Higher-degree supersingular group actions

Mathilde Chenu[1,2] and Benjamin Smith[2,1]

[1]Laboratoire d'Informatique (LIX), CNRS, École polytechnique, Institut Polytechnique de Paris, Palaiseau, France
[2]Inria, Palaiseau, France

Received: 1st June 2021 | Revised: 1st August 2021 | Accepted: 1st September 2021

Abstract *We investigate the isogeny graphs of supersingular elliptic curves over $\mathbb{F}_{p^2}$ equipped with a d-isogeny to their Galois conjugate. These curves are interesting because they are, in a sense, a generalization of curves defined over $\mathbb{F}_p$, and there is an action of the ideal class group of $\mathbb{Q}(\sqrt{-dp})$ on the isogeny graphs. We investigate constructive and destructive aspects of these graphs in isogeny-based cryptography, including generalizations of the CSIDH cryptosystem and the Delfs–Galbraith algorithm.*

Keywords: Isogeny-based cryptography, Supersingular elliptic curves, Endomorphisms
2010 Mathematics Subject Classification: 11G20, 14G50

1 INTRODUCTION

Supersingular isogeny graphs of elliptic curves over $\mathbb{F}_{p^2}$, with their rich number-theoretic and combinatorial properties, are at the heart of an increasing number of pre- and post-quantum cryptosystems. Identifying and exploiting special subgraphs of supersingular isogeny graphs is key to understanding their mathematical properties, and their cryptographic potential.

Isogeny-based cryptosystems roughly fall into two families. On one hand, we have the cryptosystems that work in the full ℓ-isogeny graph for various ℓ, including the Charles–Goren–Lauter hash [13], SIDH [30, 23, 17], SIKE [31], OSIDH [16], SQISign [24], and many more. These systems take advantage of the fact that the supersingular ℓ-isogeny graph is a large regular graph with large diameter and excellent expansion and mixing properties (indeed, it is a Ramanujan graph).

On the other hand, we have cryptosystems that work in the $\mathbb{F}_p$-subgraph supported on vertices defined over $\mathbb{F}_p$ (or with j-invariants in $\mathbb{F}_p$), such as CSIDH [11], CSI-FiSh [7], and CSURF [10]. These cryptosystems, many of which represent optimizations and extensions of pioneering work with ordinary curves due to Stolbunov [43, 47] and Couveignes [18], take advantage of the fact that the $\mathbb{F}_p$-endomorphism rings of these curves are an imaginary quadratic ring, and the ideal class group of this ring has a convenient and efficiently-computable commutative action on the $\mathbb{F}_p$-subgraph. This group action allows us to define many simple and useful cryptosystems, but it also explains the structure of the $\mathbb{F}_p$-subgraph, allowing us to use it as a cryptanalytic tool [21] and as a convenient point-of-reference when exploring structures in the full supersingular isogeny graph [1].

This paper investigates a family of generalizations of the $\mathbb{F}_p$-subgraph, one for each squarefree integer d. The key is to recognise that a curve $\mathcal{E}/\mathbb{F}_{p^2}$ has its j-invariant in $\mathbb{F}_p$ precisely when $\mathcal{E}$ is isomorphic to the conjugate curve $\mathcal{E}^{(p)}/\mathbb{F}_{p^2}$ defined by p-th powering the coefficients of $\mathcal{E}$, and the edges of the $\mathbb{F}_p$-subgraph correspond to isogenies that are compatible with these isomorphisms. In this paper, we relax the isomorphisms to d-isogenies, and consider the cryptographic consequences. We obtain a series of distinguished subgraphs of the supersingular isogeny graph, each equipped with a free and transitive action by an ideal class group.

We define (d, ϵ)-*structures*—essentially, curves with a d-isogeny to their conjugate—and the isogenies between them in §2. While (d, ϵ)-structures are defined over $\mathbb{F}_{p^2}$, in §3 we show that they have modular invariants in $\mathbb{F}_p$, and give useful parameterizations for $d = 2$ and 3. We narrow our focus to supersingular curves in §4, using the theory of orientations to set up the class group action on (d, ϵ)-structures. We give some illustrative examples of isogeny graphs of supersingular (d, ϵ)-structures in §5, before turning to cryptographic applications in §6.

Isogeny graphs of (d, ϵ)-structures are a natural setting for variants of CSIDH (and closely related cryptosystems). We give arguments for the security of such cryptosystems in §6.1. We outline a non-interactive key exchange in §6.2, generalizing CSIDH (which is the special case $d = 1$), and highlight some of the subtleties that appear when we move to $d > 1$. Optimized implementation techniques are beyond the scope of this article.

The isogeny graphs formed by (d, ϵ)-structures form interesting geographical features in the full supersingular isogeny graph. Charles, Goren, and Lauter investigated random walks that happen to hit $(d, \pm 1)$-structures in the security analysis of their hash function [13, §7]; random walks into $(\ell, \pm 1)$-structures are also key in the path-finding algorithm of [22]. Further heuristics in this direction appear in [1]. Here, we consider these vertices not in isolation,

*Corresponding Author: mathilde.chenu@inria.fr ; benjamin.smith@inria.fr

but within their own isogeny graphs; thus, we obtain a series of generalizations of the "spine" of [1], and a broad generalization of the Delfs–Galbraith isogeny-finding algorithm [21] in §6.4.

Notation and conventions. If $\mathcal{E}$ is an elliptic curve, then $\mathrm{End}(\mathcal{E})$ denotes its endomorphism ring and $\mathrm{End}^0(\mathcal{E})$ denotes $\mathrm{End}(\mathcal{E}) \otimes \mathbb{Q}$. Each elliptic curve $\mathcal{E}/\mathbb{F}_{p^2}$ has a Galois-conjugate curve $\mathcal{E}^{(p)}$, defined by p-th powering all of the coefficients in the defining equation of $\mathcal{E}$. The curve and its conjugate are connected by inseparable "Frobenius" p-isogenies $\pi_p : \mathcal{E} \to \mathcal{E}^{(p)}$ and $\pi_p : \mathcal{E}^{(p)} \to \mathcal{E}$, defined by p-th powering the coordinates (abusing notation, all inseparable p-isogenies will be denoted by π_p). Observe that $(\mathcal{E}^{(p)})^{(p)} = \mathcal{E}$, and the composition of $\pi_p : \mathcal{E} \to \mathcal{E}^{(p)}$ and $\pi_p : \mathcal{E}^{(p)} \to \mathcal{E}$ is the p^2-power Frobenius endomorphism $\pi_{\mathcal{E}}$ of $\mathcal{E}$. Conjugation also operates on isogenies: each isogeny $\phi : \mathcal{E} \to \mathcal{E}'$ defined over $\mathbb{F}_{p^2}$ has a Galois conjugate isogeny $\phi^{(p)} : \mathcal{E}^{(p)} \to \mathcal{E}'^{(p)}$, defined by p-th powering all of the coefficients in a rational map defining ϕ. We always have

$$(\phi^{(p)})^{(p)} = \phi \qquad \text{and} \qquad \pi_p \circ \phi = \phi^{(p)} \circ \pi_p \,,$$

and conjugation thus gives an isomorphism of rings between $\mathrm{End}(\mathcal{E})$ and $\mathrm{End}(\mathcal{E}^{(p)})$.

2 CURVES WITH A d-ISOGENY TO THEIR CONJUGATE

Let $p > 3$ be a prime, and d a squarefree integer prime to p. Typically, p is very large and d is very small.

We are interested in elliptic curves $\mathcal{E}/\mathbb{F}_{p^2}$ equipped with a d-isogeny $\psi : \mathcal{E} \to \mathcal{E}^{(p)}$. Given any such d-isogeny ψ, we have two returning d-isogenies:

$$\psi^{(p)} : \mathcal{E}^{(p)} \to \mathcal{E} \qquad \text{and} \qquad \widehat{\psi} : \mathcal{E}^{(p)} \to \mathcal{E} \,.$$

Definition 1. *Let $\mathcal{E}/\mathbb{F}_{p^2}$ be an elliptic curve equipped with a d-isogeny $\psi : \mathcal{E} \to \mathcal{E}^{(p)}$ to its conjugate. We say that $(\mathcal{E}, \psi)$ is a* (d, ϵ)-structure *if*

$$\widehat{\psi} = \epsilon \psi^{(p)} \quad \textit{with} \quad \epsilon \in \{1, -1\} \,.$$

Each (d, ϵ)-structure $(\mathcal{E}, \psi)$ has an associated endomorphism

$$\mu := \pi_p \circ \psi \in \mathrm{End}(\mathcal{E}) \,.$$

We say that $(\mathcal{E}, \psi)$ is ordinary resp. supersingular if $\mathcal{E}$ is ordinary resp. supersingular.[1]

Proposition 1. *If $(\mathcal{E}, \psi)$ is a (d, ϵ)-structure and μ is its associated endomorphism, then*

$$\mu^2 = [\epsilon d] \pi_{\mathcal{E}} \,.$$

If $\pi_{\mathcal{E}}$ is the Frobenius endomorphism of $\mathcal{E}$ and $t_{\mathcal{E}}$ is its trace, then there exists an integer r such that $[r]\mu = [p] + \epsilon \pi_{\mathcal{E}}$ in $\mathrm{End}(\mathcal{E})$, $dr^2 = 2p + \epsilon t_{\mathcal{E}}$ in $\mathbb{Z}$, and the characteristic polynomial of μ is $P_\mu(T) = T^2 - rdT + dp$.

Proof. We have $\psi \pi_p = \pi_p \psi^{(p)}$, so $\mu^2 = \pi_p \psi \pi_p \psi = \pi_p (\pi_p \psi^{(p)}) \psi = \pi_{\mathcal{E}} (\psi^{(p)} \psi)$. Now $\psi^{(p)} = \epsilon \widehat{\psi}$ (because $(\mathcal{E}, \psi)$ is a (d, ϵ)-structure), so $\psi^{(p)} \psi = [\epsilon d]$, and therefore $\mu^2 = [\epsilon d] \pi_{\mathcal{E}}$. For the rest: μ has degree dp, so it satisfies a quadratic polynomial $P_\mu(T) = T^2 - aT + dp$ for some integer a. The first assertion then implies $[a]\mu = \mu^2 + [dp] = [\epsilon d] \pi_{\mathcal{E}} + [dp]$. Squaring, we obtain

$$([a]\mu)^2 = [d]^2 (\pi_{\mathcal{E}}^2 + p^2) + 2[dp][\epsilon d] \pi_{\mathcal{E}} = [d]^2 (t_{\mathcal{E}} \pi_{\mathcal{E}}) + 2[dp][\epsilon d] \pi_{\mathcal{E}} = [\epsilon d] \pi_{\mathcal{E}} ([\epsilon d] t_{\mathcal{E}} + 2dp) \,,$$

so $a^2 = \epsilon d t_{\mathcal{E}} + 2dp$, hence $d \mid a^2$. But d is squarefree, so $d \mid a$, and then $r = a/d$ satisfies the given conditions. □

Remark 1. *In the situation of Proposition 1: if $\mathcal{E}$ is ordinary, then $\mathbb{Z}[\mu]$ and $\mathbb{Z}[\pi_{\mathcal{E}}]$ are orders in $\mathbb{Q}(\pi_{\mathcal{E}})$ of discriminant $d^2 r^2 - 4dp$ and $t_{\mathcal{E}}^2 - 4p^2 = r^2 (d^2 r^2 - 4dp)$, respectively, so $|r|$ is the conductor of $\mathbb{Z}[\pi_{\mathcal{E}}]$ in $\mathbb{Z}[\mu]$. (The supersingular case is treated in detail in §4.)*

Definition 2. *Let $(\mathcal{E}, \psi)$ and $(\mathcal{E}', \psi')$ be (d, ϵ)-structures. We say an isogeny (resp. isomorphism) $\phi : \mathcal{E} \to \mathcal{E}'$ is an* isogeny (resp. isomorphism) of (d, ϵ)-structures *if $\psi' \phi = \phi^{(p)} \psi$, that is, if the following diagram commutes:*

$$\begin{array}{ccc} \mathcal{E} & \xrightarrow{\psi} & \mathcal{E}^{(p)} \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi^{(p)}} \\ \mathcal{E}' & \xrightarrow{\psi'} & (\mathcal{E}')^{(p)} \end{array}$$

[1] We focus on curves defined over $\mathbb{F}_{p^2}$ because our applications involve supersingular curves, and every supersingular curve is isomorphic to a curve over $\mathbb{F}_{p^2}$. One might consider isogenies to conjugates over higher-degree extensions, but then in general we do not have the relation $\widehat{\psi} = \pm \psi^{(p)}$, which is fundamental to our results.

It is easily verified that isogenies of (d, ϵ)-structures follow the usual rules obeyed by isogenies: the composition of two isogenies of (d, ϵ)-structures is an isogeny of (d, ϵ)-structures, the dual of an isogeny of (d, ϵ)-structures is an isogeny of (d, ϵ)-structures, and every (d, ϵ)-structure has an isogeny to itself (the identity map, for example). Isogeny therefore forms an equivalence relation on (d, ϵ)-structures.

If $(\mathcal{E}, \psi)$ is a (d, ϵ)-structure with associated endomorphism μ, then

$$-(\mathcal{E}, \psi) := (\mathcal{E}, -\psi) \qquad \text{and} \qquad (\mathcal{E}, \psi)^{(p)} := (\mathcal{E}^{(p)}, \psi^{(p)})$$

are (d, ϵ)-structures with associated endomorphisms $-\mu$ and $\mu^{(p)}$, respectively. If $\phi : (\mathcal{E}, \psi) \to (\mathcal{E}', \psi')$ is an isogeny of (d, ϵ)-structures, then $\phi : -(\mathcal{E}, \psi) \to -(\mathcal{E}', \psi')$ and $\phi^{(p)} : (\mathcal{E}, \psi)^{(p)} \to (\mathcal{E}', \psi')^{(p)}$ are also isogenies of (d, ϵ)-structures. We thus have two involutions, **negation** and **conjugation**, on the category of (d, ϵ)-structures and their isogenies.

Remark 2. *The isogenies ψ and $\pi_p : \mathcal{E} \to \mathcal{E}^{(p)}$ are both in fact isogenies of (d, ϵ)-structures $(\mathcal{E}, \psi) \to (\mathcal{E}, \psi)^{(p)}$.*

Twisting. Let α be an element of $\overline{\mathbb{F}}_p \setminus \{0\}$. For each elliptic curve $\mathcal{E} : y^2 = x^3 + ax + b$, there is a curve

$$\mathcal{E}^\alpha / \mathbb{F}_{p^2}(\alpha^2) : y^2 = x^3 + \alpha^4 ax + \alpha^6 b$$

and an $\mathbb{F}_{p^2}(\alpha)$-isomorphism $\tau_\alpha : \mathcal{E} \to \mathcal{E}^\alpha$ defined by $(x, y) \mapsto (\alpha^2 x, \alpha^3 y)$. Abusing notation, we write τ_α for this map on *every* elliptic curve; with this convention, $\tau_\beta \circ \tau_\alpha = \tau_{\alpha\beta}$. If δ is a nonsquare in $\mathbb{F}_{p^2}$ then $\mathcal{E}^{\sqrt{\delta}}$ is the *quadratic twist* (which, up to $\mathbb{F}_{p^2}$-isomorphism, is independent of the choice of nonsquare δ) and $\tau_{\sqrt{\delta}}$ is the twisting isomorphism. For each isogeny $\phi : \mathcal{E} \to \mathcal{E}'$ defined over $\mathbb{F}_{p^2}$, there is an $\mathbb{F}_{p^2}(\alpha^2)$-isogeny

$$\phi^\alpha := (\tau_\alpha \circ \phi \circ \tau_{1/\alpha}) : \mathcal{E}^\alpha \longrightarrow (\mathcal{E}')^\alpha .$$

Now let $(\mathcal{E}, \psi)$ be a (d, ϵ)-structure with associated endomorphism μ. If again we choose a nonsquare δ in $\mathbb{F}_{p^2}$, and a square root $\sqrt{\delta}$ of δ in $\mathbb{F}_{p^4}$, then in general $(\mathcal{E}^{\sqrt{\delta}}, \psi^{\sqrt{\delta}})$ is *not* a $(d, \pm 1)$-structure (because conjugation and twisting generally do not commute); but $(\mathcal{E}, \psi)^{\sqrt{\delta}} := (\mathcal{E}^{\sqrt{\delta}}, \tau_{(\sqrt{\delta})^{(p-1)}} \circ \psi^{\sqrt{\delta}})$ is a $(d, -\epsilon)$-structure with associated endomorphism $\mu^{\sqrt{\delta}}$. The $\mathbb{F}_{p^2}$-isomorphism class of $(\mathcal{E}, \psi)^{\sqrt{\delta}}$ is independent of the choice of δ; we call $(\mathcal{E}, \psi)^{\sqrt{\delta}}$ the *quadratic twist* of $(\mathcal{E}, \psi)$. Note that $((\mathcal{E}, \psi)^{\sqrt{\delta}})^{\sqrt{\delta}} \cong (\mathcal{E}, \psi)$. If $\phi : (\mathcal{E}, \psi) \to (\mathcal{E}', \psi')$ is an isogeny of (d, ϵ)-structures, then $\phi^{\sqrt{\delta}}$ induces an isogeny of $(d, -\epsilon)$-structures $\phi^{\sqrt{\delta}} : (\mathcal{E}, \psi)^{\sqrt{\delta}} \to (\mathcal{E}', \psi')^{\sqrt{\delta}}$. Twisting therefore takes us from the category of (d, ϵ)-structures into the category of $(d, -\epsilon)$-structures and back again.

Example 1. *Consider the case $d = 1$. Each $(1, 1)$-structure is $\mathbb{F}_{p^2}$-isomorphic to the base-extension to $\mathbb{F}_{p^2}$ of a curve defined over $\mathbb{F}_p$ (with the 1-isogeny being $[\pm 1]$); the associated endomorphism is the p-power Frobenius endomorphism on the base-extended curve, and the integer r of Proposition 1 is the trace of the p-power Frobenius. Each $(1, -1)$-structure is the quadratic twist of a $(1, 1)$-structure: essentially, an ordinary $(1, -1)$-structure is isomorphic to a GLS curve [26]. This discussion should be compared with the remark at the end of [45, §3].*

3 PARAMETRIZATIONS AND MODULAR CURVES

For our computations, we can represent a (d, ϵ)-structure $(\mathcal{E}, \psi)$ as $(\mathcal{E}, f_\psi, \alpha)$, where f_ψ is the kernel polynomial of ψ (that is, the monic polynomial whose roots are the x-coordinates of the nonzero points in $\ker \psi$) and α is the element such that $\psi = \tau_\alpha \circ \tilde{\psi}$, where $\tilde{\psi} : \mathcal{E} \to \mathcal{E}/\ker \psi$ is the normalized "Vélu" isogeny.

We want a more space-efficient encoding of isomorphism classes of (d, ϵ)-structures, both as a canonical encoding for vertices in isogeny graphs, and for transmission of (d, ϵ)-structures used as cryptographic values.

A full development of these representations and the algorithms that operate on them is beyond the scope of this short article, but we sketch a simple compression scheme in §3.1 and §3.2, and very useful explicit parametrizations for $d = 2$ and 3 in §3.3 and §3.4, respectively. The latter were derived from explicit parametrizations of $\mathbb{Q}$-curves due to Hasegawa [29]. The associated endomorphisms for ordinary curves in these families have been used to accelerate scalar multiplication algorithms (see [45], where we also find related families for $d = 5$ and 7, and [28]) and as inputs for specialized point-counting algorithms [37].

3.1 A SIMPLE AND GENERAL COMPRESSION SCHEME

While (d, ϵ)-structures may seem to be relatively complicated objects over $\mathbb{F}_{p^2}$, their isomorphism classes can be encoded to little more than a single element of $\mathbb{F}_p$. Briefly: the key is to take the quotient by negation, which maps the set $S_{d,\epsilon}$ of isomorphism classes of (d, ϵ)-structures over $\mathbb{F}_{p^2}$ into $X_0(d)(\mathbb{F}_{p^2})$, where $X_0(d)$ is

the level-d modular curve. Then, the Atkin–Lehner involution ω_d, which maps a modular point onto its "dual", acts as conjugation on the image of $S_{d,\epsilon}$. Writing $X_0^+(d) = X_0(d)/\langle\omega_d\rangle$, we have a four-to-one map from $S_{d,\epsilon}$ onto $X_0^+(d)(\mathbb{F}_p)$, identifying the isomorphism class of $(\mathcal{E},\psi)$ with $-(\mathcal{E},\psi)$, $(\mathcal{E},\psi)^{(p)}$, and $-(\mathcal{E},\psi)^{(p)}$. We can therefore represent an element of $S_{d,\epsilon}$ as a point in $X_0^+(d)(\mathbb{F}_p)$ plus two bits (one to determine the sign, the other the conjugate). Since $X_0^+(d)$ is a curve, we can further compress the representative point in $X_0^+(d)(\mathbb{F}_p)$ to one element of $\mathbb{F}_p$ plus a few bits. This step depends strongly on the geometry of $X_0^+(d)(\mathbb{F}_p)$: for example, if $X_0^+(d)$ has genus 0 then we can rationally parametrize it, giving a simple compression of points in $X_0^+(d)(\mathbb{F}_p)$ to single elements of $\mathbb{F}_p$; if $X_0^+(d)$ is hyperelliptic, then we can compress points in $X_0^+(d)(\mathbb{F}_p)$ to a single element of $\mathbb{F}_p$ plus a "sign" bit in the usual way; and as the gonality of $X_0^+(d)$ increases, so does the number of auxiliary bits required.

3.2 EXAMPLE: COMPRESSING $(5,\epsilon)$-STRUCTURES

To illustrate the technique above in more detail, suppose we want to compress $(5,\epsilon)$-structures over $\mathbb{F}_{p^2}$ to elements of $\mathbb{F}_p$. The classical modular polynomial of level 5 is a polynomial $\Phi_5(J_0,J_1)$ with integer coefficients, of degree 6 in J_0 and J_1. It is symmetric in J_0 and J_1, so we can write

$$\Phi_5(J_0,J_1) = -F_5(J_0+J_1, J_0J_1)$$

where

$$\begin{aligned} F_5(T,N) = {} & N^5 + 40(93T + 41650211662)N^4 \\ & + 36(126415T^2 - 2996636724991200T + 12277031464804661791632)N^3 \\ & + \cdots \end{aligned}$$

is an integer polynomial of degree 6 in T and 5 in N.

In terms of modular curves: Φ_5 defines an affine model of $X_0(5)$, and the Atkin–Lehner involution on $X_0(5)$ exchanges the variables J_0 and J_1 in this model, so F_5 defines an affine model of $X_0^+(5)$, with the quotient map $X_0(5) \to X_0^+(5)$ defined by $(J_0,J_1) \mapsto (T,N) = (J_0+J_1, J_0J_1)$.

Now suppose we are given a $(5,\epsilon)$-structure $(\mathcal{E},\psi)$ over $\mathbb{F}_{p^2}$; we want to compress $(\mathcal{E},\psi)$ down to a single element of $\mathbb{F}_p$ plus a few bits. For simplicity, we will assume that $\mathcal{E}$ has no extra automorphisms.

First, there is an element γ of $\mathbb{F}_{p^2}$ such that $\psi^*(\omega_{\mathcal{E}^{(p)}}) = \gamma\omega_{\mathcal{E}}$, where $\omega_{\mathcal{E}}$ and $\omega_{\mathcal{E}^{(p)}}$ are the invariant differentials on $\mathcal{E}$ and $\mathcal{E}^{(p)}$, respectively. Fixing a sign function on $\mathbb{F}_{p^2}$, we can encode the sign of the isogeny ψ as a bit ϵ_1 determining the sign of γ. Now $(\mathcal{E},\psi)$ is determined by $(\mathcal{E},\ker\psi)$ and ϵ_1.

The pair $(\mathcal{E},\ker\psi)$ corresponds to the point $(j(\mathcal{E}), j(\mathcal{E}^{(p)})) = (j(\mathcal{E}), j(\mathcal{E})^p)$ on $X_0(5)$. Set $t = j(\mathcal{E}) + j(\mathcal{E})^p$ and $n = j(\mathcal{E})j(\mathcal{E})^p$, both in $\mathbb{F}_p$, and let ϵ_2 be a bit determining $j(\mathcal{E})$ as one of the roots in $\mathbb{F}_{p^2}$ of the quadratic $X^2 - tX + n$; then $(\mathcal{E},\psi)$ corresponds to $(\epsilon_1,\epsilon_2,t,n)$. Now let $1 \le i \le 5$ determine the position of n (in lexicographic order, say) among the (at most) 5 roots in $\mathbb{F}_p$ of the quintic $F_5(t,X)$; then $(\mathcal{E},\psi)$ corresponds to $(\epsilon_1,\epsilon_2,i,t)$.

Working in the other direction: given $(\epsilon_1,\epsilon_2,i,t)$, we compute the roots of $F_5(t,X)$ in $\mathbb{F}_p$, sort them, and let n be the i-th one; then we use ϵ_2 to choose a root α of $X^2 - tX + n$; then we construct a curve $\tilde{\mathcal{E}}$ with $j(\tilde{\mathcal{E}}) = \alpha$, and recover a 5-isogeny $\tilde{\psi} : \tilde{\mathcal{E}} \to \tilde{\mathcal{E}}^{(p)}$ using Elkies' algorithm, for example (see [44, §7 and §8]). We use ϵ_1 to correct the sign of $\tilde{\psi}$ if required by looking at the action on invariant differentials.

The encoding $(\mathcal{E},\psi) \mapsto (\epsilon_1,\epsilon_2,i,t)$ requires $\lceil \log_2(p) \rceil + 5$ bits, since $1 \le i \le 5$ can be encoded in three bits. We see that for general d, the number of extra bits depends (logarithmically) on the gonality of the modular curve $X_0^+(d)$ (i.e., its degree in N above). Using alternate models of modular curves can reduce this to some extent.

Indeed, in the case $d = 5$, the encoding above is not optimal: the curve $X_0(5)$ actually has genus 0 and can be rationally parametrized over $\mathbb{F}_{p^2}$. This lets us get down to a single element of $\mathbb{F}_p$ plus a choice of sign, as in [45, §5]; we will see examples of this for $d = 2$ in §3.3 and §3.4 below.

3.3 EXAMPLE: COMPRESSING $(2,\epsilon)$-STRUCTURES

Let Δ be a nonsquare in $\mathbb{F}_p$, and fix a square root $\sqrt{\Delta}$ in $\mathbb{F}_{p^2}$. For each u in $\mathbb{F}_p$, the curve

$$\mathcal{E}_{2,u}/\mathbb{F}_{p^2} : y^2 = x^3 - 6(5 - 3u\sqrt{\Delta})x + 8(7 - 9u\sqrt{\Delta})$$

has a rational 2-torsion point $(4,0)$, which generates the kernel of a 2-isogeny $\psi_{2,u} : \mathcal{E}_{2,u} \to {\mathcal{E}_{2,u}}^{(p)}$ defined over $\mathbb{F}_{p^2}$. The j-invariant of $\mathcal{E}_{2,u}$ is

$$j(\mathcal{E}_{2,u}) = \frac{(20 - 12u\sqrt{\Delta})^3}{(1 - u\sqrt{\Delta})(1 + u\sqrt{\Delta})^2}\,.$$

Note that when we put $u = \infty$, we get $j(\mathcal{E}_{2,\infty}) = 12^3 = 1728$.

If we use Vélu's formulae to compute the (normalized) quotient isogeny $\mathcal{E}_{2,u} \to \mathcal{E}_{2,u}/\langle(4,0)\rangle$, then the isomorphism $\mathcal{E}_{2,u}/\langle(4,0)\rangle \to {\mathcal{E}_{2,u}}^{(p)}$ is $\tau_{1/\sqrt{-2}}$. Composing, we obtain an expression for $\psi_{2,u}$ as a rational map:

$$\psi_{2,u} : (x,y) \longmapsto \left(\frac{-x}{2} - \frac{9(1+u\sqrt{\Delta})}{x-4}, \frac{y}{\sqrt{-2}}\left(\frac{-1}{2} + \frac{9(1+u\sqrt{\Delta})}{(x-4)^2}\right)\right).$$

Computing the dual isogeny $\widehat{\psi}_{2,u}$ and comparing it with ${\psi_{2,u}}^{(p)}$, we find that $(\mathcal{E}_{2,u}, \psi_{2,u})$ is a $(2,\epsilon)$-structure with $\epsilon = 1$ if $p \equiv 5,7 \pmod 8$, or -1-structure if $p \equiv 1,3 \pmod 8$. (To obtain a family of $(2,-1)$-structures when $p \equiv 5,7 \pmod 8$ or $(2,1)$-structures if $p \equiv 1,3 \pmod 8$, it suffices to take the quadratic twist.) Notice that $(\mathcal{E}_{2,u}, \psi_{2,u})^{(p)} = (\mathcal{E}_{2,-u}, \psi_{2,-u})$: that is, negating the parameter acts as conjugation on the $(2,\epsilon)$-structure.

Every $(2,\epsilon)$-structure is isomorphic to $(\mathcal{E}_{2,u}, \psi_{2,u})$ for some u in $\mathbb{F}_p \cup \{\infty\}$, up to sign (or composition with an automorphism), after a possible twist (as above) according to the congruence class of p modulo 8. The parameter u, and the isogenies $\psi_{2,u}$, therefore give us an extremely efficient compression scheme for $(2,\epsilon)$-structures.

3.4 EXAMPLE: COMPRESSING $(3,\epsilon)$-STRUCTURES

Let Δ be a nonsquare in $\mathbb{F}_p$, and fix a square root $\sqrt{\Delta}$ in $\mathbb{F}_{p^2}$. For each u in $\mathbb{F}_p$, the elliptic curve

$$\mathcal{E}_{3,u}/\mathbb{F}_{p^2} : y^2 = x^3 - 3\big(5 + 4u\sqrt{\Delta}\big)x + 2\big(2u^2\Delta + 14u\sqrt{\Delta} + 11\big)$$

has an order-3 subgroup $\{O, (3, \pm 2(1-u\sqrt{\Delta}))\}$ defined by the polynomial $x - 3$. The j-invariant of this family is

$$j(\mathcal{E}_{3,u}) = \frac{2(30+24u\sqrt{\Delta})^3}{(1-u\sqrt{\Delta})^3(1+u\sqrt{\Delta})}.$$

Putting $u = \infty$, we get $j(\mathcal{E}_{3,\infty} = 0)$.

Taking the quotient with Vélu's formulae and composing with $\tau_{1/\sqrt{-3}}$ yields an explicit 3-isogeny $\psi_{3,u} : \mathcal{E}_{3,u} \to {\mathcal{E}_{3,u}}^{(p)}$, and we find that $(\mathcal{E}_{3,u}, \psi_{3,u})$ is a $(3,1)$-structure if $p \equiv 2 \pmod 3$, or a $(3,-1)$-structure if $p \equiv 1 \pmod 3$. (To obtain a family of $(3,-1)$-structures when $p \equiv 2 \pmod 3$ or $(3,1)$-structures if $p \equiv 1 \pmod 3$, take the quadratic twist.)

Every $(3,\epsilon)$-structure is isomorphic to $(\mathcal{E}_{3,u}, \psi_{3,u})$ for some u in $\mathbb{F}_p \cup \{\infty\}$, up to sign (or composition with an automorphism), after a possible twist (as above) according to the congruence class of p modulo 3. The parameter u, and the isogenies $\psi_{3,u}$, therefore give us a convenient compression scheme for $(3,\epsilon)$-structures, which we will use in our examples in §5.

4 SUPERSINGULAR (d,ϵ)-STRUCTURES

We now come to the main focus of our investigation: supersingular (d,ϵ)-structures and their isogeny graphs.

Definition 3. *We write $\mathcal{D}_{d,\epsilon}(p)$ for the set of supersingular (d,ϵ)-structures over $\mathbb{F}_{p^2}$ up to $\mathbb{F}_{p^2}$-isomorphism, and $\Gamma(\mathcal{D}_{d,\epsilon}(p))$ for the graph on $\mathcal{D}_{d,\epsilon}(p)$ whose edges are ($\mathbb{F}_{p^2}$-isomorphism classes of) isogenies of (d,ϵ)-structures. For each prime $\ell \neq p$, we write $\Gamma_\ell(\mathcal{D}_{d,\epsilon}(p))$ for the subgraph of $\Gamma(\mathcal{D}_{d,\epsilon}(p))$ where the edges are ℓ-isogenies.*

Observe that the quadratic twist gives an isomorphism of graphs $\Gamma(\mathcal{D}_{d,\epsilon}(p)) \cong \Gamma(\mathcal{D}_{d,-\epsilon}(p))$.

Proposition 2. *Let $(\mathcal{E},\psi)$ be a (d,ϵ)-structure with associated endomorphism μ. If $\mathcal{E}$ is supersingular, then*

1. *$\mu^2 = [-dp]$.*
2. *The trace of Frobenius satisfies $t_\mathcal{E} = -2\epsilon p$, and in particular $\mathcal{E}(\mathbb{F}_{p^2}) \cong (\mathbb{Z}/(p+\epsilon)\mathbb{Z})^2$.*

Proof. With the notation of Proposition 1: The curve $\mathcal{E}$ is supersingular if and only if $p \mid t_\mathcal{E}$. Now $p \nmid d$, so $p \mid r$ by Proposition 1. The characteristic polynomial $P_\mu(T)$ of μ has discriminant $(rd)^2 - 4dp$; this discriminant cannot be positive, so $|r| \leq 2\sqrt{p/d}$. Since $p \mid r$, we have $r = 0$, so $\mu^2 = [-dp]$, and $t_\mathcal{E} = \frac{-2p}{\epsilon} = -2\epsilon p$. □

Proposition 2 tells us that if $(\mathcal{E},\psi)$ is a supersingular (d,ϵ)-structure, then ϵ is completely determined by the $\mathbb{F}_{p^2}$-isogeny class of $\mathcal{E}$. Further, $t_\mathcal{E}$ can only be $\pm 2p$: the special supersingular traces $-p$, 0, and p (corresponding to non-quadratic twists of curves of j-invariant 0 and 1728, if these are supersingular) cannot occur.

4.1 ORIENTATIONS

Proposition 2 tells us that the associated endomorphism of each supersingular (d, ϵ)-structure acts like a square root of $-dp$ in the endomorphism ring. We can make this notion more precise using *orientations*, as described by Colò and Kohel in [16] and Onuki in [38]. Before going further, we recall some generalities.

Let K be an imaginary quadratic field, $\mathcal{O}_K$ its ring of integers, and $\mathcal{O}$ an order in K. A K-orientation on an elliptic curve $\mathcal{E}/\mathbb{F}_{p^2}$ is a homomorphism $\iota : K \to \mathrm{End}^0(\mathcal{E})$; we call the pair $(\mathcal{E}, \iota)$ a *K-oriented elliptic curve*. We say ι is an $\mathcal{O}$-orientation, and $(\mathcal{E}, \iota)$ is an $\mathcal{O}$-oriented elliptic curve, if $\iota(\mathcal{O}) \subseteq \mathrm{End}(\mathcal{E})$. An $\mathcal{O}$-orientation $\iota : K \to \mathrm{End}^0(\mathcal{E})$ is *primitive* if $\iota(\mathcal{O}) = \mathrm{End}(\mathcal{E}) \cap \iota(K)$: that is, if ι is "full" in the sense that it does not extend to an $\mathcal{O}'$-orientation for any strict super-order $\mathcal{O}' \supset \mathcal{O}$.

Let $(\mathcal{E}, \iota)$ be a K-oriented elliptic curve. If $\phi : \mathcal{E} \to \mathcal{E}'$ is an isogeny, then there is an *induced K-orientation* $\phi_*(\iota)$ on $\mathcal{E}'$ defined by

$$\phi_*(\iota) : \alpha \longmapsto \frac{1}{\deg(\phi)} \phi \circ \iota(\alpha) \circ \widehat{\phi}\,.$$

Given two oriented curves $(\mathcal{E}, \iota)$ and $(\mathcal{E}', \iota')$, an isogeny $\phi : \mathcal{E} \to \mathcal{E}'$ is said to be K-oriented, or an isogeny of K-oriented elliptic curves, if $\iota' = \phi_*(\iota)$. In this case we write $\phi : (\mathcal{E}, \iota) \to (\mathcal{E}', \iota')$. If there exists a K-oriented isogeny $\tilde{\phi} : (\mathcal{E}', \iota') \to (\mathcal{E}, \iota)$ such that $\tilde{\phi} \circ \phi = [1]_{\mathcal{E}}$ and $\phi \circ \tilde{\phi} = [1]_{\mathcal{E}'}$, then we say that ϕ is a K-oriented isomorphism, and we write $(\mathcal{E}, \iota) \cong (\mathcal{E}', \iota')$. Note that $\phi : (\mathcal{E}, \iota) \to (\mathcal{E}', \iota')$ is an oriented isomorphism if and only if the underlying isomorphism of curves ϕ satisfies $\phi \circ \iota(\alpha) = \iota'(\alpha) \circ \phi$ for all α in K.

If $\phi : (\mathcal{E}, \iota) \to (\mathcal{E}', \iota')$ is a K-oriented isogeny, then ι resp. ι' is a primitive $\mathcal{O}$ resp. $\mathcal{O}'$-orientation for some order $\mathcal{O}$ resp. $\mathcal{O}'$ in K. If $\ell = \deg \phi$ is a prime not equal to p, then one of the following holds:

- $\mathcal{O} = \mathcal{O}'$, and ϕ is said to be *horizontal*; or
- $\mathcal{O} \subset \mathcal{O}'$ with $[\mathcal{O}' : \mathcal{O}] = \ell$, and ϕ is said to be *ascending*; or
- $\mathcal{O} \supset \mathcal{O}'$ with $[\mathcal{O} : \mathcal{O}'] = \ell$, and ϕ is said to be *descending*.

Let $\mathcal{O}$ be an order in a quadratic field K such that p does not split in K or divide the conductor of $\mathcal{O}$. Following [16], we let $SS_{\mathcal{O}}(p)$ denote the set of $\mathcal{O}$-oriented supersingular elliptic curves over $\overline{\mathbb{F}}_p$ up to K-oriented isomorphism. The subset of primitive $\mathcal{O}$-oriented curves (up to K-oriented isomorphism) is denoted by $SS^{\mathrm{pr}}_{\mathcal{O}}(p)$.

For any integral invertible ideal $\mathfrak{a}$ in $\mathcal{O}$ and any $\mathcal{O}$-oriented curve $(\mathcal{E}, \iota)$, we have a finite subgroup

$$\mathcal{E}[\mathfrak{a}] := \{P \in \mathcal{E} \mid \iota(\alpha)(P) = 0 \quad \forall \alpha \in \mathfrak{a}\}\,.$$

Now suppose $\mathfrak{a}$ is prime to the conductor of $\mathcal{O}$ in $\mathcal{O}_K$.[2] If $\phi_{\mathfrak{a}} : \mathcal{E} \to \mathcal{E}/\mathcal{E}[\mathfrak{a}]$ is the quotient isogeny, then $(\phi_{\mathfrak{a}})_*(\iota)$ is an $\mathcal{O}$-orientation on $\mathcal{E}/\mathcal{E}[\mathfrak{a}]$, and $\phi_{\mathfrak{a}}$ is a horizontal isogeny of $\mathcal{O}$-oriented curves. If $\mathfrak{a}$ is principal then $(\mathcal{E}/\mathcal{E}[\mathfrak{a}], (\phi_{\mathfrak{a}})_*(\iota)) \cong (\mathcal{E}, \iota)$, so the map

$$(\mathfrak{a}, (\mathcal{E}, \iota)) \mapsto (\mathcal{E}/\mathcal{E}[\mathfrak{a}], (\phi_{\mathfrak{a}})_*(\iota))$$

extends to fractional ideals and factors through the class group, and as in [16] we get a transitive group action

$$\mathrm{Cl}(\mathcal{O}) \times SS_{\mathcal{O}}(p) \longrightarrow SS_{\mathcal{O}}(p)\,.$$

Onuki [38] shows that if we restrict to a certain subset of the *primitive* $\mathcal{O}$-oriented curves, then this action is transitive and free. Let $\mathcal{J}_{\mathcal{O}}$ denote the set of j-invariants of elliptic curves $\mathcal{E}$ over $\mathbb{C}$ (not $\overline{\mathbb{F}}_p$) with $\mathrm{End}(\mathcal{E}) \cong \mathcal{O}$. All elements in $\mathcal{J}_{\mathcal{O}}$ are algebraic integers, so an elliptic curve whose j-invariant is in $\mathcal{J}_{\mathcal{O}}$ has potential good reduction at any prime ideal. Since $\mathcal{J}_{\mathcal{O}}$ is finite, we can take a number field L and a prime ideal $\mathfrak{p}$ of L above p such that for all $j \in \mathcal{J}_{\mathcal{O}}$, there exists an elliptic curve over L with good reduction at $\mathfrak{p}$ and j-invariant j. Fix an injection of the residue field of L modulo $\mathfrak{p}$ into $\overline{\mathbb{F}}_p$. Let $Ell(\mathcal{O})$ be the set of isomorphism classes of elliptic curves $\mathcal{E}$ over L with good reduction at p and j-invariants in $\mathcal{J}_{\mathcal{O}}$. For every such $\mathcal{E}$, we let $[\cdot]_{\mathcal{E}}$ be the *normalized* $\mathcal{O}$-orientation: that is, such that for any invariant differential ω on $\mathcal{E}$, $([\alpha]_{\mathcal{E}})^*\omega = \alpha\omega$ for all α in $\mathcal{O}$. Then reduction mod $\mathfrak{p}$ defines a map $\rho : Ell(\mathcal{O}) \to SS^{\mathrm{pr}}_{\mathcal{O}}(p)$ sending $\mathcal{E}$ to $(\widetilde{\mathcal{E}}, [.]_{\widetilde{\mathcal{E}}})$, where $\widetilde{\mathcal{E}}$ is the reduction of $\mathcal{E}/L$ at $\mathfrak{p}$ and $[\cdot]_{\widetilde{\mathcal{E}}}$ is the orientation such that $[\alpha]_{\widetilde{\mathcal{E}}} = [\alpha]_{\mathcal{E}} \pmod{\mathfrak{p}}$ for all α in $\mathcal{O}$.

Theorem 1 (Onuki [38, Theorem 3.4])**.** *With the notation above:* $\mathrm{Cl}(\mathcal{O})$ *acts freely and transitively on* $\rho(Ell(\mathcal{O}))$.

4.2 THE NATURAL ORIENTATION

From now on we let $K = \mathbb{Q}(\sqrt{-dp})$, and let $\mathcal{O}_K$ be the maximal order of K.

[2] Working with the class group, we can always replace ideals that are not prime to the conductor with equivalent integral ideals that are.

If $(\mathcal{E}, \psi)$ is a supersingular (d, ϵ)-structure and μ is the associated endomorphism, then

$$\begin{aligned} \iota_\psi : \mathbb{Q}(\sqrt{-dp}) &\longrightarrow \mathrm{End}^0(\mathcal{E}) \\ \sqrt{-dp} &\longmapsto \mu \end{aligned}$$

is a $\mathbb{Z}[\sqrt{-dp}]$-orientation by Proposition 2. We call this the *natural* orientation.

Lemma 1. *If $\mathcal{E}/\mathbb{F}_{p^2}$ is a supersingular elliptic curve with $\#\mathcal{E}(\mathbb{F}_{p^2}) = (p+\epsilon)^2$ and ι is a $\mathbb{Z}[\sqrt{-dp}]$-orientation on $\mathcal{E}$, then ι is the natural orientation for some (d, ϵ)-structure $(\mathcal{E}, \psi)$.*

Proof. Let $\mu := \iota(\sqrt{-dp})$ in $\mathrm{End}(\mathcal{E})$. We have $\deg(\mu) = dp$ and $p \nmid d$, so μ factors over $\mathbb{F}_{p^2}$ into the composition of a d-isogeny and a p-isogeny. Since $\mathcal{E}$ is supersingular, the p-isogeny is isomorphic to π_p, and so $\mu = \pi_p\psi$ for some d-isogeny $\psi : \mathcal{E} \to \mathcal{E}^{(p)}$. It remains to show that $\widehat{\psi} = \epsilon\psi^{(p)}$. Now $[-dp] = \mu^2 = \pi_p\psi\pi_p\psi = \psi^{(p)}\pi_p^2\psi = \psi^{(p)}\psi\pi_p^2$, and $\pi_p^2 = [-\epsilon p]$ because $\mathcal{E}$ is supersingular with $\#\mathcal{E}(\mathbb{F}_{p^2}) = (p+\epsilon)^2$, so $[d] = \epsilon\psi^{(p)}\psi$, and therefore $\widehat{\psi} = \epsilon\psi^{(p)}$. □

Lemma 2. *Let $(\mathcal{E}, \psi)$ and $(\mathcal{E}', \psi')$ be (d, ϵ)-structures with natural orientations ι_ψ and $\iota_{\psi'}$, respectively. If $\phi : \mathcal{E} \to \mathcal{E}'$ is an isogeny, then ϕ is an isogeny (resp. isomorphism) of $\mathbb{Z}[\sqrt{-dp}]$-oriented elliptic curves $(\mathcal{E}, \iota) \to (\mathcal{E}', \iota')$ if and only if it is an isogeny (resp. isomorphism) of (d, ϵ)-structures $(\mathcal{E}, \psi) \to (\mathcal{E}', \psi')$.*

Proof. Let μ resp. μ' be the associated endomorphisms of $(\mathcal{E}, \psi)$ resp. $(\mathcal{E}', \psi')$; then

$$\begin{aligned} \phi_*(\iota_\psi) = \iota_{\psi'} &\iff \phi_*(\iota_\psi)(\sqrt{-dp}) = \iota_{\psi'}(\sqrt{-dp}) && (\sqrt{-dp} \text{ generates } \mathbb{Q}(\sqrt{-dp})) \\ &\iff \phi \circ \mu \circ \widehat{\phi} = \mu'[\deg\phi] && (\text{multiplying by } \deg\phi) \\ &\iff \phi \circ \mu = \mu' \circ \phi && (\text{cancelling } \widehat{\phi}) \\ &\iff \phi \circ \pi_p \circ \psi = \pi_p \circ \psi' \circ \phi && (\text{by definition}) \\ &\iff \pi_p \circ \phi^{(p)} \circ \psi = \pi_p \circ \psi' \circ \phi && (\pi_p \circ \phi = \phi^{(p)} \circ \pi_p) \\ &\iff \phi^{(p)} \circ \psi = \psi' \circ \phi && (\text{cancelling } \pi_p) \end{aligned}$$

and the result follows on comparing definitions. □

Colò and Kohel [16] and Onuki [38] use class-group actions to study the isogeny graphs $\Gamma(SS_{\mathcal{O}}(p))$ with vertex set $SS_{\mathcal{O}}(p)$ for different orders $\mathcal{O}$. Proposition 3 allows us to transfer their results to our setting of (d, ϵ)-structures.

Proposition 3. *The graphs $\Gamma(\mathcal{D}_{d,\epsilon}(p))$ and $\Gamma(SS_{\mathbb{Z}[\sqrt{-dp}]}(p))$ are explicitly isomorphic for $\epsilon = 1$ and $\epsilon = -1$.*

Proof. This follows from Lemmas 1 and 2, once we can show that the isomorphism class of any $\mathbb{Z}[\sqrt{-dp}]$-oriented supersingular curve $(\mathcal{E}, \iota)$ over $\overline{\mathbb{F}}_p$ contains a representative over $\mathbb{F}_{p^2}$ of order $(p+\epsilon)^2$. Since $j(\mathcal{E})$ is in $\mathbb{F}_{p^2}$, after a suitable $\overline{\mathbb{F}}_p$-isomorphism we may suppose that $\mathcal{E}$ is defined over $\mathbb{F}_{p^2}$ and $\#\mathcal{E}(\mathbb{F}_{p^2}) = (p+\epsilon)^2$; and then ι is defined over $\mathbb{F}_{p^2}$ because for a supersingular elliptic curve over $\mathbb{F}_{p^2}$ all of the endomorphisms are defined over $\mathbb{F}_{p^2}$. □

Let $K = \mathbb{Q}(\sqrt{-dp})$. The order $\mathbb{Z}[\sqrt{-dp}]$ has index 2 in $\mathcal{O}_K$ if $-dp \equiv 1 \pmod 4$, and is equal to $\mathcal{O}_K$ otherwise. If $-dp \not\equiv 1 \pmod 4$, then, every natural orientation is a primitive $\mathcal{O}_K$-orientation; if $-dp \equiv 1 \pmod 4$, each natural orientation is either a primitive $\mathbb{Z}[\sqrt{-dp}]$-orientation or a primitive $\mathcal{O}_K$-orientation.

Proposition 4. *Let $(\mathcal{E}, \psi)$ be a supersingular (d, ϵ)-structure with natural orientation ι_ψ.*

1. *If $-dp \not\equiv 1 \pmod 4$, then ι_ψ is a primitive $\mathcal{O}_K$-orientation.*
2. *If $-dp \equiv 1 \pmod 4$, then ι_ψ is a primitive $\mathcal{O}_K$-orientation if the associated endomorphism μ fixes $\mathcal{E}[2]$* pointwise*, and a primitive $\mathbb{Z}[\sqrt{-dp}]$-orientation otherwise.*

Proof. By definition, ι_ψ is a $\mathbb{Z}[\sqrt{-dp}]$-orientation. To complete Case (2), it suffices to check whether the element $\iota_\psi(\frac{1}{2}(-1+\sqrt{-dp})) = \frac{1}{2}(\mu - [1])$ of $\mathrm{End}^0(\mathcal{E}) \cap \iota_\psi(K)$ is in $\mathrm{End}(\mathcal{E})$ (because $\frac{1}{2}(-1+\sqrt{-dp})$ generates $\mathcal{O}_K$, but is not in $\mathbb{Z}[\sqrt{-dp}]$). This is the case if and only if $\mu - [1]$ factors over $[2]$, if and only if μ fixes $\mathcal{E}[2]$ pointwise. □

In the light of Propositions 3 and 4, we partition $\mathcal{D}_{d,\epsilon}(p)$ into two subsets:

$$\mathcal{D}_{d,\epsilon}(p) = \mathcal{D}^{\max}_{d,\epsilon}(p) \sqcup \mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)\,,$$

where $\mathcal{D}^{\max}_{d,\epsilon}(p)$ contains the classes whose natural orientations are primitive $\mathcal{O}_K$-orientations, and $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ contains the classes whose natural orientations are primitive orientations by the order of conductor 2 in $\mathcal{O}_K$. If $-dp \not\equiv 1$

(mod 4), then $\mathcal{D}^{\max}_{d,\epsilon}(p) = \mathcal{D}_{d,\epsilon}(p)$ and $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p) = \emptyset$. If $-dp \equiv 1 \pmod 4$, then $[\mathcal{O}_K : \mathbb{Z}[\sqrt{-dp}]] = 2$, so $\mathcal{D}^{\max}_{d,\epsilon}(p)$ resp. $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ consists of the (d,ϵ)-structures where μ acts trivially resp. nontrivially on the 2-torsion.

Given Lemma 2, ℓ-isogenies of (d,ϵ)-structures are "ascending", "descending", and "horizontal" with respect to the natural orientations: we have horizontal ℓ-isogenies between vertices in $\mathcal{D}^{\max}_{d,\epsilon}(p)$ and between vertices in $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$, while $\mathcal{D}^{\max}_{d,\epsilon}(p)$ and $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ are connected by ascending and descending 2-isogenies. In the language of isogeny volcanoes, vertices in $\mathcal{D}^{\max}_{d,\epsilon}(p)$ form the "craters", and vertices in $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ the "floors".

4.3 THE CLASS GROUP ACTION

Proposition 3 translates the action of $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$ on $SS_{\mathbb{Z}[\sqrt{-dp}]}(p)$ defined above into an action on $\mathcal{D}_{d,\epsilon}(p)$. Theorem 2 makes this precise: it shows that $\mathcal{D}^{\max}_{d,\epsilon}(p)$ is a principal homogeneous space (or torsor) under $\mathrm{Cl}(\mathcal{O}_K)$, and that if $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ is not empty then it is a principal homogeneous space under $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$.

Theorem 2. *Let $K = \mathbb{Q}(\sqrt{-dp})$, with maximal order $\mathcal{O}_K$, and let $\epsilon = \pm 1$.*

- *The class group $\mathrm{Cl}(\mathcal{O}_K)$ acts freely and transitively on $\mathcal{D}^{\max}_{d,\epsilon}(p)$.*
- *If $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p) \neq 0$, then $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$ acts freely and transitively on $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$.*

Proof. Let $\mathcal{O} = \mathcal{O}_K$ or $\mathbb{Z}[\sqrt{-dp}]$. Since p does not split in K, Theorem 1 tells us that $\mathrm{Cl}(\mathcal{O})$ acts freely and transitively on $\rho(Ell(\mathcal{O})) \subseteq SS^{\mathrm{pr}}_{\mathcal{O}}(p)$. Given the isomorphism of Proposition 3, it only remains to prove that $\rho(Ell(\mathcal{O})) = SS^{\mathrm{pr}}_{\mathcal{O}}(p)$. For any $(\mathcal{E},\iota)$ in $SS^{\mathrm{pr}}_{\mathcal{O}}(p)$, Proposition 3.3 of [38] tells us that $(\mathcal{E},\iota)$ or $(\mathcal{E},\iota)^{(p)}$ is in $\rho(Ell(\mathcal{O}))$. In our case, *both* are in $\rho(Ell(\mathcal{O}))$, so the action on $SS^{\mathrm{pr}}_{\mathcal{O}}(p)$ is free: since $\mathcal{E}[\mathfrak{d}] = \mathcal{E}[d] \cap \ker\mu = \ker\psi$, the action of $\mathfrak{d} = (d, \sqrt{-dp})$ on $SS^{\mathrm{pr}}_{\mathcal{O}}(p)$ maps $(\mathcal{E},\iota)$ to $(\mathcal{E},\iota)^{(p)}$, because it maps $(\mathcal{E},\psi)$ to $(\mathcal{E},\psi)^{(p)}$. □

Corollary 1. *Let $K = \mathbb{Q}(\sqrt{-dp})$, with maximal order $\mathcal{O}_K$. If $h_K = \#\mathrm{Cl}(\mathcal{O}_K)$, then*

$$\#\mathcal{D}^{\max}_{d,\epsilon}(p) = h_K \qquad \text{and} \qquad \#\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p) = \begin{cases} h_K & \text{if } -dp \equiv 1 \pmod 8, \\ 3h_K & \text{if } -dp \equiv 5 \pmod 8, \\ 0 & \text{otherwise}. \end{cases}$$

Proof. By Theorem 2, we have $\#\mathcal{D}^{\max}_{d,\epsilon}(p) = \#\mathrm{Cl}(\mathcal{O}_K)$ and either $\#\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p) = 0$ (if $-dp \not\equiv 1 \pmod 4$) or $\#\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p) = \#\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$ (if $-dp \equiv 1 \pmod 4$). It remains to compute $\#\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$ in the case $-dp \equiv 1 \pmod 4$, where $\mathbb{Z}[\sqrt{-dp}]$ has conductor 2. In this case, the formula of [19, Theorem 7.24] simplifies to

$$\#\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}]) = \frac{\#\mathrm{Cl}(\mathcal{O}_K)}{[\mathcal{O}_K^\times : \mathbb{Z}[\sqrt{-dp}]^\times]} \left(2 - \left(\frac{-dp}{2}\right)\right),$$

where the Kronecker symbol $(-dp/2)$ is 0 if $2 \mid -dp$, 1 if $-dp \equiv \pm 1 \pmod 8$, and -1 if $-dp \equiv \pm 3 \pmod 8$. The result follows on noting that $[\mathcal{O}_K^\times : \mathbb{Z}[\sqrt{-dp}]^\times] = 1$, because $-dp$ is never -3 or -4. □

Remark 3. *The Brauer–Siegel theorem states that asymptotically, $\log_2(h_K) \sim \frac{1}{2}\log_2|\Delta_K|$, where $\Delta_K = -dp$ if $-dp \equiv 1 \pmod 4$, and $-4dp$ otherwise. (See e.g. [34, Ch. XVI] for details.)*

4.4 COMPUTING THE CLASS GROUP ACTION

Suppose we want to compute the action of (the class of) an ideal $\mathfrak{l} = (\ell, a + b\sqrt{-dp})$ on some $(\mathcal{E},\psi)$ in $\mathcal{D}_{d,\epsilon}(p)$. Following [20], we consider two approaches: "Vélu" and "modular".

In the **"Vélu"** approach, we compute a generator K_ℓ of the kernel $\mathcal{E}[\mathfrak{l}]$ of ϕ: that is, K_ℓ is a point in $\mathcal{E}[\ell]$ such that $[a]\mu(K_\ell) = -[b]K_\ell$. This point may only be defined over an extension $\mathbb{F}_{p^{2r}}$ of $\mathbb{F}_{p^2}$. We then compute the quotient isogeny $\phi : \mathcal{E} \to \mathcal{E}' := \mathcal{E}/\langle K_\ell\rangle$ using Vélu's formulæ, at a cost of $O(\ell)$ $\mathbb{F}_{p^{2r}}$-operations, or the algorithm of [6], in $\widetilde{O}(\sqrt{\ell})$ $\mathbb{F}_{p^{2r}}$-operations. Finally, we push ψ through ϕ by computing the image of its kernel subgroup and choosing the correct "sign". If we are given an $\mathbb{F}_{p^2}$-rational generator G for $\ker\psi$, then pushing ψ through ϕ essentially costs one isogeny evaluation; otherwise, this amounts to an exercise in symmetric functions, with a cost on the order of $O(d)$ isogeny evaluations. Each evaluation costs $O(\ell)$ or $\widetilde{O}(\sqrt{\ell})$ $\mathbb{F}_{p^2}$-operations. The total cost is dominated by the cost of the multiplication by the cofactor $\#E(\mathbb{F}_{p^{2r}})/\ell$ needed to find K_ℓ: we have $\log(\#E(\mathbb{F}_{p^{2r}})/\ell) = 2r\log p$, so constructing K_ℓ requires $O(r^2\log p)$ operations in $\mathbb{F}_{p^2}$.

To compute the action of $\mathfrak{l}$ on $(\mathcal{E},\psi)$, we compute $G = \gcd(\Phi_d(X, X^p), \Phi_\ell(j(\mathcal{E}), X))$ (if $d = 1$, then we take $\Phi_1(X, X^p) = X^p - X$). In general G has only two roots in $\mathbb{F}_{p^2}$, corresponding to the two ℓ-neighbours. In a

non-backtracking walk we can divide by $X - j(\mathcal{E}')$, where $(\mathcal{E}', \psi')$ is the preceding vertex, to find the next step. Otherwise, we can distinguish between the two neighbours by examining the action of μ on the ℓ-torsion. Care must be taken to identify, and to appropriately handle, the exceptional case where a neighbouring j-invariant admits multiple (d, ϵ)-structures modulo negation (as with the vertices A and C in the example of Figure 2 below).

To compute $\gcd(\Phi_d(X, X^p), \Phi_\ell(j(E), X))$, compute $F(X) := \Phi_\ell(j(\mathcal{E}), X)$ in $O(\ell)$ $\mathbb{F}_{p^2}$-operations, and then $Y := X^p \bmod F(X)$ using the square-and-multiply algorithm in $O(\ell \log p)$ $\mathbb{F}_{p^2}$-operations. We then compute $Z := \Phi_d(X, Y) \mod F$, and then $\gcd(Z, F)$, in $O(d^2\ell^2)$ $\mathbb{F}_{p^2}$-operations. Generally ℓ is polynomial in $\log p$, but typically it is even smaller, and then the dominating step is the computation of Y.

As in the ordinary case [20], the Vélu approach is more efficient when $r^2 < \ell$; in particular, when K_ℓ is defined over $\mathbb{F}_{p^2}$. If we are free to choose p, then we can optimize systems that use the action of a series of small primes ℓ_i by taking p such that the ℓ_i split in $\mathbb{Z}[\sqrt{-dp}]$ *and* $\ell_i \mid p + \epsilon$, that is, $p = c \cdot \prod_{i=1}^{n} \ell_i - \epsilon$ with c a cofactor making p prime. In the case $d = 1$, this is exactly the optimization that is key to making CSIDH practical.

Remark 4. *It would be interesting to look for an expression for the group action operating directly on the parameters in the Hasegawa families of §3.3 and §3.4.*

5 THE SUPERSINGULAR ISOGENY GRAPH

We can now describe the structure of the isogeny graph $\Gamma(\mathcal{D}_{d,\epsilon}(p))$. Factoring isogenies, it suffices to describe $\Gamma_\ell(\mathcal{D}_{d,\epsilon}(p))$ for prime ℓ. The class group actions of Theorem 2 imply the isogeny counts in Table 1.

Table 1: The number of horizontal, ascending, and descending ℓ-isogenies from each vertex in the ℓ-isogeny graph.

Prime ℓ	Conditions on (d, p)	Vertex (sub)set	Horizontal	Ascending	Descending
$\ell = 2$	$-dp \equiv 1 \pmod 8$	$\mathcal{D}^{\max}_{d,\epsilon}(p)$	2	0	1
		$\mathcal{D}^{\text{sub}}_{d,\epsilon}(p)$	0	1	0
	$-dp \equiv 5 \pmod 8$	$\mathcal{D}^{\max}_{d,\epsilon}(p)$	0	0	3
		$\mathcal{D}^{\text{sub}}_{d,\epsilon}(p)$	0	1	0
	$-dp \not\equiv 1, 5 \pmod 8$	$\mathcal{D}_{d,\epsilon}(p)$	1	0	0
$\ell > 2$	—	$\mathcal{D}_{d,\epsilon}(p)$	$1 + (-dp/\ell)$	0	0

Examples. Figures 1, 2, and 3, display ℓ-isogeny graphs on $\mathcal{D}_{3,1}(101)$, $\mathcal{D}_{3,-1}(97)$, and $\mathcal{D}_{3,1}(83)$ for various ℓ generating the class groups. These figures also form examples of the various 2-isogeny structures listed in Table 1. In each example, we use the Hasegawa parametrization of §3.4 to encode vertices.

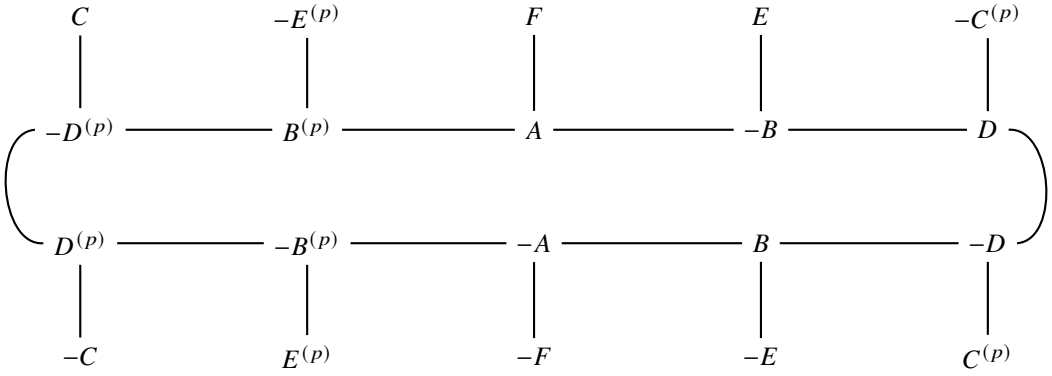

Figure 1: $\Gamma_2(\mathcal{D}_{3,1}(101))$ for $\ell = 2$. The class group of $\mathbb{Q}(\sqrt{-303})$ is isomorphic to $\mathbb{Z}/10\mathbb{Z}$, and generated by an ideal over 2 (we see this in the length-10 cycle). The correspondence between vertex labels and parameters for the degree-3 Hasegawa family of §3.4 (with $\Delta = 2$) is $A \leftrightarrow 0$, $B \leftrightarrow 6$, $C \leftrightarrow 24$, $D \leftrightarrow 25$, and $E \leftrightarrow 42$; the vertex F, which has Hasegawa parameter ∞, is $(\mathcal{E}, \psi)$ with $\mathcal{E} : y^2 = x^3 + 1$ and $\psi : (x, y) \mapsto ((67x^3 + 66)/x^2, (89x^3 + 96)\sqrt{2}y/x^3)$. Note that $A^{(p)} = -A$ and $F^{(p)} = -F$. The underlying curves of B and C are isomorphic. The vertices in the cycle are in $\mathcal{D}^{\max}_{3,1}(101)$; the others are in $\mathcal{D}^{\text{sub}}_{3,1}(101)$.

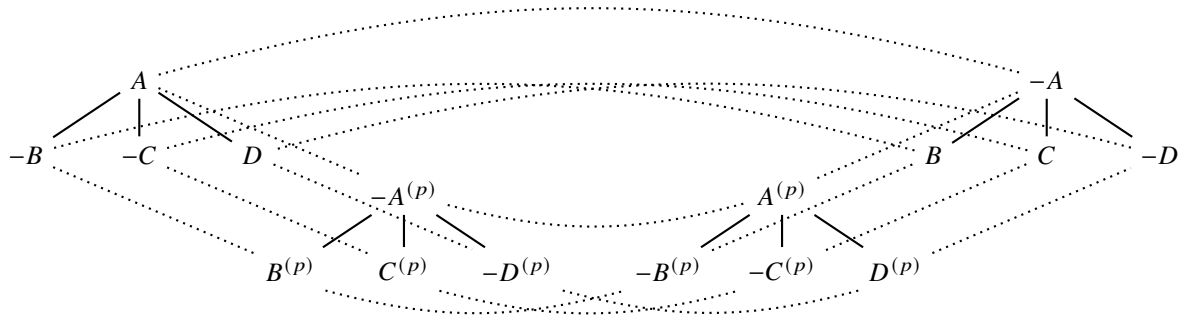

Figure 2: The isogeny graphs $\Gamma_2(\mathcal{D}_{3,-1}(97))$ (solid) and $\Gamma_5(\mathcal{D}_{3,-1}(97))$ (dotted). We have $\mathrm{Cl}(\mathbb{Q}(\sqrt{-3\cdot 97})) \cong \mathbb{Z}/4\mathbb{Z}$, generated by an ideal over 5. The 2-isogenies are ascending/descending up/down the page; the 5-isogenies are horizontal. The correspondence between vertex labels and parameters for the degree-3 Hasegawa family of §3.4 (with $\Delta = 5$) is $A \leftrightarrow 47$, $B \leftrightarrow 1$, $C \leftrightarrow 14$, and $D \leftrightarrow 22$. The underlying curves of A and C are isomorphic.

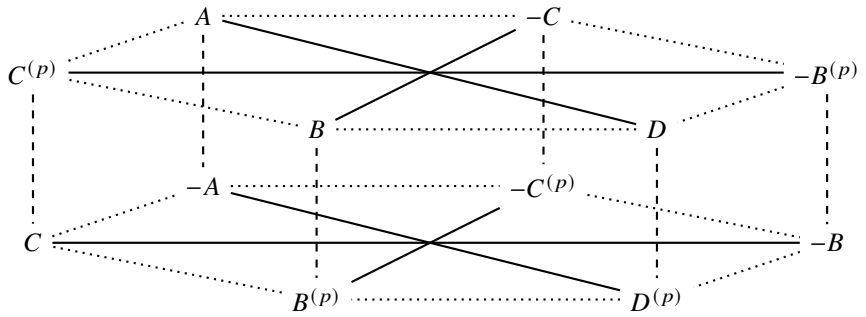

Figure 3: $\Gamma_\ell(\mathcal{D}_{3,1}(83))$ for $\ell = 2$ (solid), $\ell = 3$ (dashed) and $\ell = 5$ (dotted). All isogenies are horizontal. We have $\mathrm{Cl}(\mathbb{Q}(\sqrt{-3\cdot 83})) \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/6\mathbb{Z}$, with the $\mathbb{Z}/2\mathbb{Z}$-factor generated by the ideal above 3, and the $\mathbb{Z}/6\mathbb{Z}$-factor generated by an ideal above 5 (we see this in the length-6 cycles); the ideal above 2 is the cube of an ideal above 5. The correspondence between vertex labels and parameters for the degree-3 Hasegawa family of §3.4 (with $\Delta = 2$) is $A \leftrightarrow 0$, $B \leftrightarrow 32$, $C \leftrightarrow 40$; the vertex D, which has Hasegawa parameter ∞, is $(\mathcal{E} : y^2 = x^3 + 1, \psi)$ where ψ maps (x, y) to $(((72\sqrt{2}+14)x^3 + (39\sqrt{2}+56))/x^2, \sqrt{2}(35x^3+52)y/x^3)$. Note that $-A = A^{(p)}$.

Involutions. There are two obvious involutions on $\Gamma(\mathcal{D}_{d,\epsilon}(p))$, negation and conjugation. These are generally not the only involutions. Every prime ℓ dividing the discriminant ramifies in $\mathcal{O}_K$ (and $\mathbb{Z}[\sqrt{-dp}]$); the prime $\mathfrak{l}$ over ℓ gives an element of order 2 in $\mathrm{Cl}(\mathcal{O}_K)$ (and $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$), and thus an involution on $\Gamma(\mathcal{D}_{d,\epsilon}(p))$. Let $\mathfrak{d}_1, \ldots, \mathfrak{d}_n$ be the primes above the prime factors of d, and $\mathfrak{p}$ the prime above p; note that $[\mathfrak{d}_1]\cdots[\mathfrak{d}_n] = [\mathfrak{p}]$, because $\mathfrak{d}_1 \cdots \mathfrak{d}_n\mathfrak{p} = (\mu)$. If $-dp \equiv 1$ or $2 \pmod 4$ then $\mathrm{Cl}(\mathcal{O}_K)[2] = \langle [\mathfrak{d}_1], \ldots, [\mathfrak{d}_n], [\mathfrak{p}]\rangle$, so $\mathrm{Cl}(\mathcal{O}_K)[2] \cong (\mathbb{Z}/2\mathbb{Z})^n$. If $-dp \equiv 3 \pmod 4$, then $\mathrm{Cl}(\mathcal{O}_K)[2] = \langle [\mathfrak{a}], [\mathfrak{d}_1], \ldots, [\mathfrak{d}_n], [\mathfrak{p}]\rangle$ where $\mathfrak{a}$ is the ideal above 2, and $\mathrm{Cl}(\mathcal{O}_K)[2] \cong (\mathbb{Z}/2\mathbb{Z})^{n+1}$. In each case, the action of the ideal class $\prod_i [\mathfrak{d}_i] = [\mathfrak{p}]$ on any (d,ϵ)-structure $(\mathcal{E}, \psi)$ is realised by the isogeny $\psi : (\mathcal{E}, \psi) \to (\mathcal{E}^{(p)}, \psi^{(p)})$, and is therefore equal to the conjugation involution.

Since the group actions are free, each of the involutions that come from nontrivial 2-torsion elements in the class groups—including conjugation—has no fixed points. Negation, on the other hand, can have fixed points: for example, if $p \equiv 3 \pmod 4$ and $\mathcal{E}$ is the curve with j-invariant 1728, and i is an automorphism of degree 4, then $(\mathcal{E}, i)$ is a $(1,1)$-structure, and $(\mathcal{E}, i) \cong (\mathcal{E}, -i)$. This is the only fixed point among $(1,1)$-structures, and its existence is implied by the fact that the class number of $\mathrm{Cl}(\sqrt{-p})$ is odd when $p \equiv 3 \pmod 4$.

Remark 5. *If $-dp \equiv 5 \pmod 8$, then there is an order-3 automorphism T of $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ cycling the triplets of vertices with ascending 2-isogenies to the same vertex in $\mathcal{D}^{\max}_{d,\epsilon}(p)$. We will see that T is induced by the action of an ideal class in $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$. The ideal $\mathfrak{t} = (4, \sqrt{-dp}-1)\mathbb{Z}[\sqrt{-dp}]$ has order 3 in $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$, but capitulates to become the principal ideal (2) in $\mathcal{O}_K$ (because $\sqrt{-dp}-1 = 2\omega$, where ω is the unit $\frac{1}{2}(\sqrt{-dp}-1)$); indeed, $\mathfrak{t}$ generates the kernel of the canonical homomorphism $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}]) \to \mathrm{Cl}(\mathcal{O}_K)$. Since $\mathfrak{t}$ meets the conductor, its action on $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ is not well-defined, but we can consider the action of an equivalent ideal in the class group. Let $\prod_i \ell_i^{e_i}$ be the prime factorization of $(dp+1)/4$ (and note that each ℓ_i is odd); then $(\sqrt{-dp}-1) = \mathfrak{t}\cdot\prod_i \mathfrak{l}_i^{e_i}$ where $\mathfrak{l}_i := (\ell_i, \sqrt{-dp}-1)$; the product $\prod_i \mathfrak{l}_i^{e_i}$ is equivalent to $\mathfrak{t}$ in $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$, prime to the conductor, and its action*

on $\mathcal{D}^{\text{sub}}_{d,\epsilon}(p)$ induces the automorphism T. In the case where $d = 1$ (CSIDH), this is explained at length in [39].

Crossroads. The map $(\mathcal{E}, \psi) \mapsto \mathcal{E}$ defines a covering from $\Gamma(\mathcal{D}_{d,\epsilon}(p))$ onto a subgraph of the isogeny graph of all supersingular curves over $\mathbb{F}_{p^2}$. For $d_1 \neq d_2$ the images of $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$ and $\Gamma(\mathcal{D}_{d_2,\epsilon}(p))$ can intersect, forming "crossroads" where we can switch from walking in $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$ into $\Gamma(\mathcal{D}_{d_2,\epsilon}(p))$, and vice versa.

Definition 4. *Let $d_1 \neq d_2$ be squarefree integers such that $d_1 d_2$ is squarefree. We say that a supersingular curve $\mathcal{E}/\mathbb{F}_{p^2}$ with $\#\mathcal{E}(\mathbb{F}_{p^2}) = (p+\epsilon)^2$ is a (d_1, d_2)*-crossroad *if there exist isogenies $\psi_1 : \mathcal{E} \to \mathcal{E}_1$ and $\psi_2 : \mathcal{E} \to \mathcal{E}_2$ such that $(\mathcal{E}, \psi_1)$ is a (d_1, ϵ)-structure and $(\mathcal{E}, \psi_2)$ is a (d_2, ϵ)-structure.*

If $(\mathcal{E}, \psi)$ is a (d_1, ϵ)-structure, then we can easily check whether $\mathcal{E}$ is a (d_1, d_2)-crossroad by evaluating the classical modular polynomial Φ_{d_2} at $(j(\mathcal{E},), j(\mathcal{E})^p)$. However, (d_1, d_2)-crossroads are generally very rare. Indeed, if $\mathcal{E}$ is a (d_1, d_2)-crossroad, then it has an endomorphism of degree $d_1 d_2$ with cyclic kernel (in particular, (d_1, d_2)-crossroads with $d_1 d_2 < \frac{\sqrt{p}}{2}$ appear in the isogeny "valleys" described in [35]). We can therefore enumerate the entire set of (d_1, d_2)-crossroads over a given $\mathbb{F}_{p^2}$ by computing the set of roots j of $\Phi_{d_1 d_2}(x, x)$ in $\mathbb{F}_{p^2}$, and then checking for which j we have $\Phi_{d_1}(j, j^p) = 0$. The polynomial $\Phi_{d_1 d_2}(x, x)$ has degree $\prod_\ell (\ell + 1)$ where ℓ ranges over the prime factors of $d_1 d_2$, so there are only $O(d_1 d_2)$ (d_1, d_2)-crossroads (up to isomorphism) among the $O(\sqrt{dp})$ vertices in $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$.

But while crossroads are rare, computing the few examples is relatively easy, and computing (d_1, d_2)-crossroads gives us a useful way of quickly constructing some vertices in $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$ (and in $\Gamma(\mathcal{D}_{d_2,\epsilon}(p))$). Suppose we want to construct a vertex in $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$. Since the vertices in $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$ correspond to curves with an endomorphism subring isomorphic to $\mathbb{Z}[\sqrt{-d_1 p}]$, we might try to construct a vertex from a root in $\mathbb{F}_{p^2}$ of the Hilbert class polynomial for $\mathbb{Q}(\sqrt{-d_1 p})$; but the degree of this polynomial, which is the order of the class group, is exponential with respect to $\log p$, so this approach is infeasible for large p. Instead, we choose a small squarefree d_2 such that p does not split in the maximal order of $\mathbb{Q}(\sqrt{-d_1 d_2})$. If there exists a (d_1, d_2)-crossroad $\mathcal{E}/\mathbb{F}_{p^2}$, then its j-invariant is a root in $\mathbb{F}_{p^2}$ of a quadratic factor of the Hilbert class polynomial for $\mathbb{Q}(\sqrt{-d_1 d_2})$, because the composition of the d_1-isogeny $\mathcal{E} \to \mathcal{E}^{(p)}$ with the conjugate d_2-isogeny $\mathcal{E}^{(p)} \to \mathcal{E}$ is a cyclic endomorphism of degree $d_1 d_2$. All other vertices in $\Gamma(\mathcal{D}_{d_1,\epsilon}(p))$ can then be reached through the class group action.

6 CRYPTOGRAPHIC APPLICATIONS

The action of $\mathrm{Cl}(\mathcal{O}_K)$ on $\mathcal{D}^{\text{max}}_{d,\epsilon}(p)$ and $\mathrm{Cl}(\mathbb{Z}[\sqrt{-dp}])$ on $\mathcal{D}^{\text{sub}}_{d,\epsilon}(p)$ makes $\Gamma(\mathcal{D}_{d,\epsilon}(p))$ a natural candidate setting for group-action/HHS-based postquantum cryptosystems following Stolbunov [43, 46, 47] and Couveignes [18]. For example, for each $d > 1$, we can define a key exchange algorithm on $\mathcal{D}_{d,\epsilon}(p)$ generalizing CSIDH [11], which uses the action of $\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$ on $\mathcal{D}^{\text{sub}}_{1,1}(p)$ and CSURF [10], which uses the action of $\mathrm{Cl}(\mathbb{Q}(\sqrt{-p}))$ on $\mathcal{D}^{\text{max}}_{1,1}(p)$. Despite the prominence of orientations, the relationship between key exchange in $\mathcal{D}_{d,\epsilon}(p)$ and the OSIDH protocol [16] is distant. The $\mathcal{O}$-orientations in OSIDH involve orders $\mathcal{O}$ with massive conductors in $\mathcal{O}_K$ where $\mathcal{O}_K$ has tiny class number; here, $\mathcal{O}$ has tiny conductor and $\mathcal{O}_K$ has massive class number.

6.1 HARD PROBLEMS

The conjectural hard problems for the action of $\mathrm{Cl}(\mathcal{O}_K)$ on $\mathcal{D}_{d,\epsilon}(p)$ are vectorization (the analogue of the DLP) and parallelization (the analogue of the CDHP) from Couveigne's *Hard Homogenous Spaces* framework [18].

Definition 5 (Vectorization)**.** *Given $(\mathcal{E}, \psi)$ and $(\mathcal{E}', \psi')$ in $\mathcal{D}_{d,\epsilon}(p)$, find $\mathfrak{a} \in \mathrm{Cl}(\mathcal{O}_K)$ such that $\mathfrak{a} \cdot (\mathcal{E}, \psi) = (\mathcal{E}', \psi')$.*

Definition 6 (Parallelization)**.** *Given $(\mathcal{E}_0, \psi_0)$, $(\mathcal{E}_1, \psi_1)$, and $(\mathcal{E}_2, \psi_2)$ in $\mathcal{D}_{d,\epsilon}(p)$, compute the unique $(\mathcal{E}_3, \psi_3)$ in $\mathcal{D}_{d,\epsilon}(p)$ such that $(\mathcal{E}_3, \psi_3) = (\mathfrak{a}_1 \mathfrak{a}_2) \cdot (\mathcal{E}_0, \psi_0)$ where $(\mathcal{E}_i, \psi_i) = \mathfrak{a}_i \cdot (\mathcal{E}_0, \psi_0)$ for $i = 1$ and 2.*

Solving Vectorization immediately solves Parallelization. In the opposite direction, no classical reduction is known, but the quantum equivalence of these two problems is shown in [25].

An extensive study of the possible classical and quantum attacks on Vectorization for $d = 1$ can be found in [11]; all of these attacks extend to $d > 1$ with a slowdown at most polynomial in d for class groups of the same size, with that slowdown due to potentially more complicated isogeny evaluation and comparison algorithms. The best classical attack known on Vectorization is to use random walks in $\Gamma(\mathcal{D}_{d,\epsilon}(p))$, exactly as in the $d = 1$ case in [21], which gives a solution after an expected $O((dp)^{1/4})$ isogeny steps. Since Vectorization is an instance of the Abelian Hidden Shift Problem, the best quantum attack is Kuperberg's algorithm [32, 42, 33] using the Childs–Jao–Soukharev quantum isogeny-evaluation algorithm as a subroutine [15], adapted to push ψ through the ℓ-isogenies. This adaptation may incur a practically significant but asymptotically negligible cost; the result is a subexponential algorithm running in time $L_{dp}[1/2, \sqrt{2}]$. Even for $d = 1$, there is some debate as to the concrete

cost of this quantum algorithm, and the size of p required to provide a cryptographically hard problem instance for common security levels [5, 8, 41]. (If and) when some consensus forms on secure parameter sizes for CSIDH, the same parameter sizes should make Vectorization and Parallelization in $\mathcal{D}_{d,\epsilon}(p)$ cryptographically hard, too.

We should also consider the impact of the various involutions on $\Gamma(\mathcal{D}_{d,\epsilon}(p))$. The negation involution already exists for $d = 1$, where it essentially flips between a curve and its quadratic twist over $\mathbb{F}_p$. This involution has not yet been exploited to give an interesting speedup in solving Vectorization or Parellization in the case $d = 1$; a speedup for any d would be an interesting result. For $d > 1$, however, there is at least one new involution: namely, conjugation. We note that solving Vectorization modulo conjugation solves Vectorization, because a vertex and its conjugate are always connected by the action of an ideal of norm d. Working modulo conjugation allows us to shrink search spaces by a factor of 2, yielding a speedup by a factor of up to $\sqrt{2}$ analogous to working modulo negation when solving the classical ECDLP (as in [4]). When d has n prime factors, we get more involutions that would allow us to work with equivalence classes of 2^n vertices, shrinking the search spaces by a factor of 2^n. Prime d therefore seems the simplest and strongest case to us.

Finally, we note that if a random walk should wander into a crossroad, then we have found an isogeny to a supersingular curve with much known on its endomorphism ring. In this case, attacks analogous to that of [27] should apply. But as we have seen, crossroads are vanishingly rare; their existence should not create any weakness for schemes based on $\Gamma(\mathcal{D}_{d,\epsilon}(p))$, no more than they do for CSIDH.

6.2 NON-INTERACTIVE KEY EXCHANGE

We now describe a non-interactive key-exchange protocol based on the class group action on $\Gamma(\mathcal{D}_{d,\epsilon}(p))$, generalizing CSIDH (the case $d = 1$). The public parameters are a prime p, a prime d, an ϵ in $\{1, -1\}$, a set of primes $\{\ell_i\}_{i=1}^n$ prime to dp and splitting in $\mathbb{Q}(\sqrt{-dp})$, together with a prime ideal $\mathfrak{l}_i$ above each ℓ_i, and a "starting" vertex $(\mathcal{E}_0, \psi_0)$ in $\mathcal{D}_{d,\epsilon}(p)$ (constructed using the crossroad technique, for example). We also fix a secret keyspace $\mathcal{K} \subset \mathbb{Z}^n$ of exponent vectors such that $\#\mathcal{K} \geq 2^{2\lambda}$ to provide λ bits of security against meet-in-the-middle attacks (though smaller $\mathcal{K}$ may suffice: see [14]). The prime p must be large enough that Vectorization and Parallelization cannot be solved in fewer than 2^λ classical operations, or a comparable quantum effort.

For key generation, each user randomly samples their private key as a vector $(e_i)_{1 \leq i \leq n}$ from $\mathcal{K}$, representing the ideal class $[\mathfrak{a}] = [\prod_{i=1}^n \mathfrak{l}_i^{e_i}]$ in $\mathrm{Cl}(\mathcal{O}_K)$. Their public key is a vertex $(\mathcal{E}, \psi)$ representing $[\mathfrak{a}] \cdot (\mathcal{E}_0, \psi_0)$, which we can compute using the methods of §4.4. The public key may be compressed to a single element of $\mathbb{F}_p$ plus a few bits using the modular techniques of §3.

For key exchange, suppose Alice and Bob have key pairs $([\mathfrak{a}], (\mathcal{E}_A, \psi_A))$ and $([\mathfrak{b}], (\mathcal{E}_B, \psi_B))$, respectively. Alice receives and validates $(\mathcal{E}_B, \psi_B)$, and computes $S_{AB} = (\mathcal{E}_{AB}, \psi_{AB}) = [\mathfrak{a}] \cdot (\mathcal{E}_B, \psi_B)$; Bob receives and validates $(\mathcal{E}_A, \psi_A)$, and computes $S_{BA} = (\mathcal{E}_{BA}, \psi_{BA}) = [\mathfrak{b}] \cdot (\mathcal{E}_A, \psi_A)$. The commutativity of the group action implies that $S_{AB} \cong S_{BA}$, so Alice and Bob have a shared secret *up to isomorphism*. To obtain a unique shared value for cryptographic key derivation, they can derive a modular "compressed" representation of the shared secret as in §3 (for example, when $d = 2$ or 3, the parameter u for the family of §3.3 or §3.4 and a sign bit suffice), or simply take $j(\mathcal{E}_{AB}) = j(\mathcal{E}_{BA})$ with a minimal security loss.

Remark 6. *When ideal classes represent cryptographic secrets, it is important to compute their actions in constant time. A number of techniques have been proposed for this in the context of CSIDH [36, 40, 12, 9, 3]. Each of these methods generalizes in a straightforward way to compute class-group actions on (d, ϵ)-structures. The only real algorithmic difference when evaluating an isogeny $\phi : (\mathcal{E}, \psi) \to (\mathcal{E}', \psi')$ is that the isogeny ψ must be pushed through ϕ in constant-time as well. For $d = 2$ and 3, this amounts to pushing the x-coordinate of a single point through the isogeny, something that is already part of constant-time CSIDH implementations. For $d > 3$ the kernel polynomial of ψ can be pushed through ϕ using symmetric functions.*

6.3 KEY VALIDATION AND SUPERSINGULARITY TESTING

Public key validation is an important step in many public-key cryptosystems, notably in non-interactive key exchanges where it is a defence against active attacks. In our situation, this amounts to proving that a pair $(\mathcal{E}, \psi)$ represents an element of $\mathcal{D}_{d,\epsilon}(p)$ (or $\mathcal{D}^{\max}_{d,\epsilon}(p)$, or $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$). The first step is to check that $(\mathcal{E}, \psi)$ is a (d, ϵ)-structure: specifically, we must check that ψ is indeed an isogeny from $\mathcal{E}$ to $\mathcal{E}^{(p)}$ and that $\widehat{\psi} = \epsilon\psi^{(p)}$. This can be done with two d-isogeny computations, which costs very little when d is small.

Verifying supersingularity is more complicated. For $d = 1$ (CSIDH), we just check whether a curve over $\mathbb{F}_p$ has order $p + 1$, which can be done efficiently by probabilistically generating a point of order $m \mid p + 1$ with $m > 4\sqrt{p}$ (see [11, §5]). But this technique does not extend to $d > 1$, where we must check if $\mathcal{E}/\mathbb{F}_{p^2}$ has $(p + \epsilon)^2$ points: our valid curves have $\mathcal{E}(\mathbb{F}_{p^2}) \cong (\mathbb{Z}/(p + \epsilon)\mathbb{Z})^2$, and therefore no points with the required order $> 4p$.

Instead, for $d > 1$ we can specialize the deterministic supersingularity test of Sutherland [48]. Let $\pi_\mathcal{E}$ be the Frobenius endomorphism of $\mathcal{E}/\mathbb{F}_{p^2}$. The discriminant of $\mathbb{Z}[\pi_\mathcal{E}]$ is bounded by $4p^2$, so the conductor of $\mathbb{Z}[\pi_\mathcal{E}]$ in

O_K is bounded by $2p$; hence, if $\mathcal{E}$ is ordinary, then the maximal height of the 2-isogeny volcano containing $\mathcal{E}$ is $\log_2(p) + 1$. Sutherland's supersingularity test takes random non-backtracking 2-isogeny walks starting from each of the three 2-isogeny neighbours of $\mathcal{E}$. If $\mathcal{E}$ is ordinary, then at least one of these walks will descend the 2-isogeny volcano, and will therefore terminate (with no non-backtracking step defined over $\mathbb{F}_{p^2}$) after at most $\log_2(p) + 1$ steps. Conversely, if no walk terminates after $\log_2(p) + 1$ steps, then $\mathcal{E}$ must be supersingular.

In our case, we know that $\mathrm{End}(\mathcal{E}) \supset \mathbb{Z}[\mu] \supset \mathbb{Z}[\pi_\mathcal{E}]$, and the conductor of $\mathbb{Z}[\pi_\mathcal{E}]$ in $\mathbb{Z}[\mu]$ is the integer $|r|$ of Proposition 1, which is bounded by $2\sqrt{p/d}$. We can therefore reduce the walk length limit from $\log_2(p) + 1$ to $\frac{1}{2}(\log_2(p) - \log_2(d)) + 1$. We can also use the fact that $\mathbb{Z}[\mu] \subset \mathrm{End}(\mathcal{E})$ to ensure that we choose a "descending" path within at most two steps, and omit the other two paths. Thus, we can determine if a (d, ϵ)-structure $(\mathcal{E}, \psi)$ is supersingular for the cost of computing two d-isogenies and $(\frac{1}{2}(\log_2(p) - \log_2(d)) + 5)$ 2-isogenies.

We can determine whether $(\mathcal{E}, \psi)$ is in $\mathcal{D}^{\max}_{d,\epsilon}(p)$ or $\mathcal{D}^{\mathrm{sub}}_{d,\epsilon}(p)$ (if required, and only if $-dp \equiv 1 \pmod 4$) by computing the action of μ on the 2-torsion (at the cost of one or two d-isogeny evaluations) or by computing the 2-neighbours of $(\mathcal{E}, \psi)$ in $\Gamma_2(\mathcal{D}_{d,\epsilon}(p))$.

6.4 GENERALIZED DELFS–GALBRAITH ALGORITHMS

Let S_p be the set of supersingular curves over $\mathbb{F}_{p^2}$, up to isomorphism. The general supersingular isogeny problem is, given $\mathcal{E}_1$ and $\mathcal{E}_2$ in S_p, to compute an isogeny $\phi : \mathcal{E}_1 \to \mathcal{E}_2$.

In [21], Delfs and Galbraith use the subset of supersingular curves defined over $\mathbb{F}_p$, which we can identify with $\mathcal{D}_{1,1}(p)$, to improve classical isogeny-finding algorithms based on random walks. Their algorithm has two phases:

1. Compute a random non-backtracking isogeny walk from $\mathcal{E}_1$ resp. $\mathcal{E}_2$ until we land on a curve $\mathcal{E}'_1$ resp. $\mathcal{E}'_2$ in $\mathcal{D}_{1,1}(p)$. These walks yield isogenies $\phi_1 : \mathcal{E}_1 \to \mathcal{E}'_1$ and $\phi_2 : \mathcal{E}_2 \to \mathcal{E}'_2$. The isogeny graph on S_p has excellent mixing properties, and since $\#S_p \approx p/12$ and $\#\mathcal{D}_{1,1}(p) = O(\sqrt{p})$, this first phase takes an expected $O(\sqrt{p})$ random isogeny steps.
2. Find an isogeny $\phi' : \mathcal{E}'_1 \to \mathcal{E}'_2$ using the action of $\mathrm{Cl}(\mathbb{Q}(\sqrt{-p}))$ acting on $\mathcal{D}_{1,1}(p)$ (that is, solve Vectorization with $d = 1$). Under the Generalized Riemann Hypothesis, $\mathrm{Cl}(\mathbb{Q}(\sqrt{-p}))$ is generated by the set $\mathcal{L}$ of ideals of prime norm up to $6\log(|\Delta|)^2$, where Δ is the discriminant of $\mathbb{Q}(\sqrt{-p})$ (see [2]) though in practice we do not need so many primes. The $\mathcal{L}$-isogeny graph on $\mathcal{D}_{1,1}(p)$ is therefore connected, and we can use random walks in this subgraph to construct ϕ'. By the birthday paradox, this phase takes an expected $O(\sqrt[4]{p})$ random steps before finding the collision yielding ϕ'.

The Delfs–Galbraith algorithm exploits the action of $\mathrm{Cl}(\mathbb{Q}(\sqrt{-p}))$ on $\mathcal{D}_{1,1}(p)$ to solve the isogeny problem in S_p. We can generalize their algorithm by replacing the distinguished subgraph $\Gamma(\mathcal{D}_{1,\epsilon}(p))$ with a union of subgraphs $\sqcup_{d\in D}\Gamma(\mathcal{D}_{d,\epsilon}(p))$ where D is a set of coprime squarefree integers prime to p. In Phase 1, we now take random walks from $\mathcal{E}_1$ and $\mathcal{E}_2$ into $\sqcup_{d\in D}\mathcal{D}_{d,\epsilon}(p)$.[3] In Phase 2, if $\mathcal{E}'_1$ is in $\mathcal{D}_{d_1,\epsilon}(p)$ and $\mathcal{E}'_2$ is in $\mathcal{D}_{d_2,\epsilon}(p)$, then we need to compute a (d_1, d_2)-crossroad $\mathcal{E}'_3$ and find a path $\mathcal{E}'_1 \to \mathcal{E}'_3$ in $\mathcal{D}_{d_1,\epsilon}(p)$ and a path $\mathcal{E}'_2 \to \mathcal{E}'_3$ in $\mathcal{D}_{d_2,\epsilon}(p)$. (In particular, we should ensure that there exist supersingular (d_1, d_2)-crossroads before including d_1 and d_2 in D.)

This is not worthwhile for large d or large D. Asymptotically, $\#\mathcal{D}_{d,\epsilon}(p)$ is in $O((\sum_{d\in D}\sqrt{d})\sqrt{p})$, so the expected number of steps in Phase 1 is reduced by a factor of $O(\sum_{d\in D}\sqrt{d})$. However, the individual steps become more expensive: if we use modular polynomials to check membership of each $\mathcal{D}_{d,\epsilon}(p)$, then the number of $\mathbb{F}_{p^2}$-operations per step grows linearly with $\sum_{d\in D} d$, overwhelming the benefit of the shorter walks. Asymptotically, therefore, there is no benefit in taking large d or large D in Phase 1. (For more analysis of random walks into $(d, \pm 1)$-structures, in different contexts, see [22] and [13].)

Generalized Delfs–Galbraith can become interesting for D consisting of a few small d, however, precisely because the asymptotic $\kappa(d, p) := \#\mathcal{D}_{d,\epsilon}(p)/\#\mathcal{D}_{1,\epsilon}(p) \approx \sqrt{d}$ no longer holds. For $d < 10$, for example, we can have $\kappa(d, p)$ substantially greater than $\sqrt{d}$ (and also substantially less than 1). For example, if p is the toy SIDH-type prime $2^{52} \cdot 3^{33} - 1$, then $\kappa(5, p) \approx 4.916$. If we can test for an isomorphism or 5-isogeny to the conjugate faster than we can compute six 2-isogenies, then we can take $D = \{1, 5\}$ and walk into $\mathcal{D}_{1,\epsilon}(p) \sqcup \mathcal{D}_{5,\epsilon}(p)$ faster than walking into $\mathcal{D}_{1,\epsilon}(p)$ alone. This speedup is counterbalanced by a slowdown in Phase 2, because walking in $\Gamma(\mathcal{D}_{5,\epsilon}(p))$ costs more, and because the walks there need to be a square-root of $\kappa(5, p)$ longer—though we can work modulo conjugation to mitigate this cost.

6.5 (d, ϵ)-STRUCTURES AND SIDH GRAPHS

As we noted above, the probability of a random walk in the supersingular ℓ-isogeny graph hitting a vertex that is the image of a (d, ϵ)-structure is very low. It is even lower when we consider SIDH/SIKE graphs, which cover

[3] To measure the feasability of this attack, we need to estimate the average number of steps from a general supersingular elliptic curve $\mathcal{E}/\mathbb{F}_{p^2}$ to a curve in (the image of) $\mathcal{D}_{d,\epsilon}(p)$. This distance follows a binomial law $(m, \mathcal{P})$ where m is the number of steps and $\mathcal{P} = \sqrt{d/p}$. Hence, the probability $\mathbb{P}(X > 1)$ that we reach at least one element in $\mathcal{D}_{d,\epsilon}(p)$ after m steps from $\mathcal{E}$ is $\mathbb{P}(X > 1) = 1 - \mathbb{P}(X = 0) = 1 - (1 - \sqrt{d/p})^m$. When $d = 1$, this addresses some of the heuristic observations in [1], notably the distance to the $\mathbb{F}_p$-spine.

only a very small proportion of the full isogeny graph, resembling trees of walks of short, fixed length.

Nevertheless, when we look at specific SIKE graphs, we see that they contain sections of $\Gamma_2(\mathcal{D}_{d,\epsilon}(p))$ and $\Gamma_3(\mathcal{D}_{d,\epsilon}(p))$ for various d. For example, the starting curve in `SIKEp434` has a d-isogeny to its conjugate for $d \in D = \{5, 13, 17, 29, 37, 41\}$ (and also for much higher, but less practical values of d). If we consider the 2-isogeny graph, then we find that $\Gamma_2(\mathcal{D}_{d,\epsilon}(p))$ passes through the starting curve and continues down through the tree towards a public key for $d = 17$ and 41. Hence, if we can find a 2-isogeny path from a `SIKEp434` public key to a vertex in the image of $\mathcal{D}_{17,\epsilon}(p)$ or $\mathcal{D}_{41,\epsilon}(p)$, then we have an express route to the starting curve. Such an attack succeeds in a reasonable time with only a very small probability, but it is still devastatingly effective for a tiny proportion of `SIKEp434` keys.

REFERENCES

[1] Sarah Arpin et al. "Adventures in Supersingularland". In: *Experimental Mathematics* (to appear). URL: https://eprint.iacr.org/2019/1056.

[2] Eric Bach. *Analytic methods in the analysis and design of number-theoretic algorithms*. MIT Press, Cambridge MA, 1984.

[3] Gustavo Banegas et al. "CTIDH: faster constant-time CSIDH". In: *Transactions on Cryptographic Hardware and Embedded Systems* (to appear). URL: https://ia.cr/2021/633.

[4] Daniel J. Bernstein, Tanja Lange, and Peter Schwabe. "On the Correct Use of the Negation Map in the Pollard rho Method". In: *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings*. Ed. by Dario Catalano et al. Vol. 6571. Lecture Notes in Computer Science. Springer, 2011, pp. 128–146. DOI: 10.1007/978-3-642-19379-8_8.

[5] Daniel J. Bernstein et al. "Quantum Circuits for the CSIDH: Optimizing Quantum Evaluation of Isogenies". In: *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part II*. Ed. by Yuval Ishai and Vincent Rijmen. Vol. 11477. Lecture Notes in Computer Science. Springer, 2019, pp. 409–441. DOI: 10.1007/978-3-030-17656-3_15.

[6] Daniel J. Bernstein et al. "Faster computation of isogenies of large prime degree". In: *Fourteenth Algorithmic Number Theory Symposium*. Ed. by Steven D. Galbraith. Vol. 4. Open book series. Mathematical Sciences Publishers, 2020, pp. 39–55. URL: https://doi.org/10.2140/obs.2020.4.39.

[7] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. "CSI-FiSh: Efficient Isogeny Based Signatures Through Class Group Computations". In: *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part I*. Ed. by Steven D. Galbraith and Shiho Moriai. Vol. 11921. Lecture Notes in Computer Science. Springer, 2019, pp. 227–247. DOI: 10.1007/978-3-030-34578-5_9.

[8] Xavier Bonnetain and André Schrottenloher. "Quantum Security Analysis of CSIDH". In: *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part II*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12106. Lecture Notes in Computer Science. Springer, 2020, pp. 493–522. DOI: 10.1007/978-3-030-45724-2_17.

[9] Fabio Campos et al. "Trouble at the CSIDH: Protecting CSIDH with Dummy-Operations Against Fault Injection Attacks". In: *17th Workshop on Fault Detection and Tolerance in Cryptography, FDTC 2020, Milan, Italy, September 13, 2020*. IEEE, 2020, pp. 57–65. DOI: 10.1109/FDTC51366.2020.00015.

[10] Wouter Castryck and Thomas Decru. "CSIDH on the Surface". In: *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020, Paris, France, April 15-17, 2020, Proceedings*. Ed. by Jintai Ding and Jean-Pierre Tillich. Vol. 12100. Lecture Notes in Computer Science. Springer, 2020, pp. 111–129. DOI: 10.1007/978-3-030-44223-1_7.

[11] Wouter Castryck et al. "CSIDH: An Efficient Post-Quantum Commutative Group Action". In: *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part III*. Ed. by Thomas Peyrin and Steven D. Galbraith. Vol. 11274. Lecture Notes in Computer Science. Springer, 2018, pp. 395–427. DOI: 10.1007/978-3-030-03332-3_15.

[12] Daniel Cervantes-Vázquez et al. "Stronger and Faster Side-Channel Protections for CSIDH". In: *Progress in Cryptology - LATINCRYPT 2019 - 6th International Conference on Cryptology and Information Security in Latin America, Santiago de Chile, Chile, October 2-4, 2019, Proceedings*. Ed. by Peter Schwabe and Nicolas Thériault. Vol. 11774. Lecture Notes in Computer Science. Springer, 2019, pp. 173–193. DOI: 10.1007/978-3-030-30530-7_9.

[13] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. "Cryptographic hash functions from expander graphs". In: *Journal of Cryptology* 22.1 (2009), pp. 93–113.

[14] Jorge Chávez-Saab et al. *The SQALE of CSIDH: square-root Vélu quantum-resistant isogeny action with low exponents*. Cryptology ePrint Archive, Report 2020/1520. 2020.

[15] Andrew Childs, David Jao, and Vladimir Soukharev. "Constructing elliptic curve isogenies in quantum subexponential time". In: *Journal of Mathematical Cryptology* 8.1 (2014), pp. 1–29.

[16] Leonardo Colò and David Kohel. "Orienting supersingular isogeny graphs". In: *Journal of Mathematical Cryptolology* 14.1 (2020), pp. 414–437. DOI: 10.1515/jmc-2019-0034.

[17] Craig Costello and Benjamin Smith. "The Supersingular Isogeny Problem in Genus 2 and Beyond". In: *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020, Paris, France, April 15-17, 2020, Proceedings*. Ed. by Jintai Ding and Jean-Pierre Tillich. Vol. 12100. Lecture Notes in Computer Science. Springer, 2020, pp. 151–168. DOI: 10.1007/978-3-030-44223-1_9.

[18] Jean-Marc Couveignes. *Hard Homogeneous Spaces*. 2006. URL: http://eprint.iacr.org/2006/291.

[19] David A. Cox. *Primes of the Form $x^2 + ny^2$: Fermat, Class Field Theory, and Complex Multiplication*. 2nd. Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts. John Wiley and Sons, 2013. DOI: 10.1002/9781118400722.

[20] Luca De Feo, Jean Kieffer, and Benjamin Smith. "Towards Practical Key Exchange from Ordinary Isogeny Graphs". In: *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part III*. Ed. by Thomas Peyrin and Steven D. Galbraith. Vol. 11274. Lecture Notes in Computer Science. Springer, 2018, pp. 365–394. DOI: 10.1007/978-3-030-03332-3_14.

[21] Christina Delfs and Steven D. Galbraith. "Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$". In: *Designs, Codes and Cryptography* 78.2 (2016), pp. 425–440. DOI: 10.1007/s10623-014-0010-1.

[22] Kirsten Eisenträger et al. "Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs". In: *Fourteenth Algorithmic Number Theory Symposium*. Ed. by Steven D. Galbraith. Vol. 4. Open book series. Mathematical Sciences Publishers, 2020, pp. 215–232. DOI: https://doi.org/10.2140/obs.2020.4.215.

[23] Luca De Feo, David Jao, and Jérôme Plût. "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies". In: *Journal of Mathematical Cryptology* 8.3 (2014), pp. 209–247. DOI: 10.1515/jmc-2012-0015.

[24] Luca De Feo et al. "SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies". In: *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part I*. Ed. by Shiho Moriai and Huaxiong Wang. Vol. 12491. Lecture Notes in Computer Science. Springer, 2020, pp. 64–93. DOI: 10.1007/978-3-030-64837-4_3.

[25] Steven Galbraith et al. "Quantum Equivalence of the DLP and CDHP for Group Actions". In: *Mathematical Cryptology* 1.1 (2021), pp. 40–44. URL: https://eprint.iacr.org/2018/1199.

[26] Steven D. Galbraith, Xibin Lin, and Michael Scott. "Endomorphisms for Faster Elliptic Curve Cryptography on a Large Class of Curves". In: *Journal of Cryptology* 24.3 (2011), pp. 446–469. DOI: 10.1007/s00145-010-9065-y.

[27] Steven D. Galbraith et al. "On the Security of Supersingular Isogeny Cryptosystems". In: *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I*. Ed. by Jung Hee Cheon and Tsuyoshi Takagi. Vol. 10031. Lecture Notes in Computer Science, pp. 63–91. DOI: 10.1007/978-3-662-53887-6_3.

[28] Aurore Guillevic and Sorina Ionica. "Four-Dimensional GLV via the Weil Restriction". In: *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*. Ed. by Kazue Sako and Palash Sarkar. Vol. 8269. Lecture Notes in Computer Science. Springer, 2013, pp. 79–96. DOI: 10.1007/978-3-642-42033-7_5.

[29] Yuji Hasegawa. "Q-curves over quadratic fields". In: *Manuscripta Mathematica* 94.1 (1997), pp. 347–364. ISSN: 1432-1785. DOI: 10.1007/BF02677859.

[30] David Jao and Luca De Feo. "Towards Quantum-Resistant Cryptosystems from Supersingular Elliptic Curve Isogenies". In: *Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011, Taipei, Taiwan, November 29 - December 2, 2011. Proceedings*. Ed. by Bo-Yin Yang. Vol. 7071. Lecture Notes in Computer Science. Springer, 2011, pp. 19–34. DOI: 10.1007/978-3-642-25405-5_2.

[31] David Jao et al. *SIKE – Supersingular Isogeny Key Encapsulation*. URL: https://sike.org/.

[32] Greg Kuperberg. "A subexponential-time quantum algorithm for the dihedral hidden subgroup problem". In: *SIAM Journal of Computing* 35.1 (2005), pp. 170–188. eprint: quant-ph/0302112.

[33] Greg Kuperberg. "Another Subexponential-time Quantum Algorithm for the Dihedral Hidden Subgroup Problem". In: *8th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2013)*. Ed. by Simone Severini and Fernando Brandao. Vol. 22. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013, pp. 20–34. ISBN: 978-3-939897-55-2. DOI: 10.4230/LIPIcs.TQC.2013.20. URL: http://drops.dagstuhl.de/opus/volltexte/2013/4321.

[34] Serge Lang. *Algebraic Number Theory*. Vol. 110. Graduate Texts in Mathematics. Springer, 1994.

[35] Jonathan Love and Dan Boneh. "Supersingular Curves With Small Non-integer Endomorphisms". In: *Fourteenth Algorithmic Number Theory Symposium*. Ed. by Steven D. Galbraith. Vol. 4. Open book series. Mathematical Sciences Publishers, 2020, pp. 7–22. DOI: https://doi.org/10.2140/obs.2020.4.7.

[36] Michael Meyer, Fabio Campos, and Steffen Reith. "On Lions and Elligators: An Efficient Constant-Time Implementation of CSIDH". In: *Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019, Chongqing, China, May 8-10, 2019 Revised Selected Papers*. Ed. by Jintai Ding and Rainer Steinwandt. Vol. 11505. Lecture Notes in Computer Science. Springer, 2019, pp. 307–325. DOI: 10.1007/978-3-030-25510-7_17.

[37] François Morain, Charlotte Scribot, and Benjamin Smith. "Computing cardinalities of $\mathbb{Q}$-curve reductions over finite fields". In: *LMS Journal of Computation and Mathematics* 19.A (Aug. 2016), p. 15. DOI: 10.1112/S1461157016000267. URL: https://hal.inria.fr/hal-01320388.

[38] Hiroshi Onuki. "On oriented supersingular elliptic curves". In: *Finite Fields and their Applications* 69 (2021), p. 101777. DOI: 10.1016/j.ffa.2020.101777.

[39] Hiroshi Onuki and Tsuyoshi Takagi. "On Collisions Related to an Ideal Class of Order 3 in CSIDH". In: *Advances in Information and Computer Security*. Ed. by Kazumaro Aoki and Akira Kanaoka. Cham: Springer International Publishing, 2020, pp. 131–148. ISBN: 978-3-030-58208-1.

[40] Hiroshi Onuki et al. "A Constant-Time Algorithm of CSIDH Keeping Two Points". In: *IEICE Transactions on Fundamentals of Electronics, Communications, and Computer Science* 103-A.10 (2020), pp. 1174–1182. DOI: 10.1587/transfun.2019DMP0008.

[41] Chris Peikert. "He Gives C-Sieves on the CSIDH". In: *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part II*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12106. Lecture Notes in Computer Science. Springer, 2020, pp. 463–492. DOI: 10.1007/978-3-030-45724-2_16.

[42] Oded Regev. *A Subexponential Time Algorithm for the Dihedral Hidden Subgroup Problem with Polynomial Space*. arXiv:quant-ph/0406151. June 2004. URL: http://arxiv.org/abs/quant-ph/0406151.

[43] Alexander Rostovtsev and Anton Stolbunov. "Public-Key Cryptosystem Based on Isogenies". In: *IACR Cryptology ePrint Archive* 2006 (2006), p. 145. URL: http://eprint.iacr.org/2006/145.

[44] René Schoof. "Counting points on elliptic curves over finite fields". In: *Journal de Théorie des Nombres de Bordeaux* 7.1 (1995), pp. 219–254.

[45] Benjamin Smith. "The $\mathbb{Q}$-curve Construction for Endomorphism-Accelerated Elliptic Curves". In: *Journal of Cryptology* 29.4 (2016), pp. 806–832. DOI: 10.1007/s00145-015-9210-8.

[46] Anton Stolbunov. "Reductionist Security Arguments for Public-Key Cryptographic Schemes Based on Group Action". In: *Norsk informasjonssikkerhetskonferanse (NISK)*. Ed. by Stig F. Mjølsnes. 2009.

[47] Anton Stolbunov. "Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves". In: *Advances in Mathematics of Communications* 4.2 (2010).

[48] Andrew V. Sutherland. "Identifying supersingular elliptic curves". In: *LMS Journal of Computation and Mathematics* 15 (2012), pp. 317–325. DOI: 10.1112/S1461157012001106.