbf{v}_t - \mathbf{u}\|_2 = \left\| \frac{1}{\lambda_1} A_\ell(\mathbf{v}_{t-1} - \mathbf{u}) \right\|_2 \leq \frac{1}{\lambda_1} \lambda_2 \|\mathbf{v}_{t-1} - \mathbf{u}\|_2.$$

We also have

$$\|\mathbf{v}_0 - \mathbf{u}\|_2 = \#\mathcal{S} \left(\frac{1}{\#\mathcal{S}} - \frac{1}{\#\mathcal{G}_\ell} \right)^2 + (\#\mathcal{G}_\ell - \#\mathcal{S}) \left(\frac{1}{\#\mathcal{G}_\ell} \right)^2 = \frac{1}{\#\mathcal{S}} \left(1 - \frac{\#\mathcal{S}}{\#\mathcal{G}_\ell} \right) \leq \frac{1}{\#\mathcal{S}}.$$

Iterating, we have

$$\left| \Pr(j = j_1) - \frac{1}{\#\mathcal{G}_\ell} \right| = \|\mathbf{v}_t - \mathbf{u}\|_\infty \leq \|\mathbf{v}_t - \mathbf{u}\|_2 \leq \frac{\lambda_2^t}{\lambda_1^t} \|\mathbf{v}_0 - \mathbf{u}\|_2 \leq \frac{\lambda_2^t}{\#\mathcal{S}\lambda_1^t} \leq \frac{1}{\#\mathcal{S}} \left(\frac{2\sqrt{\ell}}{\ell + 1} \right)^t.$$

The final inequality is the bound on the spectral gap (5) for these graphs. $\qquad\square$

Corollary 9.2. *Let j_1 be a j-invariant and $\mathcal{S}$ a subset of the vertices of $\mathcal{G}_\ell$. If*

$$t > \frac{\log \#\mathcal{G}_\ell - \log \#\mathcal{S}}{\log((\ell+1)/2\sqrt{\ell})}$$

then there exists a path of length $\leq t$ in $\mathcal{G}_\ell$ starting in $\mathcal{S}$ and ending at j_1. In particular, for

$$t > \log_{\sqrt{\ell}/2} \#\mathcal{G}_\ell - \log_{\sqrt{\ell}/2} \#\mathcal{S},$$

all curves are within distance t of $\mathcal{S}$.

Proof. This follows from Proposition 9.1 by multiplying through by $\#\mathcal{G}_\ell$, to obtain

$$|\#\mathcal{G}_\ell \Pr_t(j = j_1) - 1| \leq \frac{\#\mathcal{G}_\ell}{\#\mathcal{S}} \left(\frac{2\sqrt{\ell}}{\ell + 1} \right)^t$$

and then letting t be large enough so the right side is less than 1. $\qquad\square$

9.2. The depth of j-invariants on an oriented volcano.

Each volcano of the oriented ℓ-isogeny graph includes every j-invariant of the supersingular graph (to see this, fix a j-invariant j_0 of the volcano, and then consider an ℓ^n-isogeny to the desired j-invariant j_1; this generates a path in the oriented graph). It is natural to ask how far down a volcano one must go to see all j-invariants. We may use Corollary 9.2 to answer this.

Corollary 9.3. *Let $\mathcal{V}$ be an oriented volcano, and set $\mathcal{S} := \{j \in \mathcal{G}_\ell : (E_j, \iota) \text{ is on the rim of } \mathcal{V} \text{ for some } \iota\}$. Then all j-invariants appear on $\mathcal{V}$ at depth less than or equal to*

$$\lceil \log_{\sqrt{\ell}/2} \#\mathcal{G}_\ell - \log_{\sqrt{\ell}/2} \#\mathcal{S} \rceil.$$

Remark 9.4. The proof of [5, Proposition 6.5], particularly equation (3), provides a sort of computational proof of this behaviour in the case of the Gaussian field.

It is an interesting question to ask about the cardinality of $\mathcal{S}$, which is related to the question of the multiplicity of j-invariants in $\mathrm{SS}_{\mathcal{O}}^{\mathrm{pr}}$. That is, how many $\mathcal{O}$-primitive orientations are there for a fixed curve E? Kaneko showed that a curve cannot have more than one embedding when the discriminant is at most p in absolute value [30, Theorem 2']. On the other hand, [25, Proof of Theorem 1.4] shows that for any fixed curve E and integer t, orders $\mathcal{O}$ of sufficiently large discriminant will always embed into $\mathrm{End}(E)$ at least t times. The number of curves with an $\mathcal{O}$-orientation is bounded below in [34].

9.3. **The volcanoes upon which the orientations of a fixed j-invariant lie.** Every j-invariant occurs in every volcano (as described at the beginning of Section 9.2). In [5], the algorithms depended upon the following heuristic about the relative frequency of a fixed j-invariant across different volcanoes (quoted as item ((iii)) in Section 8 above).

Heuristic 9.5 ([5, Heuristic 3.7])**.** *Let $\mathcal{O}$ be an ℓ-fundamental order in a quadratic field K. Consider the finite union $\mathcal{S}$ of $\mathcal{O}'$-cordilleras in the oriented supersingular ℓ-isogeny graph $\mathcal{G}_{K,\ell}$ for all $\mathcal{O}' \supseteq \mathcal{O}$. Let $d(v)$ denote the depth of a vertex v on its volcano. Let $j(v)$ denote its j-invariant. Define:*

- *$R_{\mathcal{V}}$, the number of edges descending from the rim of the volcano $\mathcal{V} \in \mathcal{S}$;*
- *$R_{\mathcal{S}}$, the number of edges descending from all rims in $\mathcal{S}$.*

Then for any j-invariant j_0 and any volcano $\mathcal{V} \in \mathcal{S}$, the ratio

$$\frac{\#\{v \in \mathcal{V} : j(v) = j_0, d(v) \le t\}}{\#\{v \in \mathcal{S} : j(v) = j_0, d(v) \le t\}}$$

approaches $R_{\mathcal{V}}/R_{\mathcal{S}}$ as $t \to \infty$.

However, naïve methods can only prove a weighted version of this result, where j-invariants at smaller depth are weighted more heavily. Our result also has a restriction on the rim sizes.

Theorem 9.6. *Consider a cordillera $\mathcal{C}$ in the oriented supersingular ℓ-isogeny graph $\mathcal{G}_{K,\ell}$ whose volcanoes are all isomorphic as graphs and $(\ell + 1)$-regular (e.g. a single $\mathcal{O}$-cordillera). Let $\mathcal{V}$ be one volcano of the cordillera $\mathcal{C}$ and $\mathcal{S}_{\mathcal{V}}$ the vertex set of its rim. Let $\mathcal{S}_{\mathcal{C}}$ be the set of vertices in the rims of the full cordillera. Let $d(v)$ denote the depth of the vertex v, and let $f_t(d)$ denote the function that counts the number of paths of length t starting at the rim of any volcano and terminating at any fixed vertex at depth d of that volcano (by symmetry, this is independent of the vertex but depends on the depth and rim size). If we define*

$$(11) \qquad S_{\mathcal{V},j_1,t} = \frac{\displaystyle\sum_{v \in \mathcal{V},\, j(v)=j_1} f_t(d(v))}{\displaystyle\sum_{v \in \mathcal{C},\, j(v)=j_1} f_t(d(v))} \,,$$

then

$$\left| S_{\mathcal{V},j_1,t} - \frac{\#\mathcal{S}_{\mathcal{V}}}{\#\mathcal{S}_{\mathcal{C}}} \right| \ll \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^t ,$$

where the implied constant depends on $\mathcal{C}$ and ℓ, but not t.

Proof. The numerator of (11) is the number of times we encounter $j = j_1$ among all paths of length t originating at the rim of $\mathcal{V}$. The denominator is the corresponding number for all paths of length t originating at the rims of $\mathcal{C}$. The number of paths of length t from the rim set of $\mathcal{V}$ is $(\ell + 1)\ell^{t-1}\#\mathcal{S}_{\mathcal{V}}$. So, applying Proposition 9.1 with $\mathcal{S} = \mathcal{S}_{\mathcal{V}}$, we have

$$\frac{1}{(\ell+1)\ell^{t-1}} \sum_{v \in \mathcal{V}, j(v)=j_0} f_t(d(v)) = Pr(j = j_0)\#\mathcal{S}_{\mathcal{V}} = \frac{\#\mathcal{S}_{\mathcal{V}}}{\#\mathcal{G}_\ell} + O\left(\left(\frac{2\sqrt{\ell}}{\ell+1} \right)^t \right).$$

Combining with the analogous fact for the denominator gives the result. $\square$

The restriction to $(\ell + 1)$-regularity rules out the rims having oriented curves with automorphisms (see Remark 2.7); it is merely a convenience and the result could be adjusted to accommodate these cases.

This result is insufficient for the purposes of [5] in that it doesn't rule out the possibility that some volcanoes have few occurrences of $j = j_0$ at small depth, while others have more such occurrences at greater depth. If this is indeed the case, it may be that Algorithm 6.1 of [5], which proceeds by depth, preferentially

finds certain volcanoes early on. (Nevertheless, Corollary 9.3 above provides some reassurance that all j-invariants begin to appear after a while.)

References

[1] G. Adj, O. Ahmadi, and A. Menezes. On isogeny graphs of supersingular elliptic curves over finite fields. *Finite Fields Appl.*, 55:268–283, 2019.

[2] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin. Non-backtracking random walks mix faster. *Commun. Contemp. Math.*, 09(04):585–603, 2007.

[3] T. M. Apostol. *Introduction to analytic number theory*. Undergraduate Texts in Mathematics. Springer-Verlag, New York-Heidelberg, 1976.

[4] S. Arpin, C. Camacho-Navarro, K. Lauter, J. Lim, K. Nelson, T. Scholl, and J. Sotáková. Adventures in Supersingularland. *Exp. Math.*, 0(0):1–28, 2021.

[5] S. Arpin, M. Chen, K. E. Lauter, R. Scheidler, K. E. Stange, and H. T. N. Tran. Orienteering with one endomorphism, 2022. https://arxiv.org/abs/2201.11079.

[6] S. Arpin, M. Chen, K. E. Lauter, R. Scheidler, K. E. Stange, and H. T. N. Tran. Win5 github repository, 2022. https://github.com/SarahArpin/WIN5.

[7] E. Bank, C. Camacho-Navarro, K. Eisenträger, T. Morrison, and J. Park. Cycles in the supersingular ℓ-isogeny graph and corresponding endomorphisms. In *Research Directions in Number Theory*, pages 41–66. Springer, 2019.

[8] BharatRam (https://mathoverflow.net/users/15311/bharatram). What does this connection between Chebyshev, Ramanujan, Ihara and Riemann mean? MathOverflow. https://mathoverflow.net/q/271786 (version: 2018-11-26).

[9] B. Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European J. Combin.*, 1(4):311–316, 1980.

[10] D. A. Burgess. On the quadratic character of a polynomial. *J. Lond. Math. Soc.*, s1-42(1):73–80, 1967.

[11] W. Castryck and T. Decru. An efficient key recovery attack on SIDH (preliminary version). 2022.

[12] W. Castryck, M. Houben, F. Vercauteren, and B. Wesolowski. On the decisional Diffie-Hellman problem for class group actions on oriented elliptic curves. Cryptology ePrint Archive, Paper 2022/345, 2022. https://ia.cr/2022/345.

[13] W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes. CSIDH: an efficient post-quantum commutative group action. In *Advances in Cryptology—ASIACRYPT 2018. Part III*, volume 11274 of *Lecture Notes in Comput. Sci.*, pages 395–427. Springer, Cham, 2018.

[14] W. Castryck, L. Panny, and F. Vercauteren. Rational isogenies from irrational endomorphisms. In *Advances in Cryptology—EUROCRYPT 2020. Part II*, volume 12106 of *Lecture Notes in Comput. Sci.*, pages 523–548. Springer, Cham, 2020.

[15] D. X. Charles, E. Z. Goren, and K. E. Lauter. Cryptographic hash functions from expander graphs. *J. Cryptology*, 22(1):93–113, 2009. https://eprint.iacr.org/2006/021.

[16] M. Chenu and B. Smith. Higher-degree supersingular group actions. *Math. Cryptology*, 1(1):1–15, 2021.

[17] L. Colò and D. Kohel. Orienting supersingular isogeny graphs. *J. Math. Cryptol.*, 14(1):414–437, 2020.

[18] D. A. Cox. *Primes of the form $x^2 + ny^2$*. Pure and Applied Mathematics (Hoboken). John Wiley & Sons, Inc., Hoboken, NJ, second edition, 2013.

[19] G. P. Davidoff, P. Sarnak, and A. Valette. *Elementary Number Theory, Group Theory, and Ramanujan Graphs*. Cambridge University Press, 2003.

[20] L. De Feo, C. Delpech de Saint Guilhem, T. B. Fouotsa, P. Kutas, A. Leroux, C. Petit, J. Silva, and B. Wesolowski. *SÉTA: Supersingular Encryption from Torsion Attacks*, volume 13090 of *Lecture Notes in Comput. Sci.*, pages 249–278. Springer International Publishing, Cham, 2021.

[21] L. de Feo, D. Jao, and J. Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *J. Math. Cryptol.*, 8(3):209–247, 2014.

[22] L. De Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski. Sqisign: Compact post-quantum signatures from quaternions and isogenies. In *Advances in Cryptology – ASIACRYPT 2020: 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7–11, 2020, Proceedings, Part I*, page 64–93, Berlin, Heidelberg, 2020. Springer-Verlag.

[23] V. de Quehen, P. Kutas, C. Leonardi, C. Martindale, L. Panny, C. Petit, and K. E. Stange. *Improved Torsion-Point Attacks on SIDH Variants*, volume 12827 of *Lecture Notes in Comput. Sci.*, pages 432–470. Springer International Publishing, Cham, 2021.

[24] K. Eisenträger, S. Hallgren, K. Lauter, T. Morrison, and C. Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In *Advances in Cryptology—EUROCRYPT 2018. Part III*, volume 10822 of *Lecture Notes in Comput. Sci.*, pages 329–368. Springer, Cham, 2018.

[25] N. Elkies, K. Ono, and T. Yang. Reduction of CM elliptic curves and modular function congruences. *Int. Math. Res. Not.*, 2005(44):2695–2707, 2005.

[26] M. Fouquet and F. Morain. Isogeny volcanoes and the SEA algorithm. In *Algorithmic number theory (Sydney, 2002)*, volume 2369 of *Lecture Notes in Comput. Sci.*, pages 276–291. Springer, Berlin, 2002.

[27] S. D. Galbraith, C. Petit, and J. Silva. Identification protocols and signature schemes based on supersingular isogeny problems. *J. Cryptology*, 33(1):130–175, 2020.

[28] B. H. Gross. Heights and the special values of L-series. In *Number theory (Montreal, Que., 1985)*, volume 7 of *CMS Conf. Proc.*, pages 115–187. Amer. Math. Soc., Providence, RI, 1987.

[29] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561, 2006.

[30] M. Kaneko. Supersingular j-invariants as singular moduli mod p. *Osaka J. Math.*, 26(4):849–855, 1989.

[31] M. Kempton. Non-backtracking random walks and a weighted Ihara's theorem. *Open J. Discrete Math.*, 06:207–226, 2016.

[32] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol. On the quaternion ℓ-isogeny path problem. *LMS J.Comput. Math.*, 2014.

[33] D. R. Kohel. *Endomorphism rings of elliptic curves over finite fields*. ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.)–University of California, Berkeley.

[34] A. Leroux. An effective lower bound on the number of orientable supersingular elliptic curves. Cryptology ePrint Archive, Paper 2022/357, 2022. `https://ia.cr/2022/357`.

[35] S. Li, Y. Ouyang, and Z. Xu. Endomorphism rings of supersingular elliptic curves over $\mathbb{F}_p$. *Finite Fields Appl.*, 62:101619, 24, 2020.

[36] S. Li, Y. Ouyang, and Z. Xu. Neighborhood of the supersingular elliptic curve isogeny graph at $j = 0$ and 1728. *Finite Fields Appl.*, 61:101600, 16, 2020.

[37] J. Love and D. Boneh. Supersingular curves with small noninteger endomorphisms. In *ANTS XIV—Proceedings of the Fourteenth Algorithmic Number Theory Symposium*, volume 4 of *Open Book Ser.*, pages 7–22. Math. Sci. Publ., Berkeley, CA, 2020.

[38] E. Lubetzky and Y. Peres. Cutoff on all Ramanujan graphs. *Geom. Funct. Anal.*, 26(4):1190–1216, 2016.

[39] J. M. J. Jacobson and H. C. Williams. *Solving the Pell equation*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC. Springer, New York, 2009.

[40] L. Maino and C. Martindale. An attack on sidh with arbitrary starting curve. 2022.

[41] K. McMurdy. Explicit representation of the endomorphism rings of supersingular elliptic curves. `https://phobos.ramapo.edu/~kmcmurdy/research/McMurdy-ssEndoRings.pdf`, 2014.

[42] H. Onuki. On oriented supersingular elliptic curves. *Finite Fields App.*, 69, 2021.

[43] H. Onuki, Y. Aikawa, and T. Takagi. The existence of cycles in the supersingular isogeny graphs used in SIKE. In *2020 International Symposium on Information Theory and Its Applications (ISITA)*, pages 358–362, 2020.

[44] Y. Ouyang and Z. Xu. Loops of isogeny graphs of supersingular elliptic curves at $j = 0$. *Finite Fields Appl.*, 58:174–176, 2019.

[45] L. B. Pierce and J. Xu. Burgess bounds for short character sums evaluated at forms. *Algebra Number Theory*, 14(7):1911–1951, 2020.

[46] A. K. Pizer. Ramanujan graphs and Hecke operators. *Bull. Amer. Math. Soc. (N.S.)*, 23(1):127–137, 1990.

[47] D. Robert. Breaking sidh in polynomial time. Cryptology ePrint Archive, Paper 2022/1038, 2022. `https://eprint.iacr.org/2022/1038`.

[48] J. B. Rosser and L. Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois J. Math.*, 6:64–94, 1962.

[49] J. H. Silverman. *Advanced topics in the arithmetic of elliptic curves*, volume 151 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.

[50] J. H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.

[51] K. E. Stange. Frobenius and the endomorphism ring of $j = 1728$, 2021. `http://math.colorado.edu/~kstange/papers/1728.pdf`.

[52] A. V. Sutherland. Isogeny volcanoes. In *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium*, volume 1 of *Open Book Ser.*, pages 507–530. Math. Sci. Publ., Berkeley, CA, 2013.

[53] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 9.4)*, 2022. `https://www.sagemath.org`.

[54] J.-P. Tillich and G. Zémor. Optimal cycle codes constructed from Ramanujan graphs. *SIAM J. Discrete Math.*, 10(3):447–459, 1997.

[55] J. Voight. *Quaternion algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, Cham, [2021] ©2021.

[56] B. Wesolowski. Orientations and the supersingular endomorphism ring problem. In *Advances in Cryptology—EUROCRYPT 2022*, volume 13277 of *Lecture Notes in Comput. Sci.*, pages 345–371. Springer, Cham, 2022.

[57] B. Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021*, pages 1100–1111. IEEE Computer Soc., Los Alamitos, CA, [2022] ©2022.

[58] G. Xiao, L. Luo, and Y. Deng. Constructing cycles in isogeny graphs of supersingular elliptic curves. *J. Math. Cryptol.*, 15(1):454–464, 2021.

MATHEMATICS INSTITUTE, UNIVERSITEIT LEIDEN, LEIDEN, THE NETHERLANDS
Email address: `S.A.Arpin@math.leidenuniv.nl`

SCHOOL OF COMPUTER SCIENCE, UNIVERSITY OF BIRMINGHAM, UNIVERSITY ROAD WEST, BIRMINGHAM, UK B15 2TT
Email address: `m.chen.1@bham.ac.uk`

FACEBOOK AI RESEARCH, META, SEATTLE, WA
Email address: `klauter@fb.com`

DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF CALGARY, 2500 UNIVERSITY DRIVE NW, CALGARY, ALBERTA, CANADA T2N 1N4
Email address: `rscheidl@ucalgary.ca`

Department of Mathematics, University of Colorado, Campus Box 395, Boulder, Colorado 80309-0395
Email address: kstange@math.colorado.edu

Department of Mathematical and Physical Sciences, Concordia University of Edmonton, 7128 Ada Blvd NW, Edmonton, AB T5B 4E4, Canada
Email address: hatran1104@gmail.com