

Adventures in Supersingularland

Sarah Arpin, Catalina Camacho-Navarro, Kristin Lauter,
Joelle Lim, Kristina Nelson, Travis Scholl, Jana Sotáková

September 2019

Dedicated to Alice Silverberg

Abstract

In this paper, we study isogeny graphs of supersingular elliptic curves. Supersingular isogeny graphs were introduced as a hard problem into cryptography by Charles, Goren, and Lauter for the construction of cryptographic hash functions ([CGL06]). These are large expander graphs, and the hard problem is to find an efficient algorithm for routing, or path-finding, between two vertices of the graph. We consider four aspects of supersingular isogeny graphs, study each thoroughly and, where appropriate, discuss how they relate to one another.

First, we consider two related graphs that help us understand the structure: the ‘spine’ \mathcal{S} , which is the subgraph of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ given by the j -invariants in \mathbb{F}_p , and the graph $\mathcal{G}_\ell(\mathbb{F}_p)$, in which both curves and isogenies must be defined over \mathbb{F}_p . We show how to pass from the latter to the former. The graph \mathcal{S} is relevant for cryptanalysis because routing between vertices in \mathbb{F}_p is easier than in the full isogeny graph. The \mathbb{F}_p -vertices are typically assumed to be randomly distributed in the graph, which is far from true. We provide an analysis of the distances of connected components of \mathcal{S} .

Next, we study the involution on $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ that is given by the Frobenius of \mathbb{F}_p and give heuristics on how often shortest paths between two conjugate j -invariants are preserved by this involution (mirror paths). We also study the related question of what proportion of conjugate j -invariants are ℓ -isogenous for $\ell = 2, 3$. We conclude with experimental data on the diameters of supersingular isogeny graphs when $\ell = 2$ and compare this with previous results on diameters of LPS graphs and random Ramanujan graphs.

Contents

1	Introduction	3
2	Definitions	5
2.1	Isogeny Graphs	5
2.2	Special j -invariants	7
2.2.1	Self-isogenies	8
2.2.2	Double edges	8
3	Structure of the \mathbb{F}_p-subgraph: the spine \mathcal{S}	10
3.1	Structure of the \mathbb{F}_p -Graph $\mathcal{G}_\ell(\mathbb{F}_p)$	10
3.1.1	Preliminaries	10
3.1.2	The graph $\mathcal{G}_2(\mathbb{F}_p)$ in the case of $p \equiv 1 \pmod{4}$	14
3.1.3	The graph $\mathcal{G}_2(\mathbb{F}_p)$ in the case of $p \equiv 3 \pmod{4}$	15
3.2	Passing from the graph $\mathcal{G}_\ell(\mathbb{F}_p)$ to the spine $\mathcal{S} \subset \mathcal{G}_\ell(\overline{\mathbb{F}_p})$	16
3.3	Stacking, folding and attaching for $\ell > 2$	19
3.3.1	Example: stacking, folding and attaching for $\ell = 3$	20
3.4	Stacking, folding and attaching for $\ell = 2$	24
3.5	Distances of components of the \mathbb{F}_p -subgraph \mathcal{S}	30
3.5.1	$p \equiv 7 \pmod{8}$	30
3.5.2	The number of components	31
4	Conjugate vertices, distances, and the spine	31
4.1	Distance between conjugate pairs	31
4.2	How often do shortest paths go through the \mathbb{F}_p -spine	32
4.2.1	Experimental methods	32
4.2.2	Conjugate pairs vs arbitrary pairs	33
4.2.3	Proportions varying over different residue classes	34
4.3	Distance to spine	35
4.3.1	Comparison across primes p	36
5	When are conjugate j-invariants ℓ-isogenous?	36
5.1	Motivation	36
5.2	Methods	37
5.2.1	Timing data	37
5.3	Experimental data: 2-isogenies	38
5.3.1	Primes Modulo 12	39
5.4	Experimental data: 3-isogenies	39
5.4.1	Primes Modulo 12	41
5.5	Analysis of data	41
6	Diameter	42
6.1	Diameters of Primes Modulo 12	43
7	Conclusions	44

1 Introduction

Supersingular Isogeny Graphs have been the subject of recent study due to their significance in recently proposed post-quantum cryptographic protocols. In 2006, Charles, Goren, and Lauter proposed a hash function based on the hardness of finding paths (*routing*) in supersingular isogeny graphs [CGL06]. A few years later, Jao, De Feo, and Plut proposed a key exchange based on supersingular isogeny graphs [FJP11]. The security of most cryptographic systems currently deployed today relies on either the hardness of factoring large integers of a certain form or the hardness of computing discrete logarithms in certain abelian cyclic groups. Both problems can be efficiently solved using Shor’s algorithm on a quantum computer which can handle large scale computation [Sho99]. In 2015, NIST announced a contest to standardize cryptographic algorithms that are not known to be broken by quantum computers. Now in its second round, SIKE (<https://sike.org/>, based on supersingular isogeny graphs) is still in the running for the next public key exchange standard.

While there are no known classical or quantum attacks that break the cryptographic protocols that use supersingular isogeny graphs, the graphs themselves have been relatively unstudied until recently. More study is needed before we can confidently recommend protocols which rely on the difficulty of the hard problem of finding paths in supersingular isogeny graphs.

For distinct primes p and ℓ , let $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$ denote the graph whose vertices consist of isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$ and whose edges correspond to isogenies of degree ℓ defined over $\overline{\mathbb{F}}_p$. The vertices can be labelled with the j -invariant of the curve, which is an $\overline{\mathbb{F}}_p$ -isomorphism invariant. For $p \equiv 1 \pmod{12}$ the graph $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$ is known to be a $\ell + 1$ -regular Ramanujan graph, and is one of two known families of Ramanujan graphs.

In this paper, we study two related graphs to help understand the structure of $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$. First, the full subgraph of $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$ consisting of only vertices $j \in \mathbb{F}_p$: We denote this subgraph by \mathcal{S} and call it the *spine*, which is new terminology. Second, we look at the graph $\mathcal{G}_\ell(\mathbb{F}_p)$ whose vertices are elliptic curves up to \mathbb{F}_p -isomorphism and edges are \mathbb{F}_p -isogenies of degree ℓ , already studied by Delfs and Galbraith ([DG16]). As we will need to be specific about the field of definition, we use j to denote a general j -invariant, \mathbf{j} to denote a j -invariant in \mathbb{F}_p , and \mathbf{j} to denote a j -invariant in $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$. Note that if two elliptic curves \mathbb{F}_p are twists of each other, then they share the same j -invariant. A more formal discussion of the relationship between these graphs can be found in Section 2.

There have been several approaches tried so far to attack cryptographic protocols based on supersingular isogeny graphs. One of them uses the quaternion analogue of the graph and presents an efficient algorithm for navigating between maximal orders [KLPT14]. This approach leads to the results presented in [EHL⁺18] showing that the hardness of the path finding problem is essentially equivalent to the hardness of computing endomorphism rings of supersingular elliptic curves. One of the other methods considered so far uses the structure of the $\mathcal{G}_\ell(\mathbb{F}_p)$ ([DG16]). Better quantum algorithms are known for navigating between \mathbb{F}_p -points, so paths to these points are of particular interest.

In Section 3 of this paper, we compare $\mathcal{G}_\ell(\mathbb{F}_p)$ and the spine \mathcal{S} . The main results of Section 3 show how the components of $\mathcal{G}_\ell(\mathbb{F}_p)$ (which look like volcanoes) fit together to form the spine when passing to the full graph $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$. We define the notions of *stacking*, *folding*, and *attaching* to describe how $\mathcal{G}_\ell(\mathbb{F}_p)$ becomes the spine, when isogenies defined over $\overline{\mathbb{F}}_p$ are added and j -invariants which are twists are identified. In particular for $\ell = 2$, Theorem 3.26 shows that only stacking, folding, or at most one attachment by a new edge are possible to form the spine. Theorem 3.24 gives an analogous description for $\ell = 3$. For any fixed ℓ and p , the resulting shape of the spine depends on the congruence class of p , the structure of the class group $\text{Cl}(\mathcal{O}_K)$, where $K = \mathbb{Q}(\sqrt{-p})$, and the behavior of the prime above ℓ in the class group of K , and we show how to determine it explicitly. In Section 3.5 we generate experimental data to study how the components of the spine are distributed throughout the graph and we estimate how many components there are in the spine.

Another important property of Supersingular Isogeny Graphs is that they have an involution which fixes the spine, and which sends a \mathbf{j} -invariant in \mathbb{F}_{p^2} to its Galois conjugate \mathbf{j}^p . If a \mathbf{j} -

invariant in \mathbb{F}_{p^2} has a short path to the spine, then the involution can be applied to that path to produce a path from j to its conjugate j^p . We call such paths *mirror-paths*. In Section 4.1 we study the distance between Galois conjugate pairs of vertices, that is, pairs of j -invariants of the form j, j^p . Our data suggests these vertices are closer to each other than a random pair of vertices in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. In Section 4.2 we test how often the shortest path between two conjugate vertices goes through the spine \mathcal{S} , or equivalently, contains a j -invariant in \mathbb{F}_p . We find conjugate vertices are more likely than a random pair of vertices to be connected by a shortest path through the spine. Finally, we examine the distance between arbitrary vertices and the spine \mathcal{S} in Section 4.3.

Section 5 provides heuristics on how often conjugate j -invariants are ℓ -isogenous for $\ell = 2, 3$, a question motivated by the study of mirror paths provided in Section 4.

Another known family of Ramanujan graphs are certain Cayley graphs constructed by Lubotzky-Philips-Sarnak (LPS graphs [LPS88]). The relationship between LPS graphs and supersingular isogeny graphs is studied in [CFL⁺18]. Sardari [Sar19] provides an analysis of the diameters of LPS graphs, and in Section 6 of this paper, we provide heuristics and a discussion of the diameters of supersingular 2-isogeny graphs.

Our experiments and data suggest a noticeable difference in the Supersingular Isogeny Graphs $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ depending on the congruence class of the prime $p \pmod{12}$. It has been known since the introduction of Supersingular Isogeny Graphs into cryptography [CGL06] that the congruence class of the prime p has an important role to play in the properties of the graph. In particular it was shown there how the existence of short cycles in the graph depends explicitly on the congruence conditions on p . In this paper, we extend this observation and find significant differences in the graphs depending on the congruence class of p . In summary, the data seems to suggest the following:

- $p \equiv 1, 7 \pmod{12}$:
 - smaller 2-isogeny graph diameters (Section 6.1),
 - larger number of spinal components (Section 3.5.2),
 - larger proportion of 2-isogenous conjugate pairs (Section 5.3.1),
- $p \equiv 5, 11 \pmod{12}$:
 - larger 2-isogeny graph diameters,
 - smaller number of spinal components,
 - smaller proportion of 2-isogenous conjugate pairs.

To accompany the experimental results of this paper, we have made the **Sage** code for all the computations available, along with a short discussion of the different algorithms included. The code is posted at

<https://github.com/krstnmnl/sn/Adventures-in-Supersingularland-Data>.

Acknowledgements

This paper is the result of a collaboration started at Alice Silverberg’s 60th birthday conference on open questions in cryptography and number theory <https://sites.google.com/site/silverberg2018/>. We would like to thank Heidi Goodson for her significant contributions to this project. We are grateful to Microsoft Research for hosting the authors for a follow-up visit. We would like to thank Steven Galbraith, Shahed Sharif, and Katherine E. Stange for helpful conversations. Travis Scholl was partially supported by Alfred P. Sloan Foundation grant number G-2014-13575. Catalina Camacho-Navarro was partially supported by Universidad de Costa Rica.

2 Definitions

Definition 2.1. *An elliptic curve E is a smooth, projective algebraic curve of genus 1 with a fixed point, denoted \mathcal{O}_E .*

Elliptic curves have a group law, which makes them particularly rich objects to work with. For more background on this, see [Sil09]. Elliptic curves defined over fields of characteristic $p < \infty$ come in two flavors: ordinary and supersingular. In the graphs of elliptic curves considered here, the ordinary and supersingular components are disjoint. We focus on the graph of supersingular elliptic curves, defined here as in Theorem 3.1 of [Sil09]:

Definition 2.2 (Supersingular Elliptic Curve). *Let E/K be an elliptic curve, $\text{char}(K) = p < \infty$. We say that E is supersingular if any of the following equivalent conditions hold:*

- (i) *The p^k -torsion of E is trivial for all $k \in \mathbb{Z}_{k \geq 1}$.*
- (ii) *The multiplication by p -map on E is purely inseparable and the j -invariant of E is in \mathbb{F}_{p^2}*
- (iii) *The endomorphism ring of E is isomorphic to a maximal order in a quaternion algebra.*

Elliptic curves which are not supersingular are called ordinary.

By definition, supersingular elliptic curves can all be identified with a j -invariant in \mathbb{F}_{p^2} . The j -invariants classify the $\overline{\mathbb{F}}_p$ -isomorphism classes of supersingular elliptic curves. When we consider supersingular elliptic curves with $j(E) \in \mathbb{F}_p$ up to \mathbb{F}_p -isomorphism, two \mathbb{F}_p -isomorphism classes will have a single j -invariant ([DG16, Prop. 2.3]). One can distinguish between the two curves for instance by considering Weierstrass models of the elliptic curves.

Whenever the field of definition of the j -invariant is relevant and not clear from context, we will use the following notation:

- \mathbf{j} denotes a j -invariant in \mathbb{F}_p
- \mathbf{j} denotes a j -invariant in $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$.

Otherwise, or for a general j -invariant, we denote it simply by j .

A characterizing difference between the endomorphism rings of \mathbb{F}_p j -invariants and $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ j -invariants is pointed out in [DG16, Prop. 2.4]: for $p > 3$, a supersingular elliptic curve $E/\overline{\mathbb{F}}_p$ is defined over \mathbb{F}_p if and only if $\mathbb{Z}[\sqrt{-p}] \subset \text{End}(E)$.

2.1 Isogeny Graphs

There are three graphs to consider. To introduce these graphs, we borrow the following notions from Sutherland [Sut13, Section 2.2]. We denote by $\Phi_\ell[X, Y]$ the ℓ -modular polynomial. This is a polynomial of degree $\ell + 1$ in both X and Y , symmetric in X and Y and such that there exists a cyclic ℓ -isogeny $\phi : E(j_1) \rightarrow E(j_2)$ if and only if $\Phi_\ell(j_1, j_2) = 0$. For ℓ prime, all isogenies are cyclic.

In principle, the modular polynomials can be computed and they are accessible via tables for small values of ℓ , however, their coefficients are rather large, as we see already for $\phi_2(X, Y)$:

$$\begin{aligned} \Phi_2(X, Y) = & -X^2Y^2 + X^3 + 1488X^2Y + 1488XY^2 + Y^3 - 162000X^2 + 40773375XY \\ & - 162000Y^2 + 8748000000X + 8748000000Y - 15746400000000 \end{aligned} \quad (1)$$

Definition 2.3 (Supersingular ℓ -isogeny graph over $\overline{\mathbb{F}}_p$: $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$). *The $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$ graph has vertex set consisting of the $\overline{\mathbb{F}}_p$ -isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$, labeled by their j -invariants over $\overline{\mathbb{F}}_p$. The directed edges from a vertex j correspond to (j, j') where j' is a root of the modular polynomial $\Phi_\ell(j, Y)$.*

Except possibly at vertices corresponding to $\mathbf{j} = 0, 1728$, this defines an $\ell + 1$ regular graph. These graphs are known to be Ramanujan graphs (see [CGL06] or [CFL⁺18]).

Definition 2.4 (Spine: \mathcal{S}). *The \mathbb{F}_p spine, denoted \mathcal{S} , is the full subgraph of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ consisting of all vertices with j -invariants defined over \mathbb{F}_p and all their edges in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.*

The number of vertices in \mathcal{S} can be determined from [Cox89]

$$\#\mathcal{S} = \begin{cases} \frac{1}{2}h(-4p) & \text{if } p \equiv 1 \pmod{4} \\ h(-p) & \text{if } p \equiv 7 \pmod{8} \\ 2h(-p) & \text{if } p \equiv 8 \pmod{8}. \end{cases}$$

where $h(d)$ is the class number of the imaginary quadratic field $\mathbb{Q}(\sqrt{d})$.

Definition 2.5 (Supersingular ℓ -isogeny graph over \mathbb{F}_p : $\mathcal{G}_\ell(\mathbb{F}_p)$). *The $\mathcal{G}_\ell(\mathbb{F}_p)$ has vertex set \mathbb{F}_p -isomorphism classes of j -invariants $\mathbf{j} \in \mathbb{F}_p$. The edges correspond to ℓ -isogenies defined over \mathbb{F}_p as well. As noted before, each j -invariant will appear as two distinct vertices in this graph.*

Remark 2.6. *It is worthwhile to highlight the differences between \mathcal{S} and $\mathcal{G}_\ell(\mathbb{F}_p)$:*

- \mathcal{S} has fewer vertices than $\mathcal{G}_\ell(\mathbb{F}_p)$, since the vertices are considered up to $\overline{\mathbb{F}_p}$ -isomorphism in the former and \mathbb{F}_p -isomorphism in the later.
- \mathcal{S} has (likely) more edges than $\mathcal{G}_\ell(\mathbb{F}_p)$, since we consider the edges defined over $\overline{\mathbb{F}_p}$ in the former but only those defined over \mathbb{F}_p in the later. The ‘‘appearance’’ of these edges when we move from $\mathcal{G}_\ell(\mathbb{F}_p)$ to \mathcal{S} will be discussed more thoroughly in the sequel.

Remark 2.7. *Note that we can consider these graphs can be considered to be un-directed except at the j -invariants $\mathbf{j} = 0, 1728$: Every ℓ -isogeny $\phi : E \rightarrow E'$ has a dual $\widehat{\phi} : E' \rightarrow E$ of the same degree. The only issues we run into with $\mathbf{j} = 0, 1728$ are the extra automorphisms of these curves can compose with the isogenies, affecting the regularity of the graphs at these vertices. We can still consider the graph to be undirected at these vertices, but we will not preserve the multiplicity of the edges with this relaxation.*

Definition 2.8. *We say two isogenies $\phi : E_1 \rightarrow E_2$ and $\phi' : E'_1 \rightarrow E'_2$, are equivalent over k if there exist isomorphisms $\varphi : E'_1 \rightarrow E_1$ and $\psi : E_2 \rightarrow E'_2$ over k such that $\phi' = \psi \circ \phi \circ \varphi$.*

Let us make some remarks about $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. By definition, this is an $\ell + 1$ regular graph, where we can associate an edge $\{j_1, j_2\}$ to an equivalence class of isogenies between two elliptic curve E_1 and E_2 with $j_1 = j(E_1)$ and $j_2 = j(E_2)$. Kohel [Koh96, Chapter 7] proved that very pair of supersingular elliptic curves are connected by a chain of degree ℓ isogenies, which implies that the graph is connected. If $p > 3$, the number of vertices of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ is

$$\left\lfloor \frac{p}{12} \right\rfloor + \begin{cases} 0 & \text{if } p \equiv 1 \pmod{12}, \\ 1 & \text{if } p \equiv 5, 7 \pmod{12}, \\ 2 & \text{if } p \equiv 11 \pmod{12}. \end{cases}$$

(See [Sil09, Section V.4].) The congruence condition follows from whether or not the j -invariants 0 and 1728 are supersingular or not.

Remark 2.9. *The graph $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ is a component (called the supersingular component) of the general ℓ -isogeny graph where the vertices also include the ordinary j -invariants and ℓ -isogenies between them. Isogenies preserve the properties of ‘‘being ordinary’’ and ‘‘being supersingular’’, so these vertices do not mix on connected components of the full ℓ -isogeny graph.*

It is natural to consider connections between these three graphs. Moving from $\mathcal{G}_\ell(\mathbb{F}_p)$ to \mathcal{S} identifies vertices with the same j -invariant and adds edges. To move from \mathcal{S} to $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$, we can

consider adding j -invariants in conjugate pairs: starting with \mathbf{j} in \mathcal{S} , if there is an isogeny from $E(\mathbf{j})$ to $E(j)$, there is a conjugate isogeny from $E(\mathbf{j})$ to the conjugate $E(j)^{(p)}$.

Indeed, this works for any two j -invariants j, j' : if j and j' satisfy $\Phi_\ell(j, j') = 0$ then also

$$\Phi_\ell(j^p, (j')^p) = (\Phi_\ell(j, j'))^p = 0$$

because Φ_ℓ has integer coefficients. This means that for any edge $[j, j'] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$, there is a *mirror* edge $[j^p, (j')^p]$. Constructing the graph from this perspective leads to the idea of a mirror involution on $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$:

Definition 2.10. *If j is a supersingular j -invariant, so is its \mathbb{F}_{p^2} -conjugate j^p (in the case that $j = \mathbf{j} \in \mathbb{F}_p$, $\mathbf{j}^p = j$). If there is an ℓ -isogeny $\phi : E(j_1) \rightarrow E(j_2)$ then there exists an ℓ -isogeny $\phi' : E(j_1)^{(p)} \rightarrow E(j_2)^{(p)}$. This implies that the (p) -power Frobenius map on \mathbb{F}_{p^2} gives an involution on $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. We call this the mirror involution.*

The mirror involution fixes the \mathbb{F}_p -vertices of the graph.

Definition 2.11. *We say that a path P with vertices $\{j_0, j_1, j_2, \dots, j_{n-1}, j_n\}$ (considered as an undirected path) is a mirror path if it is invariant under the mirror involution.*

There exists at least one mirror path between any two conjugate j -invariants. One way to find such a mirror path is to find a path from one j -invariant, say \mathbf{j}_0 , to an \mathbb{F}_p j -invariant. Then, conjugate that path to connect with the conjugate of \mathbf{j}_0 , which we denote \mathbf{j}_0^p . In summary, a path of the form:

$$\mathbf{j}_0 \rightarrow \mathbf{j}_1 \rightarrow \dots \rightarrow \mathbf{j}_n \rightarrow \mathbf{j} \rightarrow \mathbf{j}_n^p \rightarrow \dots \rightarrow \mathbf{j}_1^p \rightarrow \mathbf{j}_0^p$$

Another possibility is for a mirror path between conjugate j -invariants to pass through a pair of isogenous conjugate j -invariants:

$$\mathbf{j}_0 \rightarrow \mathbf{j}_1 \rightarrow \dots \rightarrow \mathbf{j}_n \rightarrow \mathbf{j}_n^p \rightarrow \dots \rightarrow \mathbf{j}_1^p \rightarrow \mathbf{j}_0^p$$

2.2 Special j -invariants

In this section, we establish a few general facts about j -invariants that require special attention.

First of all, there are j -invariants corresponding to curves with extra automorphisms that result in the undirectedness of the graph (see Remark 2.7). It is a standard fact that $j = 1728$ is supersingular if and only if $p \equiv 3 \pmod{4}$ and $j = 0$ is supersingular if and only if $p \equiv 2 \pmod{3}$. The computation can be easily argued by CM theory (see, for instance, [Igu58]). The main idea used is that the j -invariant of an elliptic curve E with CM by a quadratic order \mathcal{O} generates the ring class field of \mathcal{O} and such a curve reduces to a supersingular curve modulo p if and only if p is inert in \mathcal{O} .

Example 2.12. *Elliptic curves with j -invariant $j = 8000$ are supersingular over \mathbb{F}_p if and only if $p \equiv 5, 7 \pmod{8}$.*

Proof. The number field $\mathbb{Q}(\sqrt{-2})$ has Hilbert class polynomial $x - 8000$, meaning that the j -invariant $j = 8000$ has CM by $\mathbb{Q}(\sqrt{-2})$. This j -invariant will be supersingular whenever p is inert in $\mathbb{Q}(\sqrt{-2})$. Hence we only need to compute $\left(\frac{-2}{p}\right) = -1$, which gives the congruence conditions. \square

Example 2.13. *Elliptic curves with j -invariant $j = -3375$ are supersingular over \mathbb{F}_p if and only if $p \equiv 3, 5, 6 \pmod{7}$.*

Proof. The Hilbert class polynomial of $\mathbb{Q}(\sqrt{-7})$ is $x + 3375$. Hence $j = -3375$ will be supersingular over \mathbb{F}_p whenever p is inert in $\mathbb{Q}(\sqrt{-7})$. The rest follows from evaluating $\left(\frac{-7}{p}\right) = -1$. \square

2.2.1 Self-isogenies

Next we turn our attention to the self-loops in the graph $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$, which are easily read off from the factorization of $\Phi_\ell(X, X)$ as follows: A j -invariant j admits a self-isogeny if and only if $\Phi_\ell(j, j) = 0$.

For instance, consider the modular polynomial $\Phi_2(X, Y)$ (as seen in (1)). Now, $\Phi_2(X, X)$ factors over \mathbb{Z} as

$$\Phi_2(X, X) = -(X + 3375)^2(X - 1728)(X - 8000), \quad (2)$$

therefore, the only loops in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ are at the following vertices:

- $\mathbf{j} = -3375$ has two loops,
- $\mathbf{j} = 1728$ has one loop (we will show where this loop comes from in Example 3.8),
- $\mathbf{j} = 8000$ has one loop.

In particular, no j -invariant over $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ has a self isogeny. Note that these j -invariants may not be supersingular, so they may not appear on $\mathcal{G}_2(\overline{\mathbb{F}_p})$ for every p .

2.2.2 Double edges

We use the following general lemma about double edges in the 2-isogeny graph. Note that this lemma applies mutatis mutandis for ordinary curves, replacing $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ by the ℓ -isogeny graph of ordinary elliptic curves.

Lemma 2.14 (Double edge lemma). *If two j -invariants j_1, j_2 in the ℓ -isogeny graph $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ have a double-edge between them, then they are roots of the polynomial*

$$\text{Res}_\ell(X) := \text{Res} \left(\Phi_\ell(X, Y), \frac{d}{dY} \Phi_\ell(X, Y); Y \right) \quad (3)$$

which is a polynomial of degree bounded by $2\ell \cdot (2\ell - 1)$.

Proof of the double edge lemma. Suppose that j_1 and j_2 are two vertices in the ℓ -isogeny graph connected with a double edge. Considered as a polynomial in Y , this means

$$\Phi_\ell(j_1, Y) = (Y - j_2)^2 \cdot g(j_1, Y)$$

for some $g(j_1, Y) \in \overline{\mathbb{F}_p}[Y]$. The derivative

$$\frac{d}{dY} \Phi_\ell(j_1, Y)$$

with respect to Y then also vanishes at $Y = j_2$. This means that the polynomials $\Phi_\ell(X, Y)$ and $\frac{d}{dY} \Phi_\ell(X, Y)$ share a root when plugging in $X = j_1$. But this means that j_1 is a root of the resultant

$$\text{Res} \left(\Phi_\ell(X, Y), \frac{d}{dY} \Phi_\ell(X, Y); Y \right).$$

Since the total degree of $\Phi_\ell(X, Y)$ is 2ℓ and the total degree of $\frac{d}{dY} \Phi_\ell(X, Y)$ is $2\ell - 1$ and the resultant of two polynomials $P(X, Y)$ and $Q(X, Y)$ of total degrees d and e has generically degree $d \cdot e$, we obtain the bound

$$\#\{j : \text{there is a double edge from } j\} \leq 2\ell \cdot (2\ell - 1)$$

The bound in the lemma is not tight as we will see in the following corollary for $\ell = 2$.

Corollary 2.15 (Double edges for $\ell = 2$). *If $\ell = 2$ and there is a double edge from j in $\mathcal{G}_2(\overline{\mathbb{F}_p})$, then the j -invariant j is in the following list:*

- $\mathbf{j} = 0$,

- $\mathbf{j} = 1728$,
- $\mathbf{j} = -3375$,
- j is a root of $X^2 + 191025X - 121287375$.

Moreover, at $\mathbf{j} = 0$ we obtain a triple edge and $\mathbf{j} = 0$ is the only j -invariant that admits a triple edge in $\mathcal{G}_2(\overline{\mathbb{F}_p})$.

Proof. By lemma 2.14, double edges in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ can only occur at the roots of

$$\text{Res}_2(X) = (-1) \cdot 2^2 \cdot X^2 \cdot (X - 1728) \cdot (X + 3375)^2 \cdot (X^2 + 191025X - 121287375)^2 \quad (4)$$

For the j -invariants 0, 1728 and -3375 , we identify the double edges by factoring $\Phi_2(j, X)$.

1. For $\mathbf{j} = 0$, we have

$$\Phi_2(0, X) = (X - 54000)^3.$$

There are three outgoing 2-isogenies from $\mathbf{j} = 0$ to $\mathbf{j} = 54000$.

2. For $\mathbf{j} = 1728$, we have

$$\Phi_2(1728, X) = (X - 1728) \cdot (X - 287496)^2$$

there is always a self-2-isogeny (explained further in 3.8) and two 2-isogenies to $\mathbf{j} = 287496$. These are defined over \mathbb{F}_p for any p : from the model $y^2 = x^3 - x$, they are given by maps

$$\left(\frac{x^2 + x + 2}{x + 1}, \frac{x^2 y + 2xy - y}{x^2 + 2x + 1} \right) \quad \text{and} \quad \left(\frac{x^2 - x + 2}{x - 1}, \frac{x^2 y - 2xy - y}{x^2 - 2x + 1} \right).$$

3. For $\mathbf{j} = -3375$, we have

$$\Phi_2(-3375, X) = (X - 16581375) \cdot (X + 3375)^2$$

and so there are always two self-2-isogenies.

We also note that $j = 0$ is the only \mathbf{j} -invariant that can admit a triple edge. Indeed, since away from the vertices 1728 and 0 we can think of the graph as being undirected with 3 edges from every vertex, having a triple edge would mean having two isolated vertices in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. This is not possible. \square

The double-edges from Corollary 2.15 appear in the supersingular 2-isogeny graph only when these j -invariants are supersingular.

Remark 2.16. *The factors of the polynomial $\text{Res}_2(X)$ (as seen in (4)) are Hilbert class polynomials for imaginary quadratic fields. This is to be expected: A double edge $[j_1, j_2]$ is a 2-cycle of non-dual 2-isogenies (not equal to the multiplication map [2]). The ring $\text{End}_{\mathbb{F}_p}(E_{J_1})$ has a non-trivial element of norm 4 corresponding to this 2-cycle. The only quadratic imaginary fields that contain an element of norm 4 are $\mathbb{Q}(\sqrt{-1}), \mathbb{Q}(\sqrt{-3}), \mathbb{Q}(\sqrt{-7}), \mathbb{Q}(\sqrt{-15})$.*

Remark 2.17. *The above remark generalizes for any ℓ : the polynomial*

$$\text{Res}_\ell(X)$$

is a product of (ring) class polynomials of quadratic orders containing a nontrivial element of norm ℓ^2 .

Short cycles are considered carefully in Section 6 of [CGL06]: a sufficient condition so that there are no cycles of length 2 in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ is $p \equiv 1 \pmod{420}$.

We will use Lemma 2.14 together with Lemma 3.15 that if there is an edge in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ between two j -invariants $\mathbf{j}_1, \mathbf{j}_2 \in \mathbb{F}_p$ that corresponds to an isogeny which is not defined over \mathbb{F}_p (i.e., which is not an edge in $\mathcal{G}_\ell(\mathbb{F}_p)$), then there is a double edge between \mathbf{j}_1 and \mathbf{j}_2 in $\mathcal{G}_2(\overline{\mathbb{F}_p})$, and hence \mathbf{j}_1 and \mathbf{j}_2 are among the values listed in Lemma 2.14.

3 Structure of the \mathbb{F}_p -subgraph: the spine \mathcal{S}

In this section, we investigate the shape of the spine \mathcal{S} , which was defined in 2.4 to be the subgraph of $\mathcal{G}_\ell(\mathbb{F}_p)$ consisting of the vertices defined over \mathbb{F}_p . The motivation for studying the structure of this subgraph is the existence of attacks on SIDH that work by finding \mathbb{F}_p \mathbf{j} -invariants: the idea for such a possible attack was presented in [DG16], and a quantum attack based on this idea was given by Biasse, Jao and Sankar [BJS14]. See also [GPST16] for an overview of the security considerations for SIDH.

In Section 3.1, we start with the graph $\mathcal{G}_\ell(\mathbb{F}_p)$, the structure of which is understood well. It was studied in depth by Delfs and Galbraith in [DG16], on which we base our investigations. Moreover, a version of the \mathbb{F}_p -graph (allowing edges corresponding to ℓ -isogenies for multiple primes) has also been proposed for post-quantum cryptography in [CLM⁺18]. We recall the results of [DG16] in some detail and give a few explicit examples of their results about endomorphisms of certain elliptic curves (for instance, Example 3.6).

In Section 3.2, we discuss how the Spine \mathcal{S} can be obtained from $\mathcal{G}_\ell(\mathbb{F}_p)$ in two steps: first vertices corresponding to the same j -invariant are identified and then a few new edges are added. The possible ways the connected components can identify are given by Definition 1 and we call them *stacking*, *folding*, *attaching along a j -invariant* and *attaching by a new edge*.

In Section 3.3 we study stacking, folding and attaching for $\ell > 2$ and give an example of the theory we develop for $\ell = 3$ in Section 3.3.1. In this section we also give a complete description of stacking, folding and attaching in Theorem 3.24. We return to the case $\ell = 2$ in 3.4, giving a similarly complete theorem in 3.26, and some data on how often attachment happens. Section 3.5 contains some experimental data on the distances between the connected components of $\mathcal{S} \subset \mathcal{G}_2(\mathbb{F}_p)$.

3.1 Structure of the \mathbb{F}_p -Graph $\mathcal{G}_\ell(\mathbb{F}_p)$

3.1.1 Preliminaries

To understand the spine \mathcal{S} (Definition 2.4), we look at $\mathcal{G}_\ell(\mathbb{F}_p)$ (Definition 2.5). Recall that the vertices of $\mathcal{G}_\ell(\mathbb{F}_p)$ are all supersingular elliptic curves defined over \mathbb{F}_p , up to \mathbb{F}_p -isomorphism, and the edges in $\mathcal{G}_\ell(\mathbb{F}_p)$ are isogenies defined over \mathbb{F}_p . Keep in mind the differences between \mathcal{S} and $\mathcal{G}_\ell(\mathbb{F}_p)$, highlighted in Remark 2.6.

To see how many vertices of $\mathcal{G}_\ell(\mathbb{F}_p)$ correspond to the same j -invariant, we look at twists of elliptic curves. By Proposition 5.4 of [Sil09][Chapter X], for $\mathbf{j} \neq 0, 1728$, the set of twists is isomorphic to (assuming $p > 3$)

$$\mathbb{F}_p^*/(\mathbb{F}_p^*)^2 \cong \mathbb{Z}/2\mathbb{Z},$$

so there are two vertices corresponding to the same j -invariant \mathbf{j} .

Similarly, for $\mathbf{j} = 1728$, the set of twists is isomorphic to

$$\mathbb{F}_p^*/(\mathbb{F}_p^*)^4.$$

The j -invariant $\mathbf{j} = 1728$ is supersingular if and only if $p \equiv 3 \pmod{4}$ (equivalently, $p-1 \equiv 2 \pmod{4}$), so $(\mathbb{F}_p^*)^4 = (\mathbb{F}_p^*)^2$. Hence, there are two vertices of $\mathcal{G}_\ell(\mathbb{F}_p)$ corresponding to $\mathbf{j} = 1728$, as well. These vertices correspond to quartic twists, rather than quadratic twists, which we will use in Example 3.8.

For $\mathbf{j} = 0$, the set of twists is isomorphic to

$$\mathbb{F}_p^*/(\mathbb{F}_p^*)^6.$$

We know that $\mathbf{j} = 0$ is supersingular if and only if $p \equiv 2 \pmod{3}$ (equivalently $p-1 \equiv 1 \pmod{3}$), so $(\mathbb{F}_p^*)^6 = (\mathbb{F}_p^*)^2$. Hence there are also two vertices of $\mathcal{G}_\ell(\mathbb{F}_p)$ corresponding to $\mathbf{j} = 0$.

The structure of $\mathcal{G}_\ell(\mathbb{F}_p)$ is explained in [DG16]. They show that the $\mathcal{G}_\ell(\mathbb{F}_p)$ graph looks very similar to an isogeny graph of ordinary curves. Upon recalling some of the main definitions, we

present a simplified version of their construction results here, restricting many general results to the case $\ell = 2$.

Let $K := \mathbb{Q}(\sqrt{-p})$. Start with the definition of supersingularity for elliptic curves over \mathbb{F}_p : an elliptic curve E/\mathbb{F}_p is supersingular if and only if $\mathbb{Z}[\sqrt{-p}] \subset \text{End}_{\mathbb{F}_p}(E)$. If $p \equiv 3 \pmod{4}$, the order $\mathbb{Z}[\sqrt{-p}]$ is contained in the maximal order $\mathbb{Z}\left[\frac{1+\sqrt{-p}}{2}\right]$. To distinguish between these two possible endomorphism rings, we have the following definitions.

Definition 3.1 (Surface and Floor). *Let E be a supersingular elliptic curve over \mathbb{F}_p . We say E is on the surface (resp. E is on the floor) if $\text{End}_{\mathbb{F}_p}(E) = \mathcal{O}_K$ (resp. $\text{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[\sqrt{-p}]$). For $p \equiv 1 \pmod{4}$, surface and floor coincide.*

Definition 3.2 (Horizontal and Vertical Isogenies.). *Let φ be an ℓ -isogeny between supersingular elliptic curves E and E' over \mathbb{F}_p . If $\text{End}_{\mathbb{F}_p}(E) \cong \text{End}_{\mathbb{F}_p}(E')$ then φ is called horizontal. Otherwise, if E is on the floor and E' is on the surface, or vice versa, φ is called vertical.*

The following is Theorem 2.7 of [DG16]. We revisit this result in Sections 3.1.2 and 3.1.3.

Theorem 3.3. *Let $p > 3$ be a prime.*

1. *For $\ell > 2$, there are two horizontal isogenies from any vertex and there are no vertical isogenies, provided $\left(\frac{-p}{\ell}\right) = 1$, otherwise there are no ℓ -isogenies. Hence every connected component of $\mathcal{G}_\ell(\mathbb{F}_p)$ is a cycle.*
2. *Case $p \equiv 1 \pmod{4}$. There is one level in $\mathcal{G}_2(\mathbb{F}_p)$: all elliptic curves E have $\text{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[\sqrt{-p}]$. For $\ell = 2$: from each vertex there is one outgoing \mathbb{F}_p -rational 2-isogeny. There are $h(-4p)$ vertices on the surface (which coincides with the floor).*
3. *Case $p \equiv 3 \pmod{4}$. There are two levels in $\mathcal{G}_2(\mathbb{F}_p)$: surface and floor. For $\ell = 2$:*
 - (a) *If $p \equiv 7 \pmod{8}$, there is exactly one vertical isogeny from any vertex on the surface to a vertex on the floor, every vertex on the surface admits two horizontal isogenies and there are no horizontal isogenies between the curves on the floor. There are $h(-p)$ vertices on the floor and $h(-p)$ vertices on the surface.*
 - (b) *If $p \equiv 3 \pmod{8}$, from every vertex on the surface, there are exactly three vertical isogenies to the floor, and there are no horizontal isogenies between any vertices. There are $3 \cdot h(-p)$ vertices on the floor and $h(-p)$ vertices on the surface.*

This implies that every connected component of $\mathcal{G}_2(\mathbb{F}_p)$ is an isogeny volcano, first studied by Kohel [Koh96]. For a reference on the name and basic properties we refer to [Sut13].

Proof. Theorem 2.7 in [DG16]. We will reference the methods in the proof:

1. There is a one-to-one correspondence between supersingular elliptic curves over \mathbb{F}_p and elliptic curves defined over \mathbb{C} with CM by $\mathcal{O} \in \{\mathbb{Z}[\sqrt{-p}], \mathbb{Z}\left[\frac{1+\sqrt{-p}}{2}\right]\}$.
2. Isogeny graphs of CM curves have a volcano structure, and the edges of the volcano reduce to edges in the graph $\mathcal{G}_\ell(\mathbb{F}_p)$. Hence, there will be a volcano-like structure over \mathbb{F}_p .
3. The reduction does not add any more edges. For $\ell = 2$ we reprove the main ingredient in Lemma 3.4. Adding more edges between \mathbb{F}_p vertices would imply that $E[\ell] \subset E(\mathbb{F}_p)$ and this cannot happen for $\ell > 2$ because the curves are supersingular.

Hence we will see the volcano structure over \mathbb{F}_p . □

Let $K = \mathbb{Q}(\sqrt{-p})$, \mathfrak{p} be a prime above $\ell = 2$ in \mathcal{O}_K , and $h = \#\text{Cl}(\mathcal{O}_K)$ the class number of K . Let n be the order of \mathfrak{p} in $\text{Cl}(\mathcal{O}_K)$. The surface of any volcano in $\mathcal{G}_2(\mathbb{F}_p)$ is a cycle of precisely n vertices. There are h/n connected components (volcanoes) in $\mathcal{G}_2(\mathbb{F}_p)$, the index of $\langle \mathfrak{p} \rangle$ in $\text{Cl}(\mathcal{O}_K)$.

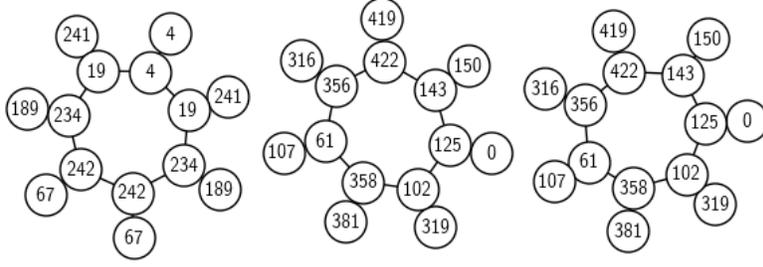


Figure 3.1: The graph $G_2(\mathbb{F}_p)$ for $p = 431$. The integer labels are the j -invariants of each curve. Each component is a volcano, with an inner ring of surface curves and the outer vertices all being curves on the floor. $431 \equiv 7 \pmod{8}$, so we are in case 2.a of Theorem 3.3. The class number of $\mathbb{Q}(\sqrt{-431})$ is $3 \cdot 7 = 21$ and the orders of the two primes above 2 are 7.

Lemma 3.4. Let p be a prime, $p > 5$. Let E be a supersingular elliptic curve defined over \mathbb{F}_p . Then

$$\text{End}_{\mathbb{F}_p}(E) = \mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right] \quad \text{if and only if} \quad E[2] \subset E(\mathbb{F}_p).$$

For $p \equiv 1 \pmod{4}$, the ring $\mathbb{Z}[(1 + \sqrt{-p})/2]$ is not an order of \mathcal{O}_K , so supersingular elliptic curves in \mathbb{F}_p do not have their full two-torsion defined over \mathbb{F}_p .

Proof. This proof is an adaptation of the techniques on page 7 of [DG16].

Let E be any supersingular elliptic curve defined over \mathbb{F}_p . Then $\#E(\mathbb{F}_p) = p + 1$ and the minimal polynomial of Frobenius π is $x^2 + p$. This implies $\pi = \pm\sqrt{-p} \in \mathbb{Z}(\sqrt{-p})$. We have

$$\mathbb{Q}(\sqrt{-p}) \supseteq \text{End}_{\mathbb{F}_p}(E) \supset \mathbb{Z}[\pi] = \mathbb{Z}[\sqrt{-p}].$$

Thus, $\text{End}_{\mathbb{F}_p}(E)$ is either isomorphic to $\mathbb{Z}[\sqrt{-p}]$ or $\mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right]$, as an order in $\mathbb{Q}(\sqrt{-p})$.

First, suppose $E[2] \subset E(\mathbb{F}_p)$. Take $P \in E[2]$. Frobenius acts as the identity on the 2-torsion:

$$\pi(P) = P \Rightarrow (1 + \pi)(P) = 0,$$

where 0 denotes the identity of E , since $-P = P$ for $P \in E[2]$. Hence $E[2] \subset \ker(1 + \pi)$. Isogenies have the universal property of a quotient, so we obtain the factorization

$$\begin{array}{ccc} E & \xrightarrow{[2]} & E \\ & \searrow^{1+\pi} & \nearrow^{\phi} \\ & & E \end{array}$$

and conclude $1 + \pi = [2]\phi$. The map $\phi = \frac{1 + \pi}{2}$ is \mathbb{F}_p -rational, since it is the quotient of \mathbb{F}_p -rational maps, so $\frac{1 + \pi}{2} \in \text{End}_{\mathbb{F}_p}(E) \cong \mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right]$.

Conversely, suppose $\text{End}_{\mathbb{F}_p}(E) \cong \mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right]$. Consider $\phi = \frac{1 - \pi}{2} \in \text{End}_{\mathbb{F}_p}(E) \cong \mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right]$, where π is Frobenius. Take any $P \in E[2]$:

$$(1 - \pi)(P) = 2\phi(P) = \phi \cdot 2(P) = 0 \implies 1 = \pi \quad \text{on } E[2].$$

Frobenius acts trivially on $E[2]$, so we have $E[2] \subset E(\mathbb{F}_p)$. □

The following corollary of Lemma 3.4 will be essential in our discussion in Section 3.4.

Corollary 3.5 (Endomorphism rings of quadratic twists). *Let $p > 3$ be a prime, let E be an elliptic curve defined over \mathbb{F}_p and let E^t denote its quadratic twist. Then*

$$\text{End}_{\mathbb{F}_p}(E) \cong \text{End}_{\mathbb{F}_p}(E^t).$$

Proof. Suppose E is given by the equation

$$E : y^2 = x^3 + ax + b.$$

Let d be a quadratic non-residue modulo p . Then the quadratic twist is given by the equation

$$E^t : y^2 = x^3 + d^2ax + d^3b$$

and the isomorphism $E \rightarrow E^t$ defined over \mathbb{F}_{p^2} is given by

$$(x, y) \mapsto \left(\frac{x}{d}, \frac{y}{d\sqrt{d}} \right).$$

2-torsion points (x, y) satisfy $y = 0$, so $E[2] \subset E(\mathbb{F}_p)$ if and only if $E^t[2] \subset E^t(\mathbb{F}_p)$. The result follows from Lemma 3.4. \square

Example 3.6 ($j = 0$ is always on the floor). *Suppose $p > 3$. Any supersingular elliptic curve E_0/\mathbb{F}_p with j -invariant 0 satisfies $\text{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[\sqrt{-p}]$.*

Proof. E_0 is supersingular if and only if $p \equiv 2 \pmod{3}$, so $p \equiv 2 \pmod{3}$. Take a short Weierstrass model $E_0 : y^2 = x^3 - d$. By inspection,

$$E[2] \subset E(\mathbb{F}_p) \iff x^3 - d \text{ splits completely} \iff \zeta_3 \in \mathbb{F}_p \iff 3|p - 1,$$

where ζ_3 denotes a 3rd root of unity. However, we have $p \equiv 2 \pmod{3}$, so $E[2]$ is not defined over \mathbb{F}_p . By Lemma 3.4 we have $\text{End}_0(E) = \mathbb{Z}[\sqrt{-p}]$. \square

Example 3.7 ($j = -3375$ is always on the floor). *Let $p > 3$ be a prime. Any supersingular elliptic curve E_{-3375}/\mathbb{F}_p with j -invariant -3375 satisfies $\text{End}_{\mathbb{F}_p}(E_{-3375}) = \mathbb{Z}[\sqrt{-p}]$.*

Proof. We have $\Phi_2(-3375, x) = (x - 16581375) \cdot (x + 3375)^2$. Suppose that either of the vertices corresponding to the j -invariant -3375 in the graph $\mathcal{G}_2(\mathbb{F}_p)$ lies on the surface, and thus has endomorphism ring isomorphic to $\mathbb{Z} \left[\frac{1+\sqrt{-p}}{2} \right]$. The \mathbf{j} -invariants on the surface have three neighbours. Since there are no loops in $\mathcal{G}_2(\mathbb{F}_p)$, this vertex would have two neighbours with j -invariants -3375 , but there cannot be three vertices corresponding to $\mathbf{j} = -3375$. Hence $\text{End}_{\mathbb{F}_p}(E_{-3375}) \cong \mathbb{Z}[\sqrt{-p}]$.

Also note that the two self-isogenies of $\mathbf{j} = -3375$ are not defined over \mathbb{F}_p . \square

If $\mathbf{j} = 1728$ and $\mathbf{j} = -3375$ are both supersingular ($p \equiv 3 \pmod{4}$ and $p \equiv 3, 5, 6 \pmod{7}$), the proof also allows us to conclude that the endomorphism ring $\text{End}_{\mathbb{F}_p}(E)$ of any supersingular elliptic curve E with j -invariant $j(E) = 16581375$ is $\mathbb{Z} \left[\frac{1+\sqrt{-p}}{2} \right]$.

Example 3.8 (The j -invariant 1728 is both on the surface and on the floor.). *Suppose $p > 3$ with $p \equiv 3 \pmod{4}$. The isogeny*

$$\begin{aligned} \phi : E_{1728} : y^2 = x^3 - x &\rightarrow y^2 = x^3 + 4x =: E_{1728}^t \\ (x, y) &\mapsto \left(\frac{x^2 + x + 2}{x + 1}, \frac{x^2y + 2xy - y}{x^2 + 2x + 1} \right) \end{aligned}$$

is a vertical 2-isogeny with kernel $(0, 0)$ of non- \mathbb{F}_p -isomorphic supersingular elliptic curves with j -invariant 1728.

Note that E_{1728}^t is a quartic twist, not a quadratic twist, so Lemma 3.5 does not apply.

Proof. The isogeny ϕ was obtained by Velu's formulas. Factoring the right-hand side of the Weierstrass equation for E_{1728} , we see $E[2] \subset E(\mathbb{F}_p)$. By Lemma 3.4,

$$\text{End}_{\mathbb{F}_p}(E_{1728}) \cong \mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right]$$

Since $p \equiv 3 \pmod{4}$, the \mathbb{F}_p points of $E_{1728}^t[2]$ are precisely $\{\mathcal{O}_E, (0, 0)\}$. Again by Lemma 3.4,

$$\text{End}_{\mathbb{F}_p}(E_{1728}^t) \cong \mathbb{Z} [\sqrt{-p}],$$

so ϕ is a vertical isogeny. \square

Example 3.8 is the only vertical isogeny between two elliptic curves with same j -invariants.

Corollary 3.9. *Let v_a, w_a be the distinct vertices in $\mathcal{G}_\ell(\mathbb{F}_p)$ corresponding to the j -invariant $\mathbf{j} = a \in \mathbb{F}_p$ for $a \neq 1728$. Then, either v_a and w_a are either both on the floor or both on the surface of $\mathcal{G}_\ell(\mathbb{F}_p)$.*

Proof. The case of $\mathbf{j} = 0$ was handled in Example 3.6. For $\mathbf{j} \neq 0, 1728$, the two vertices $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ correspond to an elliptic curve and its quadratic twist. The result follows from Corollary 3.5. \square

Another proof of this statement can be found in the appendix of [Kan89] and is obtained by a careful examination of Hilbert polynomials of discriminant $-p$ and $-4p$, considered modulo p .

Kaneko actually proves that

$$\gcd(h_{-4p}(x), h_{-p}(x)) = x - 1728,$$

which translates to the statement that $j = 1728$ is the only j -invariant that can be both on the surface and the floor. Kaneko in turn gives credit to [Ibu82], who proved the statement (and more) in purely quaternionic terms.

Now, we describe the potential shapes of $\mathcal{G}_2(\mathbb{F}_p)$. The results are given in [DG16], however, we recall these potential shapes of $\mathcal{G}_2(\mathbb{F}_p)$ to compare with those of $\mathcal{S} \subset \mathcal{G}_2(\overline{\mathbb{F}_p})$.

3.1.2 The graph $\mathcal{G}_2(\mathbb{F}_p)$ in the case of $p \equiv 1 \pmod{4}$

For $p \equiv 1 \pmod{4}$, the ring $\mathbb{Z}[\sqrt{-p}]$ is the maximal order in $\mathbb{Q}(\sqrt{-p})$ and the prime 2 is ramified.

Lemma 3.10. *Suppose that $p > 7$. Then each connected component of $\mathcal{G}_2(\mathbb{F}_p)$ is a single edge and the edges correspond to horizontal isogenies.*

Proof. Since $p \equiv 1 \pmod{4}$, the ring $\mathbb{Z}[\sqrt{-p}]$ is the ring of integers in $\mathbb{Q}(\sqrt{-p})$ and hence, any supersingular elliptic curve over \mathbb{F}_p satisfies

$$\text{End}_{\mathbb{F}_p}(E) \cong \mathbb{Z}[\sqrt{-p}]$$

All of these edges are horizontal isogenies because all the curves satisfy $\text{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[\sqrt{-p}]$.

The proof of this is already in [DG16]. We present three proofs of the first statement.

1. Since every elliptic curve over \mathbb{F}_p has $p+1$ points and $p+1 \equiv 1+1 = 2 \pmod{12}$, we see that $\#E[2] = 2$, that is, there is exactly one point of order 2 defined over \mathbb{F}_p and hence exactly one 2-isogeny defined over \mathbb{F}_p .
2. Because the ring $\mathbb{Z}[\sqrt{-p}]$ is already the maximal order, Lemma 3.4, we get that $E[2] \not\subset E(\mathbb{F}_p)$ and so there can only be one outgoing 2-isogeny just like in the previous case.
3. Since (2) is ramified in $\mathcal{O}_K = \mathbb{Z}[\sqrt{-p}]$, it has order 2 in $\text{Cl}(\mathcal{O}_K)$ (this is since $p > 7$ and so there are no elements of norm 2 in $\mathbb{Z}[\sqrt{-p}]$). We know that the volcano is a cycle with the number of edges equal to the order of the prime above 2 in $\text{Cl}(\mathcal{O}_K)$, and hence we recover cycles of length $2 - 1 = 1$. \square

3.1.3 The graph $\mathcal{G}_2(\mathbb{F}_p)$ in the case of $p \equiv 3 \pmod{4}$

We will use the construction in the proof of Theorem 3.3 to describe the shape of the components of $\mathcal{G}_2(\mathbb{F}_p)$. Since $p \equiv 3 \pmod{4}$, we have two possible orders for endomorphism rings, and

$$\mathcal{O} = \mathbb{Z}[\sqrt{-p}] \subsetneq \mathcal{O}_K = \mathbb{Z}\left[\frac{1 + \sqrt{-p}}{2}\right]$$

is an inclusion of orders of index 2. To see how the prime above 2 acts on the points in $\mathcal{G}_2(\mathbb{F}_p)$ 3.1.1, consider the splitting behavior of $(2)\mathcal{O}_K$:

1. for $p \equiv 3 \pmod{8}$ the prime 2 is inert,
2. for $p \equiv 7 \pmod{8}$ the prime 2 splits into two prime ideals.

These two congruence conditions will result in different shapes of $\mathcal{G}_2(\mathbb{F}_p)$. We also consider $j = 1728$, as the extra automorphisms affect isogenies between $j = 1728$ and its neighbors.

Case 1. $p \equiv 3 \pmod{8}$

Let $K = \mathbb{Q}(\sqrt{-p})$. 2 is inert in K , so the prime \mathfrak{p} of \mathcal{O}_K above 2 has order 1 in $\text{Cl}(\mathcal{O}_K)$. From Theorem 3.3, any component of the $\mathcal{G}_2(\mathbb{F}_p)$ will be a volcano with surface of size 1 connected to the lower-levels as a ‘claw’: There will be three edges going out of any vertex on the surface. See Figure 3.2. The elliptic curves E on the surface have endomorphism ring $\mathbb{Z}\left[\frac{1 + \sqrt{-p}}{2}\right]$ and $E[2] \subset E(\mathbb{F}_p)$, so there are three outgoing 2-isogenies defined over \mathbb{F}_p .

The volcano stops at this depth, because there are only two possible endomorphism rings: $\mathbb{Z}[\sqrt{-p}]$ and $\mathbb{Z}\left[\frac{1 + \sqrt{-p}}{2}\right]$. Therefore, the volcanoes will be *claws* for $p \equiv 3 \pmod{8}$.

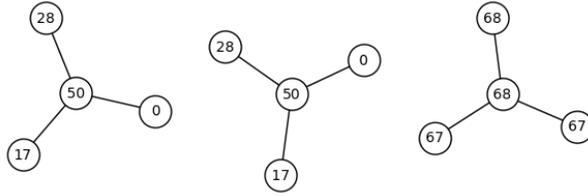


Figure 3.2: The graph $\mathcal{G}_2(\mathbb{F}_p)$ for $p = 83$: We clearly see the claw structure.

Case 2. $p \equiv 7 \pmod{8}$.

In this case, the ideal $(2)\mathcal{O}_K$ splits into two conjugate prime ideals. In general, they can have any order in the class group, but they are never principal. See Figure 3.1 for an example of this case.

Neighbours of $j = 1728$.

In Example 3.8, we saw that there is always a vertical isogeny from a vertex v_{1728} on the surface to a vertex w_{1728} on the floor of $\mathcal{G}_\ell(\mathbb{F}_p)$. Moreover, looking at the modular polynomial

$$\phi_2(1728, x) = (x - 1728) \cdot (x - 287496)^2,$$

we have the following:

1. For $p \equiv 3 \pmod{8}$, the vertices v_{287496} and w_{287496} corresponding to quadratic twists with j -invariant 287496 are on the floor, so

$$\text{End}_{\mathbb{F}_p}(E_{287496}) \cong \mathbb{Z}[\sqrt{-p}].$$

2. For $p \equiv 7 \pmod{8}$, the vertices v_{287496} and w_{287496} corresponding to quadratic twists with j -invariant 287496 are on the surface of the volcano, so

$$\text{End}_{\mathbb{F}_p}(E_{287496}) \cong \mathbb{Z} \left[\frac{1 + \sqrt{-p}}{2} \right].$$

3.2 Passing from the graph $\mathcal{G}_\ell(\mathbb{F}_p)$ to the spine $\mathcal{S} \subset \mathcal{G}_\ell(\overline{\mathbb{F}_p})$

The subgraph \mathcal{S} of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ can be obtained from the graph $\mathcal{G}_\ell(\mathbb{F}_p)$ in the following two steps:

1. Identify the vertices with the same j -invariant: these two vertices of $\mathcal{G}_\ell(\mathbb{F}_p)$ merge to a single vertex on $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Identify equivalent edges.
2. Add the edges from $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ between vertices in \mathbb{F}_p corresponding to isogenies which are defined over $\overline{\mathbb{F}_p} \setminus \mathbb{F}_p$.

One notation we use to distinguish between vertices of $\mathcal{G}_\ell(\mathbb{F}_p)$ and those of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$: Vertices of components of $\mathcal{G}_\ell(\mathbb{F}_p)$ corresponding to the j -invariant a will be denoted v_a and w_a , where v_a is a vertex on the connected component V of $\mathcal{G}_\ell(\mathbb{F}_p)$ and w_a lies on the component W (not necessarily distinct from V). Since the j -invariants uniquely determine the vertices of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$, we will use a to denote a vertex of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. It is useful to think of the vertices v_a, w_a as elliptic curves that are twists of each other.

Remember that $\mathcal{G}_\ell(\mathbb{F}_p)$ is not a subgraph of \mathcal{S} since both the vertices and edges of $\mathcal{G}_\ell(\mathbb{F}_p)$ may be merged in \mathcal{S} . Fortunately, something weaker is true, as we show in the following lemma. It turns out distinct edges from the same vertex in $\mathcal{G}_\ell(\mathbb{F}_p)$ correspond to distinct edges in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.

Lemma 3.11 (\mathbb{F}_p -edges are rigid.). *Let E be an elliptic curve with $j(E) \notin \{0, 1728\}$ defined over \mathbb{F}_p (with $p \geq 5$). Suppose that there are two ℓ -isogenies from E defined over \mathbb{F}_p . Then they are equivalent over $\overline{\mathbb{F}_p}$ if and only if they are equivalent over $\overline{\mathbb{F}_p}$.*

Proof. If the isogenies are equivalent over \mathbb{F}_q , then they are equivalent over $\overline{\mathbb{F}_p}$.

Let $\phi_1 : E \rightarrow E_1$ and $\phi_2 : E \rightarrow E_2$ be two isogenies that are defined over $\overline{\mathbb{F}_p}$. We want to show that they are equivalent over \mathbb{F}_p . By hypothesis, there exists $(\overline{\mathbb{F}_p})$ -isomorphisms $\varphi : E \rightarrow E$ and $\psi : E_1 \rightarrow E_2$ such that $\phi_2 = \psi \circ \phi_1 \circ \varphi$.

Consider the commuting square

$$\begin{array}{ccc} E & \xrightarrow{\phi_1} & E' \\ \downarrow \varphi_1 & & \downarrow \varphi_2 \\ E & \xrightarrow{\phi_2} & E' \end{array}$$

We know that the kernel of the map $\varphi_2 \circ \phi_1$ is $\ker \phi_1$. Therefore, the kernel of the map $\phi_2 \circ \varphi_1$ also is $\ker \phi_1$. This means that

$$\varphi_1(\ker(\phi_1)) = \ker(\phi_2).$$

By the hypothesis on $j(E)$, $\text{Aut}_{\mathbb{F}_p}(E) = \text{Aut}_{\overline{\mathbb{F}_p}}(E) = \{\pm 1\}$, so this is not possible as $[\pm 1] \ker \phi_1 = \ker \phi_1$. \square

We note that the proof above works if we replace q with p^n and consider isogenies and curves defined over \mathbb{F}_q , however, this will not be needed in our discussion.

Lemma 3.11 for $\ell = 2$ gives the following corollary.

Corollary 3.12. *If the neighbors of an elliptic curve E with j -invariant \mathbf{j} in $\mathcal{G}_2(\mathbb{F}_p)$ are elliptic curves with j -invariants $\mathbf{j}_1, \mathbf{j}_2$ and \mathbf{j}_3 , then the neighbors of \mathbf{j} in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ are $\mathbf{j}_1, \mathbf{j}_2$ and \mathbf{j}_3 .*

Since there are always at most 2 neighbours of any vertex in $\mathcal{G}_\ell(\mathbb{F}_p)$ for $\ell > 2$, the above corollary does not generalize. However, it is still true that if there are neighbours of v_a, w_a (recall that v_a, w_a are the two vertices in $\mathcal{G}_\ell(\mathbb{F}_p)$ that have j -invariant a) that have j -invariants b, c, d, e , then there are (not necessarily distinct) edges $[a, b], [a, c], [a, d]$ and $[a, e]$ in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.

Defined below are the four processes that can happen to the components of $\mathcal{G}_\ell(\mathbb{F}_p)$ when passing to \mathcal{S} . We will show that this list is exhaustive.

Definition 3.13 (Stacking, folding and attaching).

1. Let V and W be two distinct components of $\mathcal{G}_\ell(\mathbb{F}_p)$. We say that V and W **stack** if, when we relabel the vertices v_a by the j -invariant a , they become isomorphic as graphs.
2. A connected component V of $\mathcal{G}_\ell(\mathbb{F}_p)$ **folds** in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ if V contains vertices corresponding to both quadratic twists of every j -invariant on V . The term is meant to invoke what happens to this component when the quadratic twists are identified in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.
3. Two connected components V and W of $\mathcal{G}_\ell(\mathbb{F}_p)$ become **attached by a new edge** in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ if there is a new edge $[a, b] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$ corresponding to an isogeny between vertices $v_a \in V$ and $w_b \in W$ that is not defined over \mathbb{F}_p .
4. We say that two components V and W of $\mathcal{G}_\ell(\mathbb{F}_p)$ for $\ell > 2$ **attach along the j -invariant a** if they both contain a vertex $v_a \in V, w_a \in W$ that corresponds to j -invariant a and such that there is a neighbour v_b of v_a with j -invariant b such that the twist of w_b is not a neighbour of w_a and vice versa.

An example of the first three of the four phenomena are given by Figure 3.3 and an example of attachment along a j -invariant can be seen in Figure 3.4.

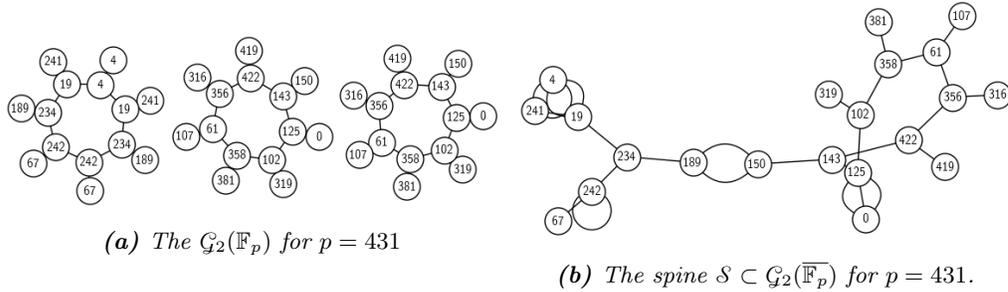


Figure 3.3: Stacking, folding and attaching by an edge for $\mathfrak{p} = 431$ and $\ell = 2$. The leftmost component of $\mathcal{G}_2(\mathbb{F}_p)$ folds, the other two components stack, and the vertices 189 and 150 get attached by a double edge.

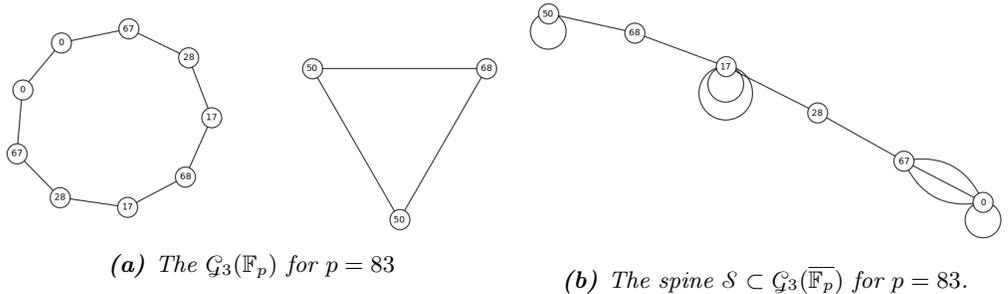


Figure 3.4: Attachment along a j -invariant for $p = 83$ and $\ell = 3$. We see that the two connected components of $\mathcal{G}_3(\mathbb{F}_p)$ are attached along the j -invariant $68 = 1728 \pmod{83}$.

Note that it can happen that an attachment is actually attaching the component V to itself. For instance, whenever there is only one component, new edges cannot attach distinct components. See Figure 3.6.

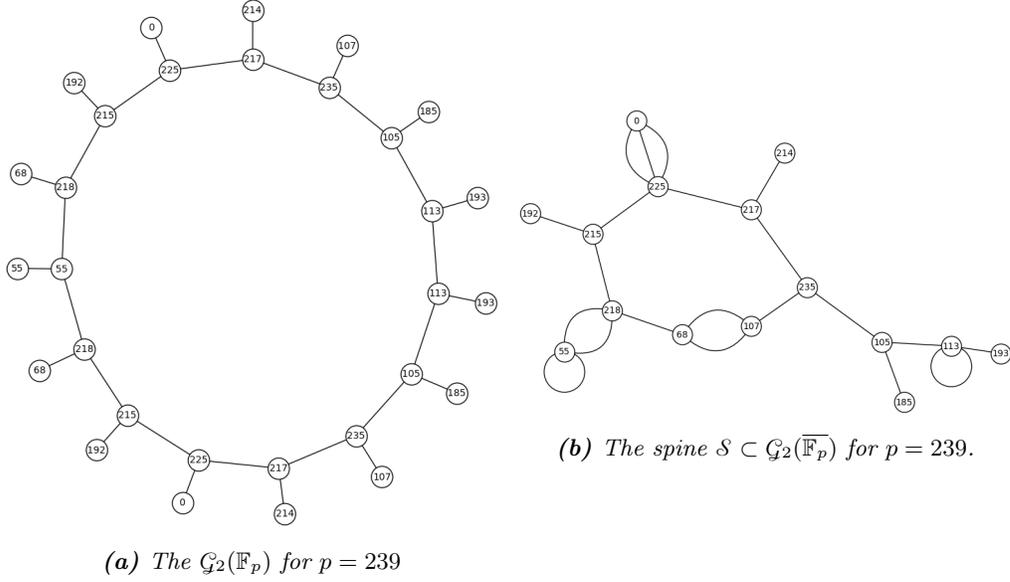


Figure 3.5: Attachment by an edge that does not attach two distinct components. The vertices with j -invariants 68 and 107 are joined by a double edge.

We now begin analyzing the new edges in \mathcal{S} that do not come from edges in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. For any elliptic curve E and ℓ prime, ℓ -isogenies are given by a cyclic subgroup of order ℓ of the ℓ -torsion points $E[\ell]$. Such a subgroup is generated by a point of exact order ℓ . An ℓ -isogeny is defined over \mathbb{F}_p if and only if its kernel is defined over \mathbb{F}_p (that is, the kernel is fixed by the Frobenius morphism).

For 2-isogenies, the kernel consists of the point at infinity, O_E and a point $P \in E[2] \setminus \{O_E\}$. The 2-isogeny with kernel $\langle P \rangle$ is defined over \mathbb{F}_p if and only if P is defined over \mathbb{F}_p .

For $\ell > 2$, the point P generating $\ker \phi$ does not have to be defined over \mathbb{F}_p , only the whole kernel needs to be fixed by the Frobenius morphism.

Remark 3.14. Let E, E' be elliptic curves with j -invariants j, j' and suppose that there is an edge $[j, j']$ in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Then there is an ℓ -isogeny $\phi : E \rightarrow E'$. Even if both j, j' are in \mathbb{F}_p , the isogeny ϕ is not necessarily defined over \mathbb{F}_p . We can see this in Figure 3.3: the 2-isogenies between 150 and 189 in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ are not defined over \mathbb{F}_p . Also, the vertex v_4 ($4 \equiv 1728 \pmod{431}$) on the floor of $\mathcal{G}_2(\mathbb{F}_p)$ has no edge to a curve with j -invariant 19, but there is an edge $[4, 19] \in \mathcal{G}_2(\overline{\mathbb{F}_p})$ coming from the two isogenies from the vertex v_4 on the surface. Moreover, Lemma 3.11 gives us that there is a double edge $[4, 19] \in \mathcal{G}_2(\mathbb{F}_p)$. This not a coincidence, as we will explain Lemma 3.15.

Lemma 3.15 (One new isogeny implies two new isogenies). Let $v_a, v_b \in \mathcal{G}_\ell(\mathbb{F}_p)$ correspond to \mathbb{F}_p -elliptic curves E_a and E_b with j -invariants a, b respectively with $a \neq 1728, 0$. Assume that there is no edge $[v_a, v_b] \in \mathcal{G}_\ell(\mathbb{F}_p)$, but there is an edge $[a, b] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Then there are two isogenies defined between $E_a \rightarrow E_b$ which are inequivalent over $\overline{\mathbb{F}_p}$ and hence a double edge $[a, b] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$.

Proof. We know that there is an ℓ -isogeny $\phi : E_a \rightarrow E$ to some elliptic curve E with j -invariant b . Since $j(E) = b$, then E is isomorphic to E_b over \mathbb{F}_{p^2} . Composing with this isomorphism, we obtain an ℓ -isogeny $\psi : E_a \rightarrow E_b$. However, ψ cannot be defined over \mathbb{F}_p , since we assumed there was no edge $[v_a, v_b] \in \mathcal{G}_\ell(\mathbb{F}_p)$.

The kernel of ψ is not defined over \mathbb{F}_p (otherwise ψ would be defined over \mathbb{F}_p), so the p -power Frobenius map $\text{Frob} : \mathbb{F}_p \rightarrow \mathbb{F}_p$ does not preserve $\ker \psi$. There is an isogeny from E_a with kernel

ψ^{Frob} . This isogeny has degree ℓ since ψ^{Frob} has order ℓ and it is not equivalent to ψ . Using the construction of isogenies from Vélú's formulae, we obtain the rational maps for defining ψ^{Frob} . In particular, the j -invariant of the target of ψ^{Frob} is necessarily $\text{Frob}(b) = b$. Hence, there are two inequivalent isogenies between E_a and E_b and hence two edges $[a, b] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$.

Note that we cannot simply compose with Frobenius, because that would give us an inseparable isogeny with degree $\ell \cdot p$. \square

The corollary below explains why for both attachment by a new edge (cf. Figure 3.3 and Figure 3.6) and attachment along a j -invariant (cf. Figure 3.4), we always see double edges.

Corollary 3.16. *Attachment of components from v_a to w_b forces a double edge $[a, b]$ in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Attachment along the j -invariant \mathbf{j} implies a double edge from \mathbf{j} in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.*

Proof. In the first case, we are adding an edge between v_a and w_b that is not defined over \mathbb{F}_p and we can directly apply Lemma 3.15.

In the second case, we assume that there is a neighbour v_b of v_j such that w_b is not a neighbour of w_j . Applying the Lemma 3.15 to the $\overline{\mathbb{F}_p}$ isogeny from w_j to v_b , we obtain a double edge $[\mathbf{j}, b] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$. \square

The next step in understanding $\mathcal{S} \subset \mathcal{G}_\ell(\overline{\mathbb{F}_p})$ is understanding the neighbours of the two vertices v_a, w_a corresponding to the same j -invariant. This is done in Lemma 3.17 for $\ell > 2$ and in Lemma 3.27 for $\ell = 2$. The case $\ell = 2$ is more involved because in this case, there exist vertical isogenies (if $p \equiv 3 \pmod{4}$), whereas for $\ell > 2$, all isogenies are horizontal.

3.3 Stacking, folding and attaching for $\ell > 2$

In this section, we consider the spine \mathcal{S} for $\ell > 2$. In this case, there are no vertical isogenies, hence the graph $\mathcal{G}_\ell(\mathbb{F}_p)$ is a union of disjoint cycles: The cycles of vertices corresponding to curves either only with endomorphism ring $\mathbb{Z}\left[\frac{1+\sqrt{-p}}{2}\right]$, or only with endomorphism ring $\mathbb{Z}[\sqrt{-p}]$.

We will avoid on the case when the graph $\mathcal{G}_\ell(\mathbb{F}_p)$ is just a disjoint union of vertices (i.e., when there are no isogenies defined over \mathbb{F}_p). It suffices to assume that $p \equiv -1 \pmod{\ell}$ (when $\ell \nmid \#E(\mathbb{F}_p) = p + 1$, there are \mathbb{F}_p -rational points of order ℓ).

Lemma 3.17 (The neighbour lemma). *Suppose that $\ell > 2$. Suppose that v_a, w_a are the two vertices in $\mathcal{G}_\ell(\mathbb{F}_p)$ corresponding to elliptic curves with j -invariant a and such that the neighbours of v_a have j -invariants b, c and the neighbours of w_a have j -invariants c, d .*

Then either $\{b, c\} = \{c, d\}$ with $b \neq c$ or there is a double edge from a in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.

Proof. Suppose that $d \neq b, c$. Since there is an edge in $[a, d] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$ corresponding to the edge $[w_a, w_d] \in \mathcal{G}_\ell(\mathbb{F}_p)$, there is an isogeny from the elliptic curve v_a to an elliptic curve with j -invariant d . This isogeny cannot be defined over \mathbb{F}_p since the neighbours of v_a in $\mathcal{G}_\ell(\mathbb{F}_p)$ have j -invariants b, c . This gives at least two edges $[a, d] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$, by Lemma 3.15. \square

Corollary 3.18. *An attachment along a j -invariant a implies a double edge from a in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$.*

Proof. See Definition 4: At least one neighbour of v_a is distinct from the neighbours of w_a . \square

The main result of this section is the following result.

Proposition 3.19 (Stacking, folding and attaching for $\ell > 2$). *While passing from $\mathcal{G}_\ell(\mathbb{F}_p)$ to \mathcal{S} , the only possible events are stacking, folding, and n attachments by a new edge and m attachments along a j -invariant with $m + 2n \leq 2\ell(2\ell - 1)$.*

Proof. Suppose that v_a is a vertex of $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ such that a does not admit a double edge in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Then the neighbours of v_a and w_a (its twist) are the same by 3.17. The connected components of v_a and w_a look the same locally at a .

Suppose further that the connected component $V \subset \mathcal{G}_\ell(\mathbb{F}_p)$ does not contain any vertex that admits a double edge. By Lemma 3.17, every vertex v_a has the same neighbours as its twist w_a , so the component either folds (if $w_a \in V$) or stacks with the component W of w_a , which is necessarily identical to V when we replace the labels of the vertices by their j -invariants.

By Definition 3, attachment happens when we add an edge that cannot be defined over \mathbb{F}_p . By 3.15, attachments necessarily imply double edge. Attachment along a j -invariant a also implies that there is a double edge from a . However, double edges can only occur at j -invariants which are roots of

$$\text{Res} \left(\Phi_\ell(X, Y), \frac{d}{dY} \Phi_\ell(X, Y); Y \right),$$

which is a polynomial of degree bounded by $2\ell(2\ell - 1)$ by Lemma 2.14. Therefore, except at vertices corresponding to j -invariants that admit double edges, the components will either stack or fold. Even in components containing vertices that admit double edges, all other pairs of vertices corresponding to the same j -invariant will either stack onto each other or, if they share a neighbour, fold onto each other. See 3.4.

Finally, for attachment by an edge $[a, b] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$, both endpoints admit a double edge in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$, hence both a and b are roots of $\text{Res}_\ell(X)$. Since the degree of $\text{Res}_\ell(X)$ is bounded by $2\ell \cdot (2\ell - 1)$, we obtain the bound. \square

For any given ℓ , we know the possible attachments: $\text{Res}_\ell(X)$ is a product of Hilbert class polynomials, so having roots in \mathbb{F}_p which give supersingular j -invariants is equivalent to satisfying certain congruence conditions on p . We can construct primes p to avoid attachments.

Typically, the polynomial $\text{Res}_\ell(X)$ will be smooth and have lots of repeated factors, so for any given choice of ℓ , the bound in Proposition 3.19 can be made more precise, which we will show in the following section.

3.3.1 Example: stacking, folding and attaching for $\ell = 3$

In this section, we study the stacking, folding and attaching behaviour for $\ell = 3$. The case $\ell = 2$ will be discussed in Section 3.4. The case $\ell = 3$ is cryptographically relevant, because of the use of $\mathcal{G}_3(\mathbb{F}_p)$ in SIDH and SIKE. Moreover, keeping ℓ small, we can give better bounds on the number of attachments and explain the results of the previous section in a more hands-on manner.

We start with factoring over \mathbb{Z} the polynomial $\text{Res}_3(x)$ introduced in (3):

$$\begin{aligned} \text{Res}_3(x) = & (-1) \cdot 3^3 \cdot x^2 \cdot (x - 8000)^2 \cdot (x - 1728)^2 \cdot (x + 32768)^2 \cdot (x^2 - 52250000x + 12167000000)^2 \\ & \cdot (x^2 - 1264000x - 681472000)^2 \cdot (x^2 + 117964800x - 134217728000)^2. \end{aligned}$$

The irreducible factors are Hilbert class polynomials of discriminants $-3, -8, -4, -11, -32, -20$ and -35 , respectively. Removing the repeated factors, we see that there are at most 10 vertices at which a double edge can occur.

Double-edges also arise in loops (double self-3-isogenies), which accounts for some of the factors of $\text{Res}_3(x)$. We find the self loops by factoring the modular 3-isogeny polynomial:

$$\Phi_3(x, x) = (-1) \cdot x \cdot (x - 54000) \cdot (x - 8000)^2 \cdot (x + 32768)^2.$$

At $\mathbf{j} = 8000$ and $\mathbf{j} = -32768$, there are two self-3-isogenies and no attachment at these vertices.

Example 3.20 (Neighbours of the vertices with loops). *In this example, we determine the $\mathcal{G}_3(\mathbb{F}_p)$ neighbours of $\mathbf{j} = 0, 8000, 54000$ and -32768 . This is done by factoring $\Phi_3(\mathbf{j}, x)$:*

1. $\mathbf{j} = 0$: $\Phi_3(0, x) = (x + 12288000)^3 \cdot x$. From this we conclude, that there is an isogeny v_0, w_0 that is defined over \mathbb{F}_p (the factor x has multiplicity 1, indicating this is not a double-edge and thus cannot appear only over \mathbb{F}_{p^2}). Hence, the neighbours of v_0 are w_0 and $v_{-12288000}$ and the neighbours of w_0 are $w_{-12288000}$ and v_0 . Moreover, the edges to -12288000 that are not defined over \mathbb{F}_p will not be attaching edges.

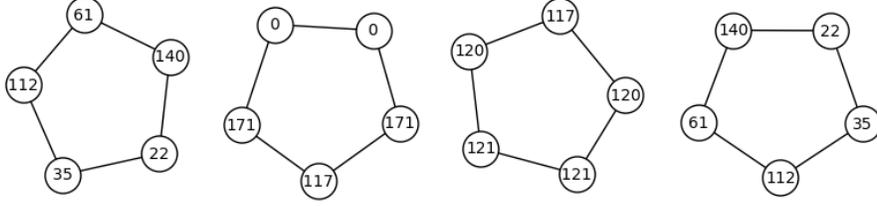


Figure 3.6: The graph $\mathcal{G}_3(\mathbb{F}_p)$ for $p = 179$. We see that the neighbours of vertices with j -invariant 0 both have j -invariant $-12288000 \bmod 179 = 171$.

As an aside, we note that since all isogenies in $\mathcal{G}_\ell(\mathbb{F}_p)$ for $\ell > 2$ are horizontal and $\text{End}_{\mathbb{F}_p}(E_0) \cong \mathbb{Z}[\sqrt{-p}]$, it follows that $\text{End}_{\mathbb{F}_p} E_{-12288000} \cong \mathbb{Z}[\sqrt{-p}]$.

2. $\mathbf{j} = 54000$: There is one self-3-isogeny which arises from a 3-isogeny ψ between (non-isomorphic) quadratic twists with $j = 54000$. Explicitly, let $E_{54000} : y^2 = x^3 - 15x + 22$, $E'_{54000} : y^2 = x^3 - 135x - 594$. ψ is given:

$$\psi = \left(\frac{x^3 - 6x^2 + 33x - 56}{x^2 - 6x + 9}, \frac{x^3y - 9x^2y + 3xy + 13y}{x^3 - 9x^2 + 27x - 27} \right)$$

that reduces modulo any prime p to a 3-isogeny over \mathbb{F}_p .

Moreover, $\phi_3(54000, x)$ factors as

$$(x - 54000) \cdot (x^3 - 151013228706000x^2 + 224179462188000000x - 187999470568800000000),$$

so (for p large enough) the j -invariant 54000 cannot admit a double edge.

3. $\mathbf{j} = 8000$: factor $\Phi_3(8000, x) = (x^2 - 377674768000x + 232381513792000000)(x - 8000)^2$. There is a double loop $[8000, 8000] \in \mathcal{G}_3(\overline{\mathbb{F}_p})$. Neither of the loops occur over \mathbb{F}_p : both cannot occur over \mathbb{F}_p because there are no double edges in $\mathcal{G}_3(\mathbb{F}_p)$. If only one of them came from an isogeny over \mathbb{F}_p , we could use Lemma 3.15 to get a third loop, which is not possible (for large p).
4. $\mathbf{j} = -32768$: $\Phi_3(-32768, x) = (x^2 + 37616060956672 \cdot x - 56171326053810176) \cdot (x + 32768)^2$. Repeating the argument we gave above for $\mathbf{j} = 8000$, the self loops cannot come from isogenies over \mathbb{F}_p .

To conclude: For $\mathbf{j} = 0, 54000$, the self-3-isogeny comes from an isogeny between the twists in $\mathcal{G}_3(\mathbb{F}_p)$, and for $\mathbf{j} = 8000, -32768$, the double self-3-isogenies are not defined over \mathbb{F}_p .

The following lemma shows that we can distinguish attachment along a j -invariant and an attachment by a new edge looking at the neighbours of the given j -invariants.

Lemma 3.21 (Attaching for $\ell = 3$). 1. Let a be an attaching j -invariant. Then the neighbours of v_a have the same j -invariant b and induce a double edge $[a, b] \in \mathcal{G}_3(\overline{\mathbb{F}_p})$ and the neighbours of w_a have the same j -invariant c and induce a double edge $[a, c] \in \mathcal{G}_3(\overline{\mathbb{F}_p})$, with $b \neq c$.

2. Let $[a, b]$ be an attaching edge in $\mathcal{G}_3(\overline{\mathbb{F}_p})$. Suppose that the neighbours of v_a have j -invariants c, d . Then the neighbours of w_a have j -invariants c, d . Necessarily $c \neq d$.

Proof. 1. This follows from Lemma 3.15: Suppose the neighbouring vertices w_b, w_c of w_a have j -invariants b, c . Suppose that v_c , the twist of w_c , is not a neighbour of v_a . Lemma 3.15 applied to the pair v_a, w_c gives a double edge $[a, c] \in \mathcal{G}_3(\mathbb{F}_p)$.

Similarly, there is a neighbour v_d of v_a with j -invariant d such that w_d is not a neighbour of w_a and we obtain a double edge $[a, d]$ in $\mathcal{G}_3(\mathbb{F}_p)$. Since there are only 4 edges from a

in $\mathcal{G}_3(\mathbb{F}_p)$ and since we assumed that at least one of the neighbours of v_a had a different j -invariant than the neighbours of w_a (and vice versa), we necessarily have that both v_a and w_a have two neighbours with the same j -invariant.

2. Suppose that there is a new (double) edge $[a, b]$, not coming from an edge in $\mathcal{G}_3(\mathbb{F}_p)$. Let v_a and w_a be the twists corresponding to $\mathbf{j} = a$. Let v_c, v_d be the neighbours of v_a . The edges from a in $\mathcal{G}_3(\mathbb{F}_p)$ are $[a, b], [a, b], [a, c]$ and $[a, d]$. Since we assumed that the new edge does not come from the $\mathcal{G}_3(\mathbb{F}_p)$, the neighbours of w_a cannot have j -invariant b and are necessarily w_c, w_d . \square

Corollary 3.22 (Neighbours of twists). *For every $a \in \mathbb{F}_p$ that is not a root of $\text{Res}_3(x)$, the neighbours of v_a and its twist w_a have the same j -invariants b, c with $b \neq c$.*

Proof. If at least one of the neighbours of w_a had a different j -invariant than the neighbours of v_a , it would be an attaching j -invariant (every vertex in $\mathcal{G}_3(\mathbb{F}_p)$ has only two neighbours). The result follows from Lemma 3.21. \square

Example 3.23. *Let us work out the above lemmas for $p = 71$. $\mathcal{G}_3(\mathbb{F}_p)$ is given in Figure 3.7. The supersingular j -invariants are $0, 17, 24, 40(\equiv 54000 \pmod{71}), 41, 48(\equiv 8000 \pmod{71}), 66$.*

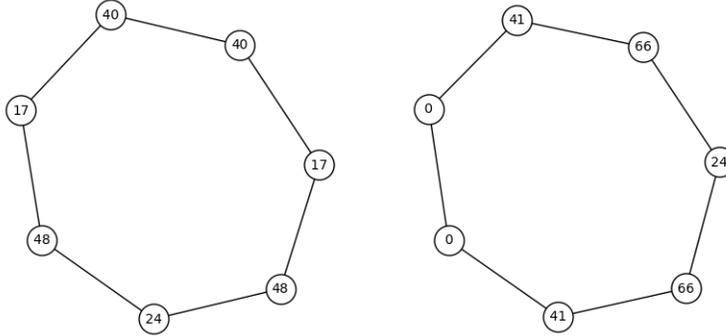


Figure 3.7: The $\mathcal{G}_3(\mathbb{F}_p)$ for $p = 71$

The polynomial $\text{Res}_3(x)$ factors in \mathbb{F}_p :

$$\begin{aligned} \text{Res}_3(x) = & (44) \cdot x^2 \cdot (x - 66)^2 \cdot (x - 48)^2 \cdot (x - 42)^2 \cdot (x - 41)^2 \\ & \cdot (x - 40)^2 \cdot (x - 34)^2 \cdot (x - 25)^2 \cdot (x - 24)^2 \cdot (x - 17)^2. \end{aligned}$$

All the vertices of \mathcal{S} are roots of $\text{Res}_3(x)$, admitting a double edge. The ones that correspond to a double self-loop are the roots of

$$\frac{\partial \Phi_3(x, x)}{\partial x} = (65) \cdot (x - 48) \cdot (x - 34) \cdot (x^3 + 54x^2 + 54x + 54).$$

Namely, there is a double self-loop at 48 (because 34 is not a supersingular j -invariant for $p = 71$). Finally, there are single loops at 0 and 40, as these are zeroes of $\Phi_3(x, x)$ with multiplicities 1.

For the neighbours of 0 and $24 \equiv 1728 \pmod{71}$, we count the edges from these special vertices as one. Moreover, we see that only the edges $[40, 66]$ and $[17, 41]$ are the attaching edges.

The following theorem is a specialization of Proposition 3.19. We are mainly interested in the cryptographic applications, so we restrict to the case $p \equiv 3 \pmod{4}$. Then the class numbers $h(-p)$ and $h(-4p) = 3 \cdot h(-p)$ are both odd. With our assumption $p \equiv 2 \pmod{3}$, we have $p \equiv 11 \pmod{12}$.

Theorem 3.24 (Stacking, folding and attaching for $\ell = 3$). *Let p be a prime, $p \equiv 11 \pmod{12}$. When passing from $\mathcal{G}_3(\mathbb{F}_p)$ to the spine $\mathcal{S} \subset \mathcal{G}_3(\mathbb{F}_p)$,*

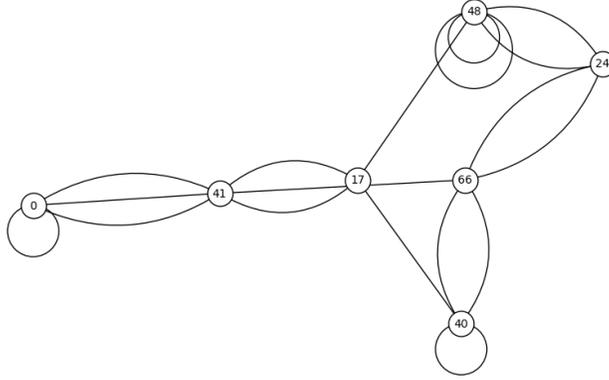


Figure 3.8: The $\mathcal{S} \subset \mathcal{G}_3(\mathbb{F}_p)$ for $p = 71$

1. all components that do not contain 0 or 54000 stack,
2. there are two distinct connected components V and W that contain a j -invariant 1728, one of them contains both vertices with j -invariant 0 and the other one both vertices with j -invariant 54000. V and W fold and get attached at the j -invariant 1728.
3. At most 8 vertices admit new edges, attaching at most 4 pairs of components by a new edge.

Proof. In Lemmas 3.21 and 3.17, we showed that j -invariants that attach by a new edge and j -invariants that do not admit a double edge look the same in the graph $\mathcal{G}_3(\mathbb{F}_p)$: that is, if the vertex v_a has neighbours v_b, v_c , then the vertex w_a has neighbours w_b, w_c . We do not need to treat these vertices with a separate case.

Suppose that there is a component V that does not stack. This either means that there is a vertex v_a whose neighbours are different than those of w_a , in this case a is an attaching j -invariant and we will treat this case below.

Or, there is a j -invariant a such that both the vertices v_a, w_a are in the component V . We know that V is a cycle. The vertices v_a, w_a divide the cycle in two halves, choose either half H . Look at the neighbours of v_a and w_a . If they have the same j -invariant b , replace a with b and continue moving along the halves of the cycle, until either of the following happens:

- (i) v_a and w_a are neighbours in $\mathcal{G}_3(\mathbb{F}_p)$ and hence induce a loop in $\mathcal{G}_\ell(\mathbb{F}_p)$.
- (ii) The only neighbour of v_a and w_a is a vertex v_j with j -invariant \mathbf{j} . This is necessarily an attaching j -invariant because the neighbours of w_j cannot have j -invariants a .
- (iii) The neighbour v_b of v_a has j -invariant b and the neighbour w_c of w_a has j -invariant c , for $b \neq c$. Then either the other neighbour of w_a is w_b or a is an attaching j -invariant. Suppose a is not an attaching j -invariant. Continuing along the whole cycle V in the direction of the edge $[v_a, v_b]$, and symmetrically in the direction of $[w_a, w_b]$, we will reach a point when the neighbour of some v_c is not a neighbour of the w_c . This happens when class number is odd because we cannot get the same sequence in the half that has odd length and in the half that has even length. Here we also obtain an attaching j -invariant.

We now discuss what happens with the components that contain an attaching j -invariant. The proof is a similar argument to the one above. Starting at any attaching j -invariant $\mathbf{j} \in V$ (there could be multiple), we know that its neighbours v_a, w_a have the same j -invariant by 3.21. By walking away from \mathbf{j} , we will at some point reach either

- (i) a pair of vertices v_b, w_b that are connected and induce a loop in $\mathcal{G}_3(\overline{\mathbb{F}_p})$. The component then folds.

- (ii) A pair of vertices v_b, w_b such that the neighbour in the direction away from \mathbf{j} of v_b is v_c and of w_b is w_d for some $c \neq d$. But then b is an attaching j -invariant and hence the neighbours of v_b and w_b have different j -invariants by Lemma 3.21. But we assumed that they come from the chain from \mathbf{j} and so the neighbours of v_b and w_b in the direction of \mathbf{j} have the same j -invariant. This is a contradiction.
- (iii) A single vertex v_b from ‘both sides’. But since the class number is odd, this gives us a contradiction.

The above then shows that any component V of $\mathcal{G}_3(\mathbb{F}_p)$ that contains an attaching j -invariant \mathbf{j} contains precisely one attaching j -invariant, folds and the ‘opposite vertices’ (the vertices v_b, w_b that are the furthest away from \mathbf{j}) are connected by an \mathbb{F}_p -isogeny, hence inducing a loop in $\mathcal{G}_3(\mathbb{F}_p)$.

Example 3.20 showed that the only possible opposite vertices are j -invariants 0 and 54000. For $p \equiv 3 \pmod{4}$, there are two components containing 1728: By Section 3.1.3, one of the vertices corresponding to 1728 is on the floor and the other one is on the surface, so they are on different components of $\mathcal{G}_3(\mathbb{F}_p)$. One of these vertices is on the same component of $\mathcal{G}_3(\mathbb{F}_p)$ as the vertices with j -invariant 0 and the other one will contain both vertices with j -invariant 54000. \square

Remark 3.25. 1. *The proof above shows that*

$$\text{End}_{\mathbb{F}_p}(E_{54000}) = \mathbb{Z}[\sqrt{-p}]$$

whenever this j -invariant is supersingular.

- 2. *The proof above holds for any p such that the order of the prime above 3 in $\text{Cl}(\mathcal{O}_K)$ is odd, which is necessarily the case for $p \equiv 3 \pmod{4}$.*
- 3. *It is possible to extend the proof for primes p such that the order of the prime above 3 in $\text{Cl}(\mathcal{O}_K)$ is even, however, one needs to consider the case of cyclic graphs like Figure 3.9 and correspond to case 3 in the proof of Theorem 3.24.*

It should be possible to argue that the two distinct paths from v_a to w_a cannot collapse onto one loop if one adapts the proof of Lemma 3.11 because a composition of cyclic isogenies with no backtracking will again be a cyclic isogeny.

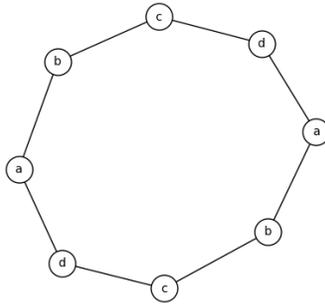


Figure 3.9: *In case iii for primes p such that the prime above 3 has even order, one needs to disprove the situation depicted above.*

3.4 Stacking, folding and attaching for $\ell = 2$

We identify how the components of $\mathcal{G}_2(\mathbb{F}_p)$ come together in $\mathcal{S} \subset \mathcal{G}_2(\overline{\mathbb{F}_p})$. Vertical 2-isogenies are possible, in contrast to the $\ell = 3$ case from Section 3.3.1.

The main theorem of this section is the following:

Theorem 3.26 (Stacking, folding and attaching). *Only stacking, folding or at most 1 attachment by a new edge are possible. In particular, no attachments by a j -invariant are possible.*

Recall that we form the graph $\mathcal{S} \subset \mathcal{G}_2(\overline{\mathbb{F}_p})$ from $\mathcal{G}_2(\mathbb{F}_p)$ in two steps: identify vertices corresponding to the same j -invariant and identify the edges, then add new edges.

We will show that:

1. the neighbours of the two vertices that correspond to twists have the same j -invariants (Proposition 3.27) and this will imply that only stacking and folding is possible.
2. At most one component folds, and for $p \equiv 3 \pmod{4}$ this is the component containing $\mathbf{j} = 1728$ (Proposition 3.31).
3. Attaching of components by a new edge happens between at most one pair of vertices, and those vertices are roots of the Hilbert class polynomial of $\mathbb{Q}(\sqrt{-15})$ (Proposition 3.30).

We begin with some results on the neighbours of vertices corresponding to the same j -invariants. From Corollary 3.5, we know that (except for 1728), twists have isomorphic endomorphism rings and hence lie on the same level in the volcano. More is true:

Proposition 3.27. *Let \mathbf{j} be a supersingular j -invariant and let $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ be two distinct vertices in $\mathcal{G}_2(\mathbb{F}_p)$ corresponding to elliptic curves with j -invariant \mathbf{j} . If $j \neq 1728$, then the two vertices corresponding to the same j -invariants have the same neighbours, that is:*

1. if $p \equiv 1 \pmod{4}$ and the neighbour of $v_{\mathbf{j}}$ is $v_{\mathbf{j}'}$ then the neighbour of $w_{\mathbf{j}}$ is $w_{\mathbf{j}'}$,
2. if $p \equiv 3 \pmod{4}$ and if
 - (a) the vertices $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ are both on the floor, and $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ are each attached to a vertex with j -invariant \mathbf{j}' ,
 - (b) the vertices $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ are both on the surface. $v_{\mathbf{j}}$ has three neighbours with distinct j -invariants a, b, c and $w_{\mathbf{j}}$ has three neighbours with the same distinct j -invariants a, b, c .

The neighbours of $j = 1728$ are given in Section 3.1.3.

Proof of Proposition 3.27. 1. For $p \equiv 1 \pmod{4}$, any connected component of $\mathcal{G}_2(\mathbb{F}_p)$ is an edge. If $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ are on the same connected component, the result follows immediately.

If $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ are on different connected components, denote the neighbour of $v_{\mathbf{j}}$ as v_a and the neighbour of $w_{\mathbf{j}}$ is w_b . If $a = b$, the result holds. If $a \neq b$, Lemma 3.15 applied to the pair $v_{\mathbf{j}}, w_b$ gives a double edge $[\mathbf{j}, a]$. Similarly, there is a double edge $[\mathbf{j}, b]$ in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. There are only 3 edges from \mathbf{j} in $\mathcal{G}_2(\overline{\mathbb{F}_p})$, so we obtain a contradiction.

2. Suppose now that $p \equiv 3 \pmod{4}$. By Corollary 3.9, either both $v_{\mathbf{j}}, w_{\mathbf{j}}$ lie on the surface or they both lie on the floor of their respective components. Considering these two cases:

- (a) Case 1: The vertex $v_{\mathbf{j}}$ is on the surface of $\mathcal{G}_2(\mathbb{F}_p)$ component V and $w_{\mathbf{j}}$ is on the surface of component W . Since $v_{\mathbf{j}}$ is on the surface of V , it has three (not necessarily distinct) neighbours v_a, v_b and v_c .

By Lemma 3.11, the three neighbors of $v_{\mathbf{j}}$ in V give the three neighbors of $v_{\mathbf{j}}$ in $\mathcal{G}_2(\overline{\mathbb{F}_p})$: Any neighbor of $w_{\mathbf{j}}$ in W has to be one of w_a, w_b or w_c . Any set of neighbors of $v_{\mathbf{j}}$ in V (counted with multiplicity) is a subset of the neighbors of $w_{\mathbf{j}}$. Since v_a and w_a are both floor vertices and $a \neq b, c$, the vertices corresponding to b and c are on the surface. Suppose $b = c$. Since there are only two vertices with j -invariant b , $w_{\mathbf{j}}$ is attached to the same two j -invariants v_b, w_b as $v_{\mathbf{j}}$ is. Then we see a cycle on the surface of length 4, and this is a contradiction since for $p \equiv 3 \pmod{4}$, the class number $h(-p)$ is odd. Hence $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ have the same set of neighbors and those neighbours are all distinct.

- (b) Case 2: $v_{\mathbf{j}}$ and $w_{\mathbf{j}}$ are on the floors of their respective $\mathcal{G}_2(\mathbb{F}_p)$ components, V and W . Let v_a denote the neighbour of $v_{\mathbf{j}}$, where v_a lies on the surface. Let w_b denote the surface neighbor of $w_{\mathbf{j}}$. Suppose $a \neq b$: We will show this leads to a contradiction. Lemma 3.15 applied to the pair $v_{\mathbf{j}}, w_b$ gives a double edge $[\mathbf{j}, b]$ in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. Similarly,

we obtain a double edge $[\mathbf{j}, a]$ in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ as well. This would mean that there are four inequivalent edges from \mathbf{j} in the graph $\mathcal{G}_2(\overline{\mathbb{F}_p})$, which is not possible so $a = b$. \square

Corollary 3.28 (Isogenies for twists). *Let $\phi : E \rightarrow E'$ be an \mathbb{F}_p -isogeny of degree 2, $j(E), j(E') \neq 1728$. Then, there is an \mathbb{F}_p -isogeny of degree 2 between the quadratic twists $E^t \rightarrow (E')^t$.*

Proof. Suppose $\phi : E \rightarrow E'$ as in the statement, with $j(E) = \mathbf{a}, j(E') = \mathbf{b}$ and E corresponds to the vertex $v_{\mathbf{a}}$. ϕ corresponds to an edge $[v_{\mathbf{a}}, v_{\mathbf{b}}] \in \mathcal{G}_2(\mathbb{F}_p)$. Let $w_{\mathbf{a}}$ be the vertex in $\mathcal{G}_2(\mathbb{F}_p)$ corresponding to the quadratic twist E^t . Proposition 3.27 gives a neighbour $w_{\mathbf{b}}$ of $w_{\mathbf{a}}$ such that $[w_{\mathbf{a}}, w_{\mathbf{b}}] \in \mathcal{G}_2(\mathbb{F}_p)$.

If $w_{\mathbf{b}}$ corresponds to the twist $(E')^t$, then the edge $[w_{\mathbf{a}}, w_{\mathbf{b}}]$ gives the desired \mathbb{F}_p -isogeny.

If, instead, $w_{\mathbf{b}} = v_{\mathbf{b}}$, there are two edges $[v_{\mathbf{a}}, v_{\mathbf{b}}], [w_{\mathbf{a}}, v_{\mathbf{b}}] \in \mathcal{G}_2(\mathbb{F}_p)$. Suppose $z_{\mathbf{b}}$ is the vertex of $\mathcal{G}_2(\mathbb{F}_p)$ corresponding to $(E')^t$. Since we assumed $\mathbf{b} \neq 1728$, Proposition 3.27 gives that $z_{\mathbf{b}}$ also has two neighbours with j -invariants \mathbf{a} . This means there must be edges $[z_{\mathbf{b}}, v_{\mathbf{a}}]$ and $[z_{\mathbf{b}}, w_{\mathbf{a}}]$ in $\mathcal{G}_2(\mathbb{F}_p)$, and $[z_{\mathbf{b}}, w_{\mathbf{a}}]$ gives the desired \mathbb{F}_p -isogeny $E^t \rightarrow (E')^t$. \square

Corollary 3.29 (Attachment along a j -invariant for $\ell = 2$). *Attachment along a j -invariant cannot happen for $\ell = 2$.*

Proof. Proposition 3.27 shows that, except at 1728, the neighbours of the twists are exactly the same. Attachment along a j -invariant (Definition 4) only happens when at least one of the neighbours is distinct.

At $\mathbf{j} = 1728$, we saw in 3.1.3 that the twists are connected by a 2-isogeny in $\mathcal{G}_2(\mathbb{F}_p)$. \square

By a new edge in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ we mean an edge that does not come from an edge in $\mathcal{G}_2(\mathbb{F}_p)$.

Proposition 3.30 (Possible new edges and attachments). *A new edge in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ between \mathbb{F}_p j -invariants can only be added between vertices whose j -invariants correspond to the roots of*

$$f(X) = X^2 + 191025X - 121287375$$

in \mathbb{F}_p , provided these are supersingular \mathbb{F}_p j -invariants not equal to $-3375, 1728$ or 0 .

Attachment cannot happen at $\mathbf{j} = 0, 1728$ or -3375 .

Proof. Let $v_{\mathbf{a}}, w_{\mathbf{b}} \in \mathcal{G}_2(\mathbb{F}_p)$ correspond to j -invariants \mathbf{a}, \mathbf{b} , respectively, such that there is no edge in $\mathcal{G}_2(\mathbb{F}_p)$ between $v_{\mathbf{a}}$ and $w_{\mathbf{b}}$, but there is an edge $[a, b]$ in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. By Lemma 3.15, we obtain a two inequivalent edges $[\mathbf{a}, \mathbf{b}], [\mathbf{a}, \mathbf{b}] \in \mathcal{G}_2(\overline{\mathbb{F}_p})$. By Lemma 2.14, \mathbf{a}, \mathbf{b} must both be one of $0, 1728, -3375$ or an \mathbb{F}_p root of $f(X) = X^2 + 191025X - 121287375$. However, no new edges can occur at the j -invariants $0, 1728$ and -3375 (see the discussion after the proof of Lemma 2.14):

1. For $\mathbf{j} = 0$ is already connected to its only neighbor $\mathbf{j} = 54000$ in $\mathcal{G}_2(\mathbb{F}_p)$, as there are no isolated points in $\mathcal{G}_2(\mathbb{F}_p)$.
2. For $\mathbf{j} = 1728$, all 2-isogenies are defined over \mathbb{F}_p .
3. For $\mathbf{j} = -3375$, there are always two self-loops. Attachment is not possible, as it would require two additional inequivalent outgoing 2-isogenies, giving 4 edges at -3375 in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. \square

Proposition 3.31 (Folding happens for the component containing $\mathbf{j} = 1728$). *Let $p \equiv 3 \pmod{4}$ be prime. The connected component $V \in \mathcal{G}_2(\mathbb{F}_p)$ containing the vertices corresponding to $\mathbf{j} = 1728$ is symmetric over a reflection passing through the vertices v_{1728} lying on the surface of V and w_{1728} lying on the floor of V . In particular, the component V folds when we pass from $\mathcal{G}_2(\mathbb{F}_p)$ to \mathcal{S} .*

To understand this symmetry, picture the surface of the component V as a perfect circle with equidistant vertices and all the edges to the floor are perpendicular to the surface. Then V is symmetric with respect to the line extending the edge $[v_{1728}, w_{1728}]$.

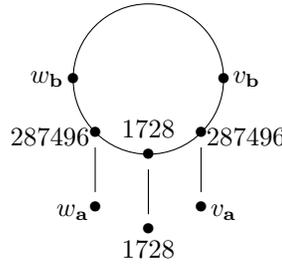
This symmetry is mentioned in Remark 5 of [CLM⁺18], albeit without proof or reference.

Proof. For $p \equiv 3 \pmod{4}$, the possible shapes of the component V are described in Section 3.1.3.

1. Case $p \equiv 3 \pmod{8}$. V is a claw (see Figure 3.2) and the proof of Proposition 3.30 shows that there is one 2-isogeny down from the surface vertex corresponding to $\mathbf{j} = 1728$ to each vertex with j -invariant 287496 and the other vertex corresponding j -invariant 1728. The claw V is clearly symmetric and folds as described.
2. Case $p \equiv 7 \pmod{8}$. In this case, $h(-p)$ is odd.

We may assume that $h(-p) > 1$ (otherwise we are in the claw situation discussed above).

$v = v_{287496}$ and $w = v_{287496}^t$ are both on the surface. By Proposition 3.27, their neighbours have the same j -invariants, say: 1728, \mathbf{a}, \mathbf{b} . Say the neighbours of v are $v_{\mathbf{a}}, v_{\mathbf{b}}$ and the neighbours of w are $w_{\mathbf{a}}, w_{\mathbf{b}}$. Assume that $v_{\mathbf{a}}$ is on the floor. Since $\mathbf{a} \neq 1728$, Lemma 3.9 tells us $w_{\mathbf{a}}$ is also on the floor. Thus, both $v_{\mathbf{b}}$ and $w_{\mathbf{b}}$ are on the surface and the symmetry is preserved.



Continuing in this manner, because $h(-p)$ is odd, we will arrive at a pair of vertices $v_{\mathbf{c}}, w_{\mathbf{c}}$ that share an edge, accounting for all of the vertices in the component. The symmetry holds. \square

Remark 3.32. Proposition 3.31 shows, for $p \equiv 7 \pmod{8}$, the 2-isogeny between the pair of vertices $v_{\mathbf{c}}, w_{\mathbf{c}}$ corresponding to the same \mathbf{j} -invariant \mathbf{c} at the end of the process will be precisely one loop at \mathbf{c} in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. The only vertices with precisely one self-isogeny in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ are $\mathbf{j} = 8000$ and $\mathbf{j} = 1728$. Since $v_{\mathbf{c}}, w_{\mathbf{c}}$ are on the surface of $\mathcal{G}_2(\overline{\mathbb{F}_p})$, $\mathbf{c} = 8000$ (see Section 2.2.1). There is an \mathbb{F}_p -rational 2-power isogeny between any two supersingular elliptic curves with \mathbf{j} -invariants 1728 and 8000.

Corollary 3.33 (Folding). Suppose $V \subset \mathcal{G}_2(\overline{\mathbb{F}_p})$ is a component which folds in $\mathcal{S} \subset \mathcal{G}_2(\overline{\mathbb{F}_p})$.

1. If $p \equiv 1 \pmod{4}$, then V is a single edge between two vertices with j -invariant 8000.
2. If $p \equiv 3 \pmod{4}$, then V contains both the vertices corresponding to $\mathbf{j} = 1728$.

Proof. 1. If $p \equiv 1 \pmod{4}$, then V is an edge: $[v_{\mathbf{a}}, v_{\mathbf{b}}]$. Folding happens if and only if $\mathbf{a} = \mathbf{b}$, resulting in a self-2-isogeny in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. For $p \equiv 1 \pmod{4}$, the only vertices with self-2-isogenies are $\mathbf{j} = -3375, 8000$, when these j -invariants are supersingular (see Section 2.2.1).

For $j = 8000$, there is a 2-isogeny from the curve with j -invariant 8000 given by the equation $E : y^2 = x^3 - 4320x - 96768$ to its twist $y^2 = x^3 - 17280x - 774144$. The latter is a twist of E by 2, and 8000 is only supersingular for $p \equiv 5 \pmod{8}$, so 2 is a nonsquare modulo p .

For $j = -3375$, there are two self-loops in $\mathcal{G}_2(\overline{\mathbb{F}_p})$, and at least one of them not defined over \mathbb{F}_p . Applying Lemma 3.15 to this loop, we conclude that neither of these loops are defined over \mathbb{F}_p and folding does not happen for the edge containing -3375 .

2. If $p \equiv 3 \pmod{4}$, let V be a component that folds. The surface has $h(-p)$ vertices and this class number is odd. We assume that V folds, so every vertex in it gets identified with the vertex corresponding to its twist. By Corollary 3.9, for $j \neq 1728$, the two vertices are either both on the surface or both on the floor. Since there are odd number of vertices on the surface, there cannot only be pairs of twists on the surface, so V must contain the two vertices corresponding to $\mathbf{j} = 1728$, one on the floor and the other on the surface. \square

Now, we prove Theorem 3.26.

Proof of Theorem 3.26. Recalling the possible events when passing from $\mathcal{G}_2(\mathbb{F}_p)$ to $\mathcal{S} \subset \mathcal{G}_2(\overline{\mathbb{F}_p})$. We identify the vertices with the same j -invariants, causing:

1. Folding: Vertices corresponding to twists of the same j -invariant lie on the same component and get identified when we pass to $\mathcal{G}_2(\overline{\mathbb{F}_p})$.
2. Stacking: two isomorphic volcanoes (not just as graphs, but with vertices corresponding to the same j -invariants) have the twist vertices identified.
3. Attaching along a j -invariant: Corollary 3.29 shows this is not possible.

First, let $p \equiv 1 \pmod{4}$. The components of $\mathcal{G}_2(\mathbb{F}_p)$ are edges. Corollary 3.33 shows that the edge containing the two vertices with j -invariant 8000 folds (if 8000 is a supersingular j -invariant for p , i.e. $p \equiv 5 \pmod{8}$). For the other edges, Proposition 3.27 says that for any edge $[v_a, v_b] \in \mathcal{G}_2(\mathbb{F}_p)$ the twists w_a, w_b also give an edge $[w_a, w_b] \in \mathcal{G}_2(\mathbb{F}_p)$. Moreover, Proposition 3.30 gives that there is at most 1 attachment among these edges.

For $p \equiv 3 \pmod{4}$, take any component V of $\mathcal{G}_2(\mathbb{F}_p)$ and any vertex v_a on the surface of V , $a \neq 1728$. Choose a neighbour v_b of v_a . Continue along the surface in the direction of the edge $[v_a, v_b]$ and consider the sequence j -invariants of neighbours $\mathcal{V} = \{v_i\}$ until we reach a vertex with j -invariant a . Similarly, on the component W containing the edge w_a, w_b , consider the sequence of j -invariants of the neighbours $\mathcal{W} = \{w_i\}$ until we reach a vertex with j -invariant a (every surface is a cycle, so this will happen in finitely many steps). We have the following possible outcomes:

1. For some i , we find that $v_i \neq w_i$. This means that the curve i away from v_a on V has a different neighbour than its twist, which is i away from w_a . But this can only happen for $w_i = 1728$ and hence the component folds by Proposition 3.31.
2. The sequences are equal, but \mathcal{V} stops at the twist w_a and \mathcal{W} stopped at the curve v_a . Then v_a, w_a are on the same component V and the cycle on the surface has length $2 \cdot \text{length}(\mathcal{V})$. As $h(-p)$ is odd, this is not possible.
3. The sequences \mathcal{V} and \mathcal{W} are the same and the components V and W are isomorphic as graphs upon replacing labels of vertices by their j -invariants. In this case, the components V and W stack.

Finally, in Proposition 3.30, we showed that at most one attachment is possible. □

Finally, we study the possible attachments given by the roots of the polynomial $f(X) = X^2 + 191025X - 121287375$. Because the polynomial $f(X)$ is the Hilbert class polynomial of $\mathbb{Q}(\sqrt{-15})$, roots of $f(X)$ in $\overline{\mathbb{F}_p}$ give supersingular j -invariant if and only if p is inert in $\mathbb{Q}(\sqrt{-15})$. By factoring the discriminant of $f(X)$

$$191025^2 + 4 \cdot 121287375 = 36975700125 = 3^6 \cdot 5^3 \cdot 7^4 \cdot 13^2,$$

we see that there is a root in \mathbb{F}_p if and only if $p \equiv \pm 1 \pmod{5}$. Combining with the congruence condition that $\left(\frac{-15}{p}\right) = -1$, we obtain that there the roots of $f(X)$ are j -invariants of a supersingular elliptic curves defined over \mathbb{F}_p :

1. $p \equiv 1 \pmod{4}$ and $p \equiv 1 \pmod{3}$ and $p \equiv \pm 1 \pmod{5} \rightarrow p \equiv 1, 24 \pmod{60}$
2. $p \equiv 3 \pmod{4}$ and $p \equiv 2 \pmod{3}$ and $p \equiv \pm 1 \pmod{5} \rightarrow p \equiv 11, 59 \pmod{60}$

We have an additional result about when attachment occurs, as a corollary to Proposition 3.30:

Corollary 3.34 (Attachment happens for $p \not\equiv 7 \pmod{8}$). *Suppose that $p \not\equiv 7 \pmod{8}$ and suppose that \mathbf{j} and \mathbf{j}' are two distinct \mathbb{F}_p -roots of $f(X) = X^2 + 191025X - 121287375$ (it suffices to assume $p > 101$). Then, the new edge $[\mathbf{j}, \mathbf{j}'] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$ is an attaching edge.*

Rephrased, this means that attachment happens whenever it can happen (i.e., when the roots of $f(X)$ are in \mathbb{F}_p) for $p \not\equiv 7 \pmod{8}$.

Proof. First, let $p \equiv 1 \pmod{4}$. The $\mathcal{G}_2(\mathbb{F}_p)$ components are horizontal edges. Suppose that the j -invariant \mathbf{j} admits a double edge $[\mathbf{j}, \mathbf{j}'] \in \mathcal{G}_\ell(\overline{\mathbb{F}_p})$ that it is not an attaching edge, i.e., there is an edge $[v_j, v_{j'}]$ in $\mathcal{G}_2(\mathbb{F}_p)$. By Lemma 3.15, there is then a triple edge $[\mathbf{j}, \mathbf{j}]$. This is only possible if $\mathbf{j} = 0$. For \mathbf{j} or \mathbf{j}' to be equal to 0, we would need X to be a factor of $f(X)$. Since $121287375 = 3^6 \cdot 5^3 \cdot 11^3$, for $p > 11$ attachment happens whenever it can.

Next, let $p \equiv 3 \pmod{4}$. The components of $\mathcal{G}_2(\mathbb{F}_p)$ are claws. If the double edge is not between two different components, then v_j and $v_{j'}$ are on the same claw (for some choice of the twists). Assume, $\mathbf{j} \neq 1728$, they both lie on the floor.

Let v_a be the unique surface vertex of V (see Figure 3.10). This gives us two distinct loops in

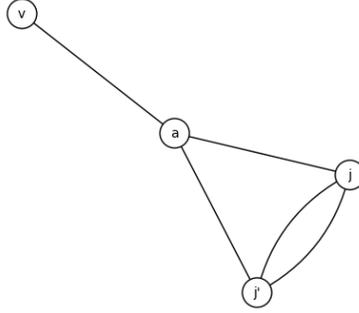


Figure 3.10: The double edge from \mathbf{j} to \mathbf{j}' .

$\mathcal{G}_2(\overline{\mathbb{F}_p})$ of length 3.

These correspond to endomorphisms of norm 8 in $\text{End}_{\overline{\mathbb{F}_p}}(E_j)$.

We check for the existence of such an endomorphism using the modular polynomial $\Phi_8(X, X)$: We need to check whether the roots of $f(X)$ can simultaneously be the roots of the polynomial

$$\begin{aligned} \Phi_8(X, X) = & (-1) \cdot (X - 16581375)^2 \cdot (X - 287496)^2 \cdot (X + 3375)^2 \cdot (X^2 - 52250000X + 12167000000) \\ & \cdot (X^3 + 3491750X^2 - 5151296875X + 12771880859375)^2 \\ & \cdot (X^3 + 39491307X^2 - 58682638134X + 1566028350940383)^2. \end{aligned}$$

Take the resultant

$$\text{Res}(f(X), \Phi_8(X, X)) = (-1) \cdot 3^{72} \cdot 5^{36} \cdot 7^{48} \cdot 11^{34} \cdot 13^{24} \cdot 37^{10} \cdot 41^2 \cdot 43^8 \cdot 59^2 \cdot 71 \cdot 89^2 \cdot 101^2.$$

For primes $p > 101$, this will be nonzero, and there is no such a loop in $\mathcal{G}_2(\overline{\mathbb{F}_p})$, hence attachment happens. In the factorization of the resultant, there is one prime $p \equiv 11, 59 \pmod{60}$ and $3 \pmod{4}$. For $p = 11$, we only have one connected component of $\mathcal{G}_2(\mathbb{F}_p)$, for $p = 59$, attachment happens. \square

In the case $p \equiv 7 \pmod{8}$, attachments that can happen do *not* necessarily. We checked this for all primes $p \equiv 7 \pmod{8}$ between 50,000 and 100,000 such that the primes above 2 do not generate the class group (in this case, there is only one component in $\mathcal{G}_2(\mathbb{F}_p)$, see the following Section 3.5). There are 217 such primes, and for 41 of them the attachment happens. However, there are 12 primes p for which the attachment can happen ($p \equiv 11$ or $59 \pmod{60}$) but there is no attachment:

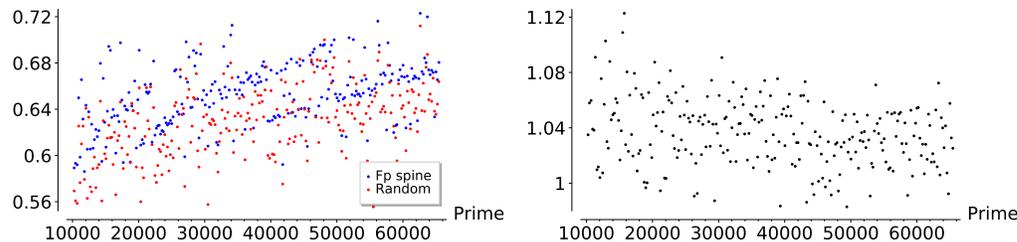
$$53639, 58511, 66959, 71879, 72431, 72551, 79151, 86711, 88919, 90239, 93911, 99719.$$

For $p = 53639$, the two roots of $f(X)$ are $\mathbf{j} = 30505$ and $\mathbf{j} = 46665$. There are two elliptic curves with these j -invariants on the same component of $\mathcal{G}_2(\mathbb{F}_p)$ which are 48 edges apart.

3.5 Distances of components of the \mathbb{F}_p -subgraph \mathcal{S} .

In the above section, we have fully described how the spine \mathcal{S} is formed by passing from $\mathcal{G}_2(\mathbb{F}_p)$ to $\mathcal{G}_\ell(\mathbb{F}_p)$. A natural question is how the spine \mathcal{S} sits inside the graph $\mathcal{G}_\ell(\mathbb{F}_p)$.

For primes $p \equiv 1 \pmod{4}$, the subgraph is given by single edges (with a possibility of a few isolated vertices and one component of size 4), as we proved in Section 3.2. These components seem to be distributed the same way in the graph as random vertices: we compare the mean of the distances of the components with the distances between random vertices (100 random choices), normalized by the diameter. We compared these distances for 254 primes with $p \equiv 1 \pmod{4}$ from 10253 to 65437. The primes were chosen to be spaced with a gap of at least 200. Our results are shown in Figure 3.11.



(a) Normalized distance between \mathcal{S} components (blue) and random pairs (red) (b) Ratio of distance between components vs distance between random pairs

Figure 3.11: Comparison of distances of \mathcal{S} components versus distances between random vertices for $p \equiv 1 \pmod{4}$.

We do not know how to explain that the average distances between components seem to be larger than distance between two random points in $\mathcal{G}_2(\mathbb{F}_p)$.

3.5.1 $p \equiv 7 \pmod{8}$

We start with the following easy lemma.

Lemma 3.35. *Let $p \equiv 7 \pmod{8}$ and set $K := \mathbb{Q}(\sqrt{-p})$. Let \mathfrak{p} denote a prime ideal of \mathcal{O}_K above (2) and suppose that $\langle \mathfrak{p} \rangle = \text{Cl}(\mathcal{O}_K)$.*

Then, the $\mathcal{G}_2(\mathbb{F}_p)$ is has only one connected component with

$$\# \text{Cl}(\mathcal{O}_K) = h(-p)$$

vertices on the surface and from every vertex on the surface, there is exactly one isogeny down.

A fortiori, the spine $\mathcal{S} \subset \mathcal{G}_\ell(\mathbb{F}_p)$ is connected.

Proof. An immediate consequence of 3.1.3. □

It is interesting to know that the converse of this lemma is not true: If primes above (2) do not generate the class group, it is still possible for the \mathbb{F}_p subgraph of $\mathcal{G}_\ell(\mathbb{F}_p)$ to be connected, thanks to attaching.

In the range $50,000 < p < 10000$, there are 217 primes for which \mathfrak{p} does not generate the class group.

We have seen the following:

1. for 12 primes 57119, 59471, 61871, 64439, 70439, 76871, 85199, 88799, 91631, 92399, 92951, and 96671 the spine \mathcal{S} is nonetheless connected.
2. for 57 out of those 66 primes there will be exactly 2 connected components of \mathcal{S} and those will be at most 6 apart (with diameter being about 15). For 29 of these primes, 51287, 51383, 53639, 54559, 54767, 58511, 59063, 63439, 63799, 65831, 66863, 67751, 69191,

70607, 72679, 74759, 76159, 79151, 80783, 82799, 83471, 84559, 85847, 86711, 90239, 91823, 95959, 99079, and 99719, these components are exactly 2 apart.

The diameter is between 14 and 16. The graphs have between 5400 and 8300 vertices.

These are the distances of non-normalized. The average distance of two random vertices for 2-isogeny graphs of this size is around 9. This is approximately 0.6 times the diameter of the graphs. This number grows slowly (for primes $p \approx 500,000$, the average distance of two random vertices is about 0.7 times the diameter) and we expect it to converge to the diameter, however, we don't know how quickly.

We also computed the average of the mean distances of connected components of $\mathcal{S} \subset \mathcal{G}_2(\mathbb{F}_p)$ for these primes. The mean is 4.3395, with standard deviation 1.1092, and the maximum is 7.000 and the minimum 2.333, which indicates that the components tend to be close to each other.

3.5.2 The number of components

We estimate the number of connected components of \mathcal{S} , under the assumption $h(D) \approx \sqrt{D}$. By Theorem 3.26, the number of vertices of \mathcal{S} is approximately half (respectively, one fourth) of the size of \mathcal{S} if $p \equiv 1 \pmod{4}$ (resp., $p \equiv 3 \pmod{8}$) and depends on the order of the prime lying above 2 for $p \equiv 7 \pmod{8}$.

prime mod 8	shape of $\mathcal{G}_2(\mathbb{F}_p)$	$\#\mathcal{S}$	\approx number of components
1 mod 4	edges	$\frac{1}{2}h(-4p)$	$\frac{1}{4}h(-4p)$
3 mod 8	claw	$2h(-p)$	$\frac{1}{4} \cdot 2h(-p) = \frac{1}{2}h(-p)$
7 mod 8	volcanoes (2 levels, size $\text{ord}(\mathfrak{p}_2)$)	$h(-p)$	$\frac{1}{2 \cdot \text{ord}(\mathfrak{p}_2)} \cdot h(-p) \ll \frac{1}{2}h(-p)$

4 Conjugate vertices, distances, and the spine

We examine several distances of cryptographic interest. In Section 4.1 we study the distance between Galois conjugate pairs of vertices, that is, pairs of j -invariants of the form j, j^p . Our data suggests these vertices are closer to each other than a random pair of vertices in $\mathcal{G}_2(\overline{\mathbb{F}_p})$. In Section 4.2 we test how often the shortest path between two conjugate vertices goes through the spine \mathcal{S} , or equivalently, contains a j -invariant in \mathbb{F}_p . We find conjugate vertices are more likely than a random pair of vertices to be connected by a shortest path through the spine. Finally, we examine the distance between arbitrary vertices and the spine \mathcal{S} in Section 4.3.

4.1 Distance between conjugate pairs

Isogeny-based cryptosystems such as cryptographic hash functions and key exchange rely on the difficulty of computing paths (*routing*) in the supersingular graph $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Our experiments with $\ell = 2$ show that two random conjugate vertices are "closer" than two random vertices.

We tested the distances of conjugate vertices as follows. First for a given prime p , we constructed the graph $\mathcal{G}_2(\overline{\mathbb{F}_p})$. Then we computed the distances $\text{dist}(j_1, j_2)$ between all pairs $j_1, j_2 \in \mathcal{G}_2(\overline{\mathbb{F}_p})$. These values were organized into two lists:

$$C_p = [\text{dist}(j, j^p) : j \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p]$$

$$A_p = [\text{dist}(j_1, j_2) : j_1, j_2 \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p].$$

The distributions C_p and A_p for $p = 19489$ are shown as histograms in Figure 4.1. We call the pairs from C_p *conjugate* pairs and pairs from A_p *arbitrary* pairs.

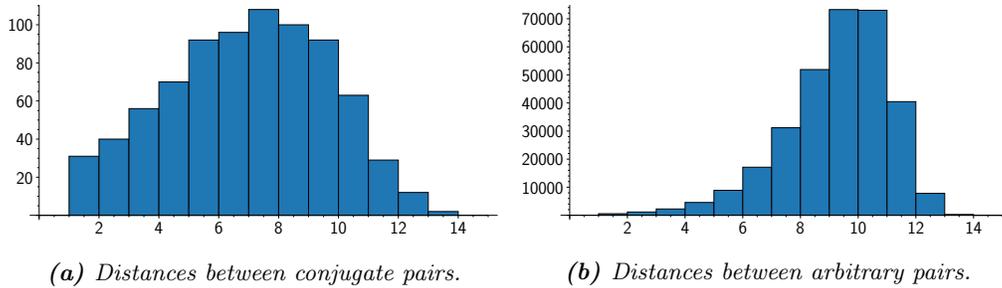


Figure 4.1: Distances measured between conjugate pairs and arbitrary pairs of vertices not in \mathbb{F}_p for the prime $p = 19489$.

For a larger prime, it is too costly to iterate over all vertices. Instead, we took a random sample of 1000 conjugate and arbitrary pairs. The data collected for the prime $p = 1000003$ is shown in Figure 4.2.

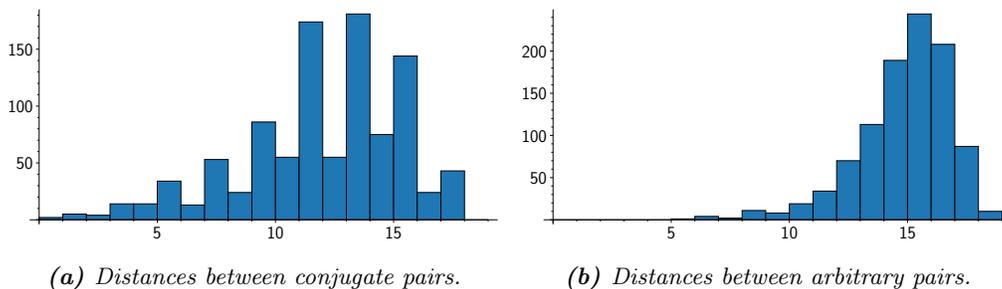


Figure 4.2: Distances between 1000 randomly sampled pairs of arbitrary and conjugate vertices for the prime $p = 1000003$.

From our data, it seems likely that distances between conjugate vertices have a different distribution than distances between arbitrary vertices. However, more study on a broader sample of primes is needed.

Remark 4.1. In Figure 4.2, we see a clear bias towards paths of odd length (that is, odd number of edges). This is due to the fact that conjugate j -invariants often admit a shortest path that is a mirror path (Definition 2.11). These paths do not usually go through the spine \mathcal{S} , so they have an even number of vertices and an odd number of edges. This topic is studied further in Section 4.2.

4.2 How often do shortest paths go through the \mathbb{F}_p -spine

It was shown in [DG16] that if one navigates to the spine \mathcal{S} , one obtains a subexponential attack on the path finding problem. This attack, however, uses L -isogenies, where L is a set of small primes. We study the situation when one only uses $L = \{2\}$. When $j' = j^p$, any path from j to the spine \mathcal{S} can then be mirrored to obtain a path of equal length from j^p to the same point of the spine, and hence a path between j and j^p passing through the spine. This notion motivates the following definition:

Definition 4.2. A pair of vertices are **opposite** if there exists a shortest path between them that passes through the \mathbb{F}_p spine.

4.2.1 Experimental methods

We tested how often a shortest path between two conjugate vertices went through the spine \mathcal{S} . Shortest paths are not necessarily unique, so it is not enough to compute a shortest path and check

whether passes through the spine. We used the built-in function of Sage ([The19]) to perform our computations. For efficiency, we did not compute all the shortest paths. Instead, to verify whether a pair j_1, j_2 is opposite, we run over all vertices in $\mathbf{j} \in \mathbb{F}_p$ and check whether there is a \mathbf{j} such that

$$\text{dist}(j_1, j_2) = \text{dist}(j_1, \mathbf{j}) + \text{dist}(\mathbf{j}, j_2).$$

For smaller primes (< 5000) we computed the proportions for all pairs of vertices in $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$. For larger primes, we randomly selected 1000 pairs of points j_1, j_2 in $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ and checked whether each of the pairs $(j_1, j_2), (j_1, j_1^p)$ were opposite.

4.2.2 Conjugate pairs vs arbitrary pairs

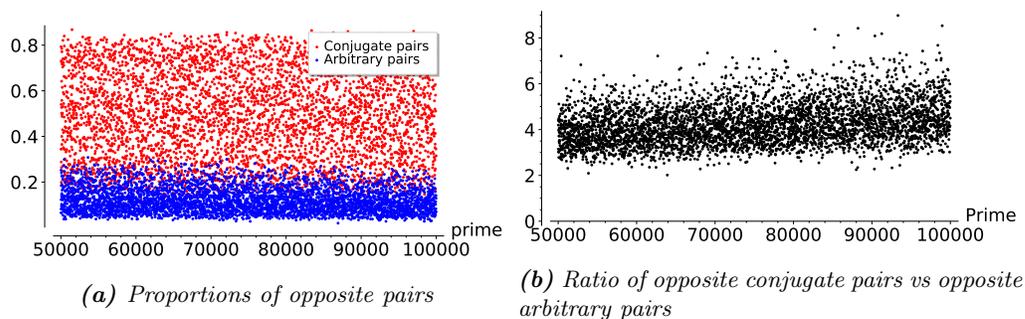


Figure 4.3: Data for random sample of 1000 pairs of conjugate and arbitrary pairs.

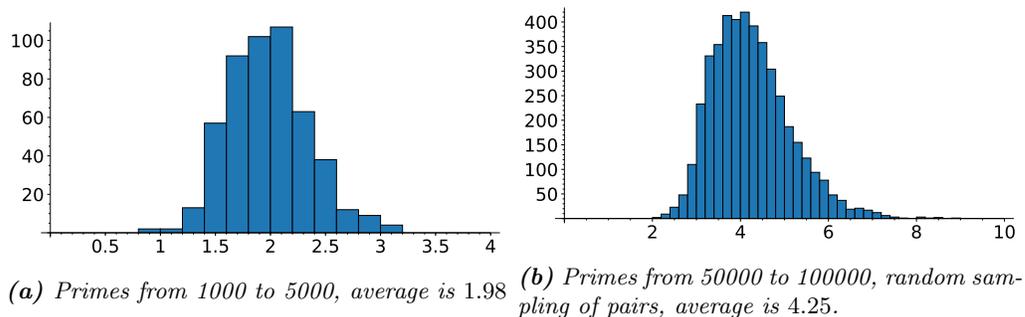


Figure 4.4: Histogram of primes with proportion of opposite conjugate pairs divided by the proportion of opposite arbitrary pairs as in (5).

Our data suggests that conjugate vertices are more likely to be opposite than arbitrary vertices. For a random sampling of pairs over primes between 50000 and 100000, we observe that

$$\text{average} \left(\frac{\#\text{opposite conjugate pairs}}{\#\text{opposite arbitrary pairs}} \right) \approx 4.25 \quad (5)$$

The ratio seems to increase with the size of the prime, as seen in Figures 4.3 and 4.4. This leads to the following observation: Due to the mirror involution, to build the graph $\mathcal{G}_\ell(\mathbb{F}_p)$, one can start with the spine \mathcal{S} and keep adding edges along with their mirror edges. This might suggest that the spine is central to the graph. However, the shortest paths between arbitrary pairs of vertices are less likely to pass through the spine, contradicting that perspective.

4.2.3 Proportions varying over different residue classes

We observe that the proportion of pairs of opposite vertices varies based on the residue class of p . In this section, we consider arbitrary pairs of vertices. From the data, as shown in Figure 4.5a, the proportion is higher for primes $p \equiv 2 \pmod 3$ compared to $p \equiv 1 \pmod 3$ and higher for primes $p \equiv \pm 2 \pmod 5$ compared to $p \equiv \pm 1 \pmod 5$.

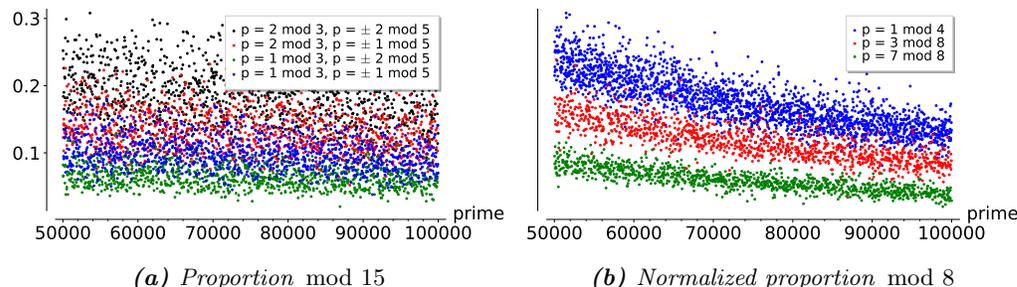


Figure 4.5: Proportion of opposite pairs out of a random sample of 1000 pairs.

Based on our results, we suggest that the size and connectedness of the \mathbb{F}_p spine could be key factors affecting the proportion of opposite pairs.

1. Size of \mathbb{F}_p spine: when the number of \mathbb{F}_p points is higher, pairs are more likely to have shortest paths through these points.
 - To consider this effect, we study each proportion divided by the number of \mathbb{F}_p points for the prime p . After normalizing the proportions, we no longer see clear differences when considering residue classes mod 3 and mod 5. This suggests that the underlying cause of the difference was the size of the \mathbb{F}_p spine.
 - However, the normalized proportions as shown in Figure 4.5b appear to fall into three classes $p \equiv 1 \pmod 4$, $p \equiv 3 \pmod 8$ and $p \equiv 7 \pmod 8$. One possible cause for this is the connectedness of the \mathbb{F}_p spine.
2. Connectedness of \mathbb{F}_p spine: when the \mathbb{F}_p spine is less connected to itself, pairs are more likely to have shortest paths through \mathcal{S} .
 - From the table in Section 3.5.2, the spine is the least connected when $p \equiv 1 \pmod 4$, and can be highly connected when $p \equiv 7 \pmod 8$. This could explain the difference in proportions when normalized by the size of \mathcal{S} .
 - For example, we consider the cases $p_1 = 19991$ ($p_1 \equiv 7 \pmod 8$, \mathcal{S} is connected, $|\mathcal{S}| = 199$) and $p_2 = 19993$ ($p_2 \equiv 1 \pmod 4$, \mathcal{S} is maximally disconnected, $|\mathcal{S}| = 30$). We would expect $199/30 > 6$ times more opposite pairs in the p_1 case. However, for 1000 random pairs, 266 pairs were opposite for p_1 compared to 112 pairs for p_2 .

To further study whether differences occurring in the normalized proportion mod 8 were due to the connectedness of the \mathbb{F}_p spine or other structures of $\mathcal{G}_2(\overline{\mathbb{F}_p})$, we took a random subgraph of the same size as \mathcal{S} and obtained the proportion of pairs with a shortest path passing through the random subgraph. We took the average of these results over 10 random subgraphs for each prime between 1000 and 5000.

From the data in Figure 4.6, there is less distinction mod 8 for random subgraphs. This suggests that the connectedness of \mathcal{S} is the dominant factor affecting the normalized proportion.

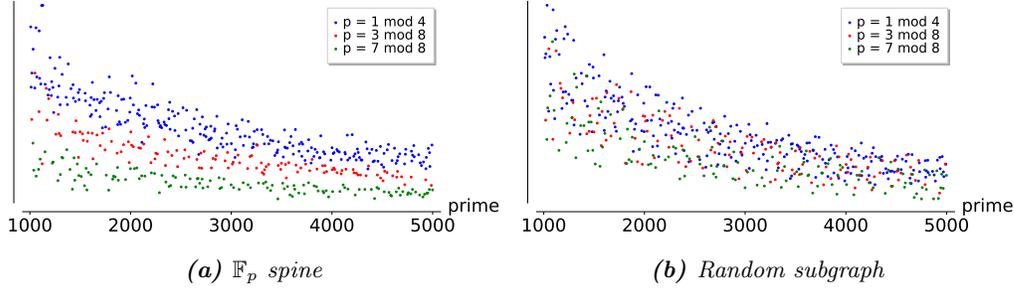


Figure 4.6: Normalized proportion of pairs with a shortest path through the subgraph specified.

4.3 Distance to spine

In this section, we compare the *distance from a random vertex to the spine*, with the *distance from a random vertex to a random subgraph of the same size as the spine*. We observe that if the spine is connected, then the distance to the spine seems greater than the distance to a random subgraph. This agrees with the intuition that a small connected subgraph (remember that the spine has size $O(\sqrt{p})$) will be further from most vertices than a random subgraph, which will have many connected components uniformly distributed throughout the graph.

We tested the distances as follows. For a value of p , we constructed the graph $\mathcal{G}_2(\overline{\mathbb{F}_p})$, the spine $S_0 := \mathcal{S}$, and chose several random subgraphs S_1, \dots, S_n . We define the distance between a vertex j and a subgraph S_i to be

$$\text{dist}(j, S_i) = \min\{\text{dist}(j, j') : j' \in S_i\}.$$

We computed lists $d_i = [\text{dist}(j, S_i) : j \in \mathcal{G}_2(\overline{\mathbb{F}_p})]$ in order to measure how dispersed S_i is in $\mathcal{G}_2(\overline{\mathbb{F}_p})$.

Distances were computed for two primes, $p = 19991$ and $p = 19993$. Histograms of the distributions of the d_i are given in Figure 4.7. For $p = 19991$, the subgraph \mathcal{S} is connected, whereas for $p = 19993$, \mathcal{S} is maximally disconnected because $19993 \equiv 1 \pmod{12}$ (see Lemma 3.10).

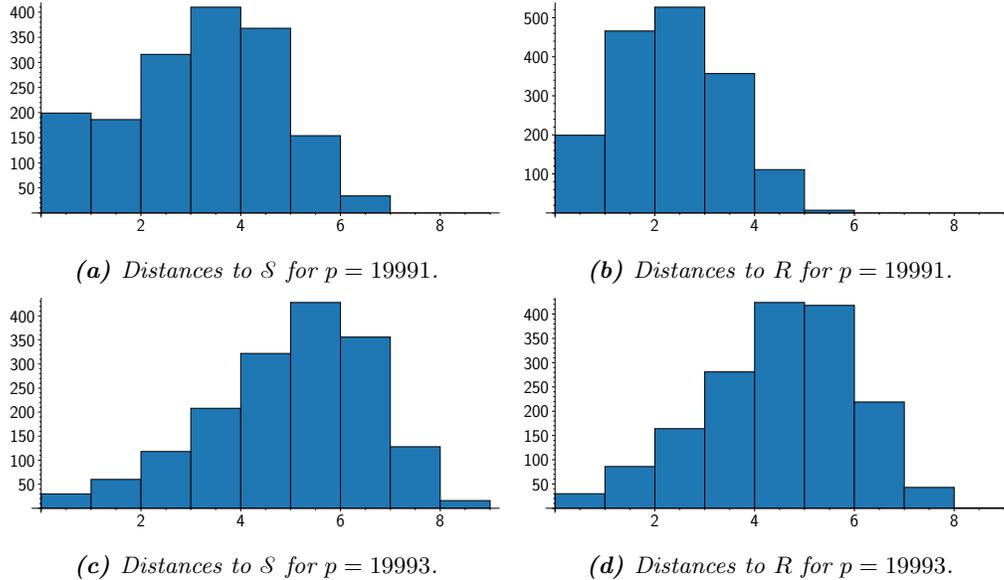


Figure 4.7: Distances to the spine \mathcal{S} compared to distances to a random subgraph of the same size. The subgraph \mathcal{S} is connected for $p = 19991$ and a union of disconnected edges for $p = 19993$.

The significant difference between the two primes shown in Figure 4.7 can also be explained

by the number of vertices in \mathcal{S} . Since $\mathcal{G}_2(\overline{\mathbb{F}_p})$ is a 3-regular graph, for a random vertex j , there are at most $3 \cdot 2^{d-1}$ vertices of distance d away from j (and this limit is achieved if there are no collisions on the paths leaving j). If $\mathcal{G}_2(\overline{\mathbb{F}_p})$ has N vertices and H is a random subgraph with M vertices, then the expected distance to H from a random vertex should be $\approx \log_2(N/M)$.

For $p = 19991$, $|\mathcal{S}| = 199$, so we expect the average distance to \mathcal{S} to be 3.06. For $p = 19993$, $|\mathcal{S}| = 30$, so we expect the average distance to \mathcal{S} to be 5.80.

4.3.1 Comparison across primes p

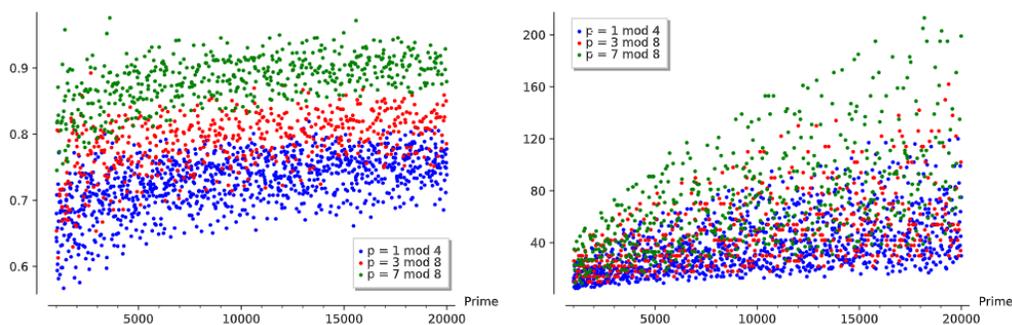
In order to compare the distances to \mathcal{S} across different primes and account for the expected average distance based on the size of \mathcal{S} we consider normalized distances as follows:

$$d_p = (\text{average distance to } \mathcal{S} \text{ for prime } p) / \log_2(|\mathcal{G}_2(\overline{\mathbb{F}_p})|/|\mathcal{S}|)$$

Recall that $\log_2(|\mathcal{G}_2(\overline{\mathbb{F}_p})|/|\mathcal{S}|)$ is the expected distance to the spine from a random vertex. We observed that the average distances were lower than the expected distance based on the connectedness of \mathcal{S} . There are also clear differences in the distributions of d_p when considering residue classes of p modulo 8. This is shown in Figure 4.8a. In particular, the data mod 8 matches our findings on the proportion of opposite pairs, see Figure 4.5b.

However, the different behaviour of d_p for the different congruence classes mod 8 can be explained by the size of the spine. If the size of the spine $|\mathcal{S}|$ is large, we will need fewer steps to reach the spine from a random vertex v . Hence, when counting the paths of length 2^d from v , we will encounter less backtracking and the estimate is more precise. Looking at Figure 4.8b, we see that for $p \equiv 7 \pmod 8$, the size of the spine is the largest, and for $p \equiv 1 \pmod 4$, the size of the spine is the smallest.

We also tested this within a fixed congruence class: for primes with $p \equiv 7 \pmod 8$ and $15,000 < p < 20,000$, the mean distance to the spine is 4.040 with standard deviation 0.413 if $|\mathcal{S}| < 100$ and mean 3.007 with standard deviation 0.335 if $|\mathcal{S}| > 100$.



(a) Comparison varying $p \pmod 8$

(b) Size of spine \mathcal{S} varying $p \pmod 8$

Figure 4.8: Normalized average distances to the \mathbb{F}_p spine versus the size of the spine.

5 When are conjugate j -invariants ℓ -isogenous?

5.1 Motivation

In Section 4.2 we studied paths between conjugate j -invariants in $\mathcal{G}_2(\overline{\mathbb{F}_p})$ that go through the spine \mathcal{S} . On the other hand, if j and j^p are 2-isogenous, then the shortest path between them has length one and does not go through \mathcal{S} . This leads us to the natural question:

Question 1: How often are conjugate $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ j -invariants ℓ -isogenous, for $\ell = 2, 3$?

5.2 Methods

For varying primes p , we want to collect data on how often conjugate $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ j -invariants are ℓ -isogenous, for $\ell = 2, 3$. We used two main approaches for this, and here we compare their efficiency. The first one corresponds to a modified Breadth-First Search (BFS) algorithm. The second one is based on the supersingular j -invariant polynomial and the modular polynomial.

Breadth-First Search. Breadth-First Search (BFS) is an algorithm for exploring a graph starting at a fixed vertex. From the starting vertex, the algorithm explores all of the neighbor nodes at a given depth before moving on to the nodes at the next depth.

In our experiments we generated the graph $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ as follows. First we find a supersingular j -invariant j_0 over \mathbb{F}_p using the CM method described in [Brö09]. It works by finding the smallest prime q such that p is inert in $\mathbb{Q}(\sqrt{-q})$ and then finding a root of the Hilbert class polynomial for $\mathbb{Q}(\sqrt{-q})$ over \mathbb{F}_p . In practice this step is very efficient for small p .

Next we generate the graph using BFS. BFS takes $O(|V| + |E|)$ steps when performed on a graph with $|V|$ vertices and $|E|$ edges. Because $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ is an ℓ -regular graph with $\approx p/12$ vertices, this will run in $O(p)$ steps. Each step requires finding the neighbors of a particular vertex j in the graph. This is done by finding the roots (with multiplicities) of $\Phi_\ell(j, x) \in \mathbb{F}_{p^2}[x]$. So the total time complexity for this step is $\tilde{O}(p)$.

We modified this algorithm to collect the ℓ -isogenous conjugate pairs of $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ j -invariants as we explore the graph, so the time complexity of the algorithm is essentially the same as the complexity for exploring the graph.

Supersingular j -invariant polynomial with Φ_ℓ . Another method to calculate the proportion of ℓ -isogenous conjugate pairs uses the supersingular j -invariant polynomial. Sage has a built-in command to compute this polynomial. It uses the fact that an elliptic curve over $\overline{\mathbb{F}_p}$ given by the Legendre equation $y^2 = x(x-1)(x-\lambda)$ is supersingular if and only if λ is a root of the polynomial

$$H(t) = \sum_{i=0}^m \binom{m}{i}^2 t^i \tag{6}$$

for $m = \frac{p-1}{2}$ (see [Sil09, Section V.4]). If j is the j -invariant of an elliptic curve as above, then the polynomial

$$F(s, t) = st^2(t-1)^2 - 2^8(t^2 - t + 1)^3$$

vanishes at (j, λ) . Sage then computes the resultant $R(s)$ of $H(t)$ and $F(s, t)$ with respect to t , factors it over \mathbb{F}_p and defines the supersingular j -invariant polynomial as the product of these factors counted only once. If $(s - 1728)$ or s are factors, they are excluded.

Once we obtain this polynomial, we proceed as follows:

1. Compute the roots of the supersingular j -invariant polynomial over \mathbb{F}_{p^2} .
2. Sort the roots over $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ by conjugate pairs and count the number of such pairs.
3. Use the ℓ^{th} modular polynomial Φ_ℓ to determine which conjugate pairs are ℓ -isogenous.

For a prime p , the number of supersingular j -invariants over \mathbb{F}_{p^2} is $\lfloor \frac{p}{12} \rfloor + \varepsilon$ for $\varepsilon \in \{0, 1, 2\}$ [Sil09, Thm V.4.1] and the supersingular j -invariant polynomial is thus a polynomial of degree $\lfloor \frac{p}{12} \rfloor$.

5.2.1 Timing data

We expected the BSF algorithm to be faster than the supersingular j -polynomial one, since the latter must factor a polynomial of degree $(p-1)/2$. To experimentally verify the difference in running time, we used both algorithms to find the 2-isogenous conjugate pairs for 37 primes between 103 and 95471. The resulting data is displayed in Figure 5.1.

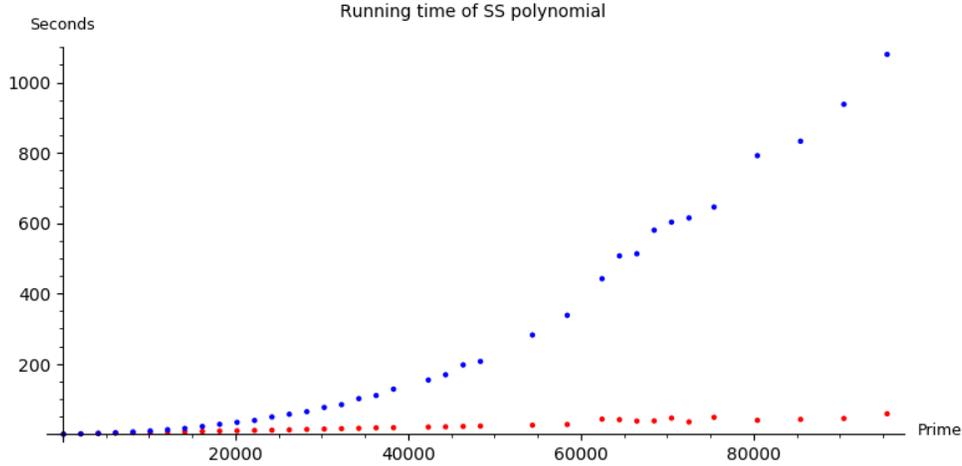


Figure 5.1: Timing data for the BFS (red) and Supersingular polynomial algorithms (blue), showing the time it took each algorithm to find all 2-isogenous conjugate pairs for that prime.

5.3 Experimental data: 2-isogenies

We collected data on supersingular j -invariants over \mathbb{F}_{p^2} for all primes $5 \leq p \leq 100193$ (a total of 9,605 primes). For each p , we collected all of the $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ j -invariants and counted those that are also 2-isogenous. The plot shown in Figure 5.2 shows the proportion of conjugate pairs that are 2-isogenous.

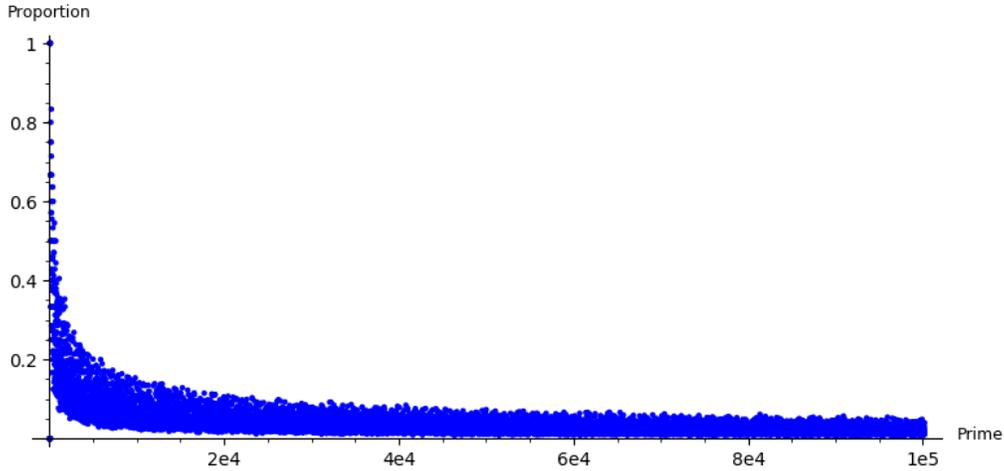


Figure 5.2: Proportion of 2-Isogenous Pairs of Conjugate j -Invariants in $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$. Points are of the form (p, y) where p is a prime and y is the proportion of conjugate pairs of j -invariants which are 2-isogeneous.

With a few exceptions, all of the proportions computed are positive and strictly less than 1. The small primes (roughly $p < 5000$) have a wide range of proportions, between 0 and 1. This is expected due to the small number of points on their $\mathcal{G}_\ell(\overline{\mathbb{F}_p})$ graphs. For example: there are some primes p such that all of the pairs of conjugates are 2-isogenous. On the other hand, if $\mathbb{F}_{p^2} \setminus \mathbb{F}_p = \emptyset$, which can happen for small primes, then the proportion will be trivially zero. Notably, the only examples of p for which the proportion is zero are $p = 101, 131$.

To avoid small prime phenomena, we focused on analyzing the data we collected for $10007 \leq$

$p \leq 100193$ (a total of 8378 primes). When referring to this data, we will use the phrase “main data”. When referring to all of the data collected for $5 \leq p \leq 100193$, we use the phrase “all data”.

The graph of proportions for the main data can be found in Figure 5.3.

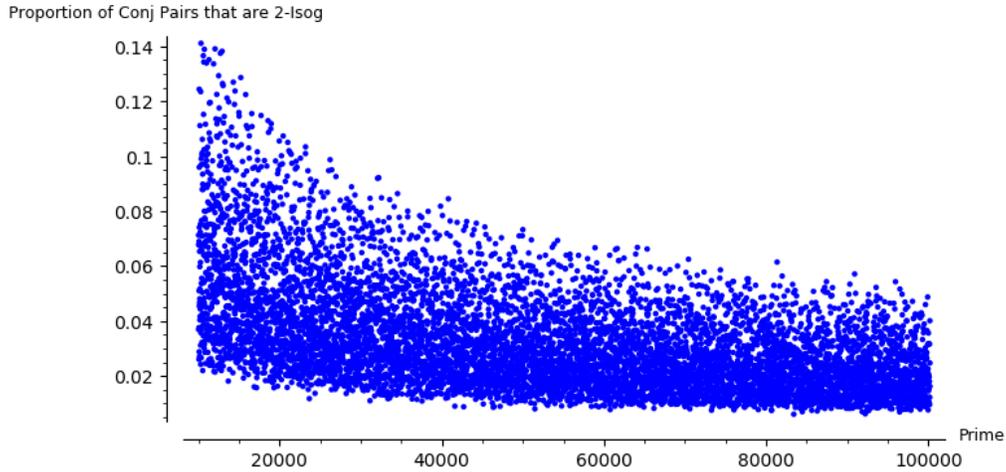


Figure 5.3: Proportion of 2-isogenous conjugate pairs in $\mathcal{G}_2(\overline{\mathbb{F}}_p)$ for $p > 10000$

In this collection of data, for primes $10007 \leq p \leq 100193$, we found there to be a mean proportion of 0.032780 with standard deviation of 0.019134.

We then sorted the data by congruence conditions to look for patterns. The biggest difference appeared when we re-sorted the data according to the congruence class of the primes modulo 12.

5.3.1 Primes Modulo 12

In Table 1, we summarize the differences between the different congruence classes modulo 12. Note the similar, higher means for $p \equiv 1, 7 \pmod{12}$ and the similar, lower means for $p \equiv 5, 11 \pmod{12}$.

	$p \equiv 1 \pmod{12}$	$p \equiv 5 \pmod{12}$
Total # of primes:	2079	2104
Mean:	0.043551	0.021969
Standard Deviation:	0.019815	0.010206
	$p \equiv 7 \pmod{12}$	$p \equiv 11 \pmod{12}$
Total # of primes:	2101	2094
Mean:	0.043375	0.022244
Standard Deviation:	0.020140	0.010512

Table 1: Proportions of 2-isogenous conjugates, $10007 \leq p \leq 100193$, sorted by $p \pmod{12}$

These distributions are skewed according to the congruence class, as we can also see from the graph in Figure 5.4.

There appears to be a correlation between primes $p \equiv 1, 7 \pmod{12}$ and between primes $p \equiv 5, 11 \pmod{12}$. A two-sample t -test confirms these correlations at the 99.8% level.

5.4 Experimental data: 3-isogenies

We collected data on the supersingular j -invariants over \mathbb{F}_{p^2} for all the primes $5 \leq p \leq 100193$ (a total of 9,605 primes) and computed the proportion of conjugate pairs that are also 3-isogenous.

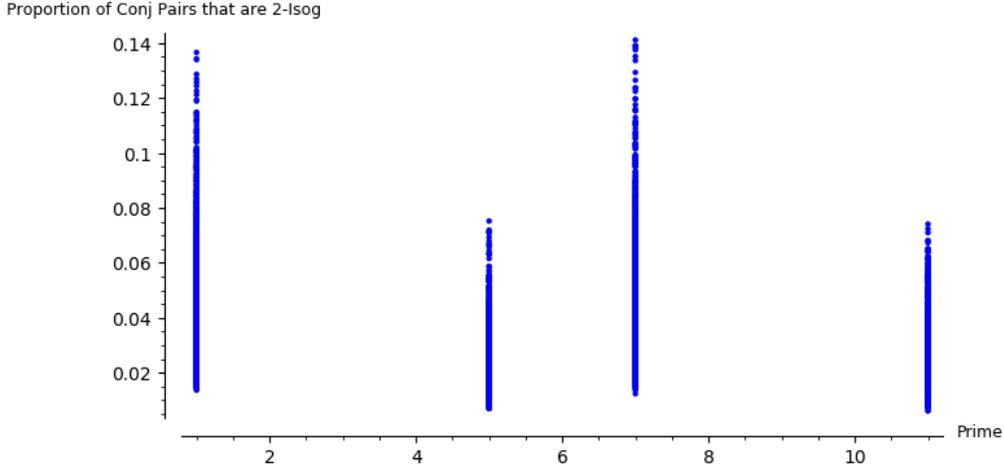


Figure 5.4: Proportions of 2-isogenous conjugates, $10007 \leq p \leq 100193$, sorted by $p \bmod 12$

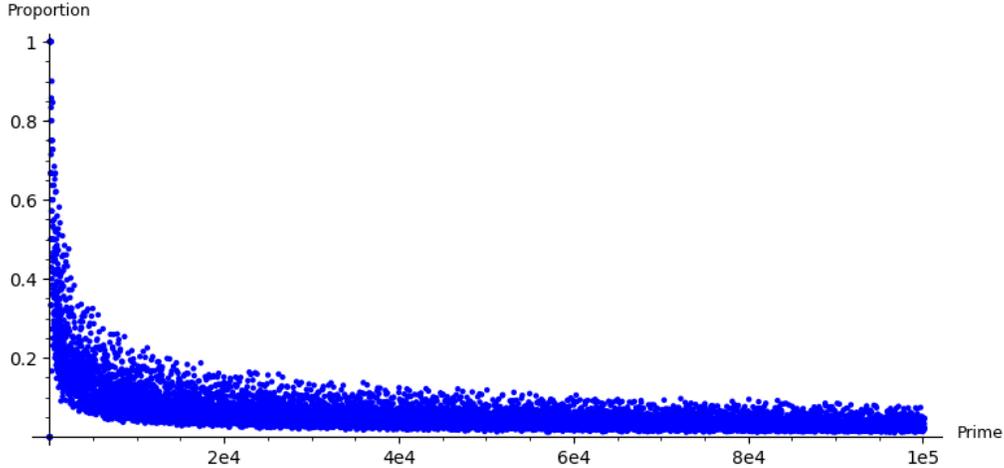


Figure 5.5: Proportion of 3-Isogenous Pairs of Conjugate j -Invariants in $\mathbb{G}_\ell(\overline{\mathbb{F}_p})$

We present this data in the same format as the 2-isogeny data presented in 5.3.

In Figure 5.5, observe the proportions of conjugate pairs of primes for all of the primes we collected data on.

Again, we observe some small prime phenomena (proportions of 1 and 0 for p small). However, in the 3-isogeny case we do not have nontrivial examples of primes p for which the proportion of 3-isogenous conjugates is 0: if there exist conjugate j -invariants in $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$, then there is at least one pair of 3-isogenous conjugates. (Recall that the two counterexamples to this statement in the 2-isogeny case were $p = 101, 131$.)

To avoid small prime phenomena, we again focused on analyzing the data we collected for $10007 \leq p \leq 100193$ (a total of 8378 primes). Again, when referring to this data, we will use the phrase “main data”. When referring to all of the data collected for $5 \leq p \leq 100193$, we use the phrase “all data”.

The graph of proportions for the main data can be found in Figure 5.6. In this collection of data, for primes $10007 \leq p \leq 100193$, we found there to be a mean proportion of 0.047306 with standard deviation of 0.026568.

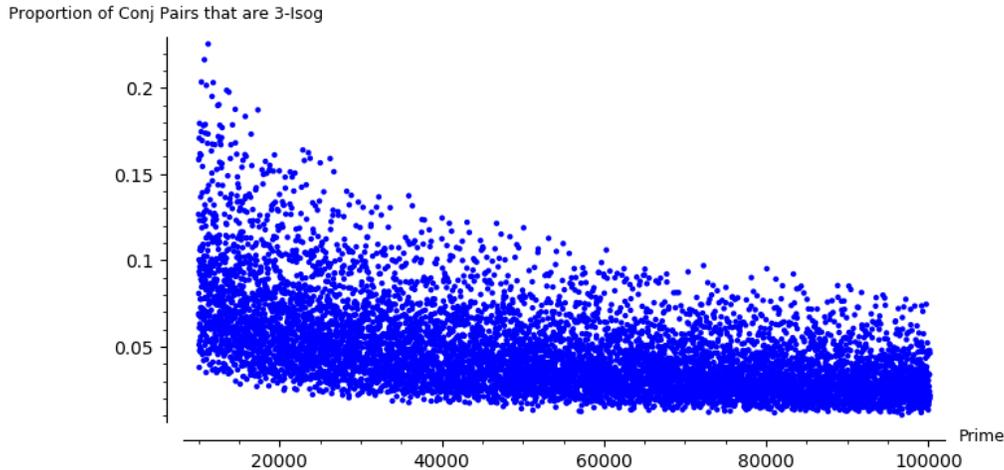


Figure 5.6: Proportion of 3-isogenous conjugate pairs in $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$ for primes $p > 10000$

As in the 2-isogeny case, we again sorted the data by congruence conditions to look for patterns. The biggest difference appeared when we re-sorted the data according to the congruence class of the primes modulo 12.

5.4.1 Primes Modulo 12

In Table 2, we summarize the differences between the different congruence classes modulo 12. Note the similar and higher means for $p \equiv 1, 5 \pmod{12}$ and the similar and lower means for $p \equiv 7, 11 \pmod{12}$.

	$p \equiv 1 \pmod{12}$	$p \equiv 5 \pmod{12}$
Total # of primes:	2079	2104
Mean:	0.058526	0.059034
Standard Deviation:	0.029488	0.029729
	$p \equiv 7 \pmod{12}$	$p \equiv 11 \pmod{12}$
Total # of primes:	2101	2094
Mean:	0.035620	0.036107
Standard Deviation:	0.016369	0.016706

Table 2: Proportions of 3-isogenous conjugates for $10007 \leq p \leq 100193$, sorted by $p \pmod{12}$

These distributions are skewed according to the congruence class, as we can also see from the graph in Figure 5.7.

There appears to be a correlation between primes $p \equiv 1, 5 \pmod{12}$ and between primes $p \equiv 7, 11 \pmod{12}$. A two-sample t -test confirms these correlations at the 99.8% level.

5.5 Analysis of data

Our experimental data suggests that, at least for $\ell = 2, 3$ and with the exception of a few small primes, the proportion of conjugate pairs that are ℓ -isogenous is a small positive number. In particular, all of the primes $p \neq 101, 131$ with supersingular j -invariants in $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$ observed have at least one such pair. This motivates the following two questions:

Question 2: For $p > 131$, is there always at least one pair of ℓ -isogenous conjugate j -invariants on $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$?

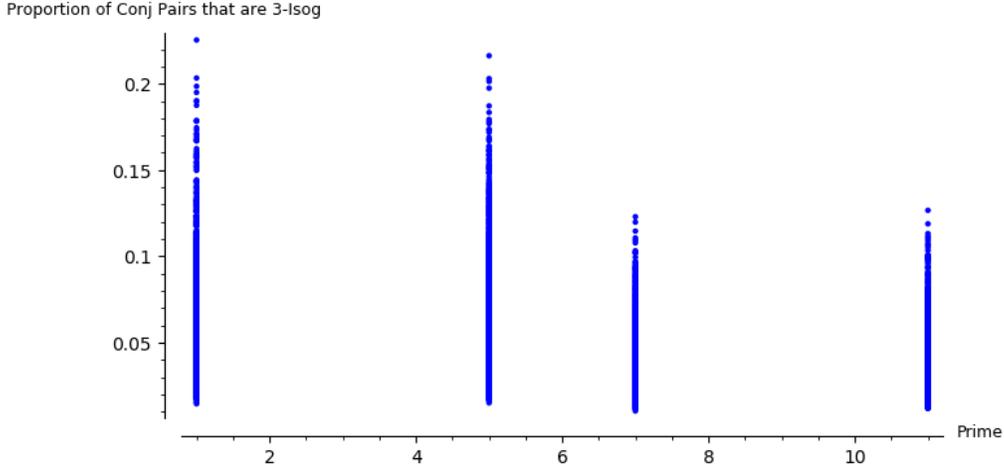


Figure 5.7: Proportions of 3-isogenous conjugates for $10007 \leq p \leq 100193$, sorted by $p \bmod 12$

Question 3: For large p , is there a nontrivial lower and/or upper bound for the proportion of ℓ -isogenous conjugate j -invariants on $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$?

There is a significant difference on the average of the proportion ℓ -isogenous conjugate pairs when we look at the congruence class of modulo 12. We see that this number tends to be smaller when $p \equiv 5, 11 \pmod{12}$ than when $p \equiv 1, 7 \pmod{12}$.

Question 4: How does the proportion of ℓ -isogenous conjugate j -invariants on $\mathcal{G}_\ell(\overline{\mathbb{F}}_p)$ relate to the conjugacy class of $p \bmod 12$?

6 Diameter

Numerical experiments in [Sar19] estimated the diameters of k -regular LPS Ramanujan graphs and random Cayley graphs to be asymptotically $(4/3) \log_{k-1} n$ and $\log_{k-1} n$ respectively, where n is the number of vertices. In this section, we present data on the diameters of the supersingular 2-isogeny graphs, which are 3-regular on approximately $p/12$ vertices (precisely $\lfloor p/12 \rfloor + 0, 1$ or 2 vertices, depending on p).

We can see a lower bound

$$\log_2 \left(\lfloor \frac{p}{12} \rfloor \right) - \log_2(3) + 1$$

on the diameter as follows. Starting from a random vertex and taking a walk of length n , the walk reaches at most $3 \cdot 2^{n-1}$ vertices as endpoints (exactly that number if there are no collisions). Since there are $\lfloor \frac{p}{12} \rfloor + \epsilon$ vertices in the graph, with $\epsilon = 0, 1, 2$, the diameter cannot be less than the smallest n_0 such that

$$3 \cdot 2^{n_0-1} \geq \lfloor \frac{p}{12} \rfloor.$$

This lower bound is shown in green in Figure 6.1 below.

Our numerical data suggests the diameter of the supersingular 2-isogeny graph do *not* grow like $(4/3) \log_2(p/12)$, contrary to the behaviour of LPS graphs. This can be seen from the blue line in Figure 6.1, which has been shifted vertically to fit the data as well as possible, but has too large a slope to match the shape of the distribution. We found $O(\log_2(p/12))$ (the red line in Figure 6.1) to be a better fit, suggesting the 2-isogeny graph might behave more like random Cayley graphs.

We collected graph data for the diameters of the supersingular 2-isogeny graphs $\mathcal{G}_2(\overline{\mathbb{F}}_p)$ for 3387 primes p . We used the built-in Sage "diameter" function ([The19]) on graphs. The implementation

can be found in the walk.sage worksheet available on our github repository. We collected the data in batches, taking snapshots of the possible diameters of $\mathcal{G}_2(\overline{\mathbb{F}}_p)$ for ranges of primes. The smallest prime we have data for is $p = 1009$ and the largest is $p = 4010173$. This data is displayed in figure 6.1.

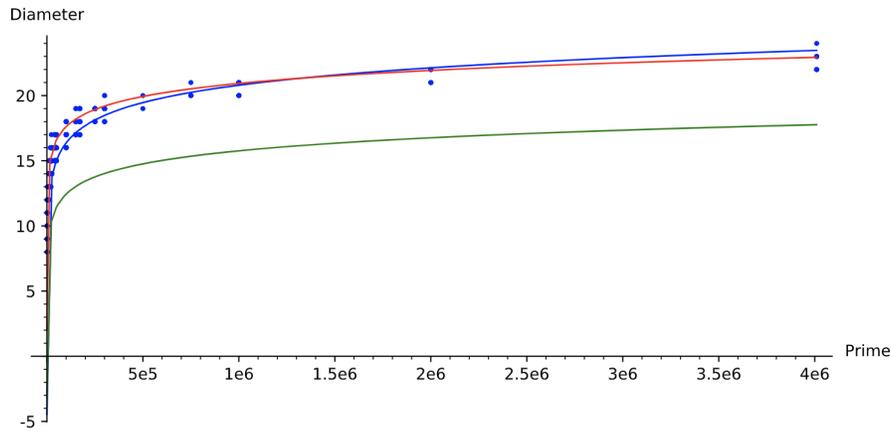


Figure 6.1: Diameters of 2-isogeny graph over $\overline{\mathbb{F}}_p$, with $y = \log_2(p/12) + \log_2(12) + 1$ (red) and $y = \frac{4}{3} \log_2(p/12) - 1$ (blue).

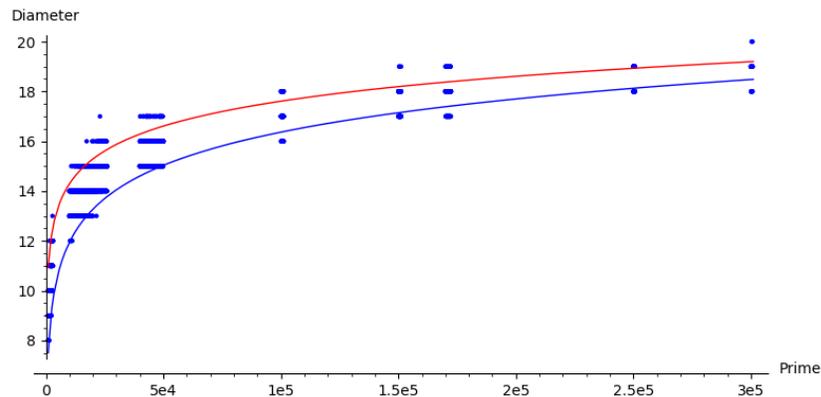


Figure 6.2: Cropped and enlarged graph of Figure 6.1, for the data collected on 3313 primes p with $1009 \leq p \leq 300361$.

6.1 Diameters of Primes Modulo 12

Recall that the number of spinal components and 2-isogenous conjugate pairs is dependent on the congruence class of p modulo 8. Motivated by this, we investigated the behaviour of the diameter as p varies modulo 8. We found a slight, but noticeable, bias for primes congruent to 5 and 11 modulo 12 to have a 2-isogeny graph of larger diameter compared with primes congruent to 1 or 7 modulo 12.

This is visible in Figures 6.3 and 6.4. Notice that in Figure 6.4, the scatter plot points tend to be slightly higher than the graph of $y = \log_2(p/12) + \log_2(12) + 1$, whereas those in Figure 6.3 tend to be more evenly distributed above and below $y = \log_2(p/12) + \log_2(12) + 1$.

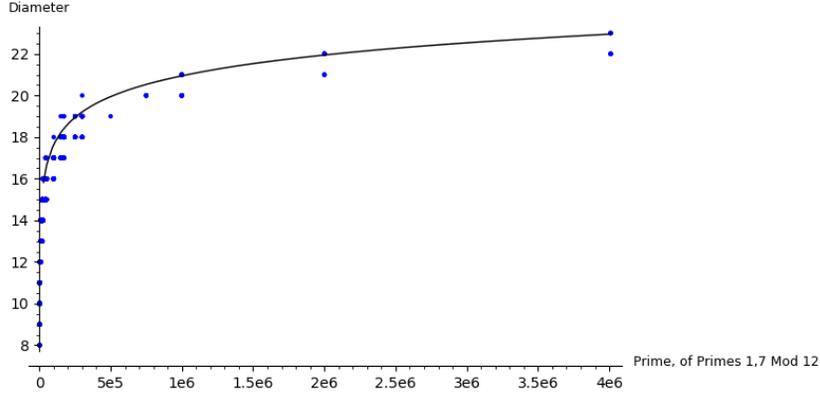


Figure 6.3: Diameters of $\mathcal{G}_2(\overline{\mathbb{F}}_p)$ for $p \equiv 1, 7 \pmod{12}$, with $y = \log_2(p/12) + \log_2(12) + 1$

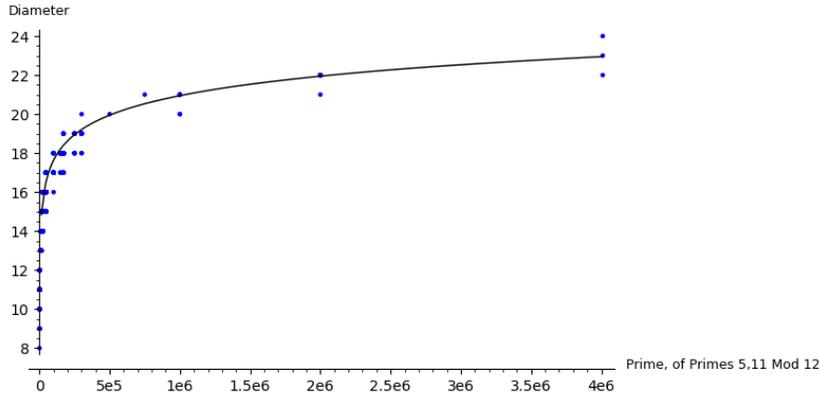


Figure 6.4: Diameters of 2-isogeny graph over $\overline{\mathbb{F}}_p$, for $p \equiv 5, 11 \pmod{12}$, with $y = \log_2(p/12) + \log_2(12) + 1$

Table 3 confirms the visible bias.

average diameter for $100,000 < p < 300,000$			
1 mod 12	17.2190476190476	5 mod 12	17.8761061946903
7 mod 12	17.7346938775510	11 mod 12	17.9919354838710
average diameter for $300,000 < p < 500,000$			
1 mod 12	18.4000000000000	5 mod 12	18.9230769230769
7 mod 12	18.8235294117647	11 mod 12	19.1000000000000

Table 3: Average diameters sorted by primes modulo 12. The first data set contains around 100 primes in each congruence class, the latter between 10 to 17 primes.

7 Conclusions

We determined how the connected components of $\mathcal{G}_\ell(\mathbb{F}_p)$ merge together to give the spine $\mathcal{S} \subset \mathcal{G}_\ell(\overline{\mathbb{F}}_p)$. For any specific ℓ and any p , one can determine the resulting shape explicitly if one knows the structure of the class group $\text{Cl}(\mathcal{O}_K)$.

For $\ell = 2$, we gave heuristics on the distances of the connected components of \mathcal{S} , paths that pass through the spine, the proportion of conjugate pairs, and the diameters of graphs $\mathcal{G}_2(\overline{\mathbb{F}}_p)$.

We saw differences between the congruence classes modulo 12. In summary, the data suggests the following, although more careful analysis is needed to confirm:

- $p \equiv 1, 7 \pmod{12}$:
 - smaller 2-isogeny graph diameters
 - larger number of spinal components
 - larger proportion of 2-isogenous conjugate pairs
- $p \equiv 5, 11 \pmod{12}$:
 - larger 2-isogeny graph diameters
 - smaller number of spinal components
 - smaller proportion of 2-isogenous conjugate pairs

References

- [BJS14] Jean-François Biasse, David Jao, and Anirudh Sankar. *A quantum algorithm for computing isogenies between supersingular elliptic curves*. Springer, 2014.
- [Brö09] Reinier Bröker. Constructing Supersingular Elliptic Curves. *Journal of Combinatorics and Number Theory* 1, 1:269–273, 2009.
- [CFL⁺18] Anamaria Costache, Brooke Feigon, Kristin Lauter, Maike Massierer, and Anna Puskas. Ramanujan graphs in cryptography. Cryptology ePrint Archive, Report 2018/593, 2018. <https://eprint.iacr.org/2018/593>.
- [CGL06] Denis Charles, Eyal Goren, and Kristin Lauter. Cryptographic hash functions from expander graphs. Cryptology ePrint Archive, Report 2006/021, 2006. <https://eprint.iacr.org/2006/021>.
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. Cryptology ePrint Archive, Report 2018/383, 2018. <https://eprint.iacr.org/2018/383>.
- [Cox89] David Cox. *Primes of the form $x^2 + ny^2$* . John Wiley and Sons, Inc., New York, 1989.
- [DG16] C. Delfs and S. D. Galbraith. Computing isogenies between supersingular elliptic curves over \mathbb{F}_p . *Des. Codes Cryptography*, 78(2):425–440, 2016. <https://arxiv.org/pdf/1310.7789.pdf>.
- [EHL⁺18] Kirsten Eisentraeger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. Cryptology ePrint Archive, Report 2018/371, 2018. <https://eprint.iacr.org/2018/371>.
- [FJP11] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. Cryptology ePrint Archive, Report 2011/506, 2011. <https://eprint.iacr.org/2011/506>.
- [GPST16] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. Cryptology ePrint Archive, Report 2016/859, 2016. <https://eprint.iacr.org/2016/859>.
- [Ibu82] Tomoyoshi Ibukiyama. On maximal orders of division quaternion algebras over the rational number field with certain optimal embeddings. *Nagoya Math. J.*, 88:181–195, 1982.

- [Igu58] Jun-Ichi Igusa. Class number of a definite quaternion with prime discriminant. *Proceedings of the National Academy of Sciences of the United States of America*, 44(4):312–314, 1958.
- [Kan89] Masanobu Kaneko. Supersingular j -invariants as singular moduli mod p . *OSAKA JOURNAL OF MATHEMATICS*, 26, 12 1989.
- [KLPT14] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ -isogeny path problem. *LMS J. Comput. Math.*, 17(suppl. A):418–432, 2014.
- [Koh96] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkely, 1996.
- [LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, Sep 1988.
- [Sar19] Naser T. Sardari. Diameter of Ramanujan Graphs and Random Cayley Graphs. *Combinatorica*, 39:427–446, 2019. <https://doi.org/10.1007/s00493-017-3605-0>.
- [Sho99] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev.*, 41(2):303–332, 1999.
- [Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves, 2nd Edition*. Springer-Verlag, New York, N.Y., 2009.
- [Sut13] Andrew 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.
- [The19] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 8.7)*, 2019. <https://www.sagemath.org>.