The Supersingular Endomorphism Ring and One Endomorphism Problems are Equivalent

Aurel Page[1(✉)] and Benjamin Wesolowski[2]

[1] Univ. Bordeaux, CNRS, INRIA, Bordeaux INP, IMB, UMR 5251, 33400 Talence, France
aurel.page@inria.fr85

[2] ENS de Lyon, CNRS, UMPA, UMR 5669, Lyon, France

Abstract. The supersingular Endomorphism Ring problem is the following: given a supersingular elliptic curve, compute all of its endomorphisms. The presumed hardness of this problem is foundational for isogeny-based cryptography. The One Endomorphism problem only asks to find a single non-scalar endomorphism. We prove that these two problems are equivalent, under probabilistic polynomial time reductions.

We prove a number of consequences. First, assuming the hardness of the endomorphism ring problem, the Charles–Goren–Lauter hash function is collision resistant, and the SQIsign identification protocol is sound for uniformly random keys. Second, the endomorphism ring problem is equivalent to the problem of computing arbitrary isogenies between supersingular elliptic curves, a result previously known only for isogenies of smooth degree. Third, there exists an unconditional probabilistic algorithm to solve the endomorphism ring problem in time $\tilde{O}(p^{1/2})$, a result that previously required to assume the generalized Riemann hypothesis.

To prove our main result, we introduce a flexible framework for the study of isogeny graphs with additional information. We prove a general and easy-to-use rapid mixing theorem.

1 Introduction

The endomorphism ring problem lies at the foundation of isogeny-based cryptography. On one hand, its presumed hardness is necessary for the security of all cryptosystems of this family (see for instance the reductions in [Wes22a]). On the other hand, many cryptosystems of this family can be proven secure if this problem (or some variant) is hard (the earliest example being [CLG09]). Isogeny-based cryptography takes its name from the *isogeny problem*. An isogeny is a certain kind of map between two elliptic curves, and the isogeny problem consists in finding such a map, given the two curves. Formalising the meaning of "finding an isogeny" can lead to several versions of the isogeny problem, the most prominent being the *ℓ-isogeny path problem*. In isogeny-based cryptography, one typically restricts to supersingular elliptic curves, for which this problem is believed to be hard.

M. Joye and G. Leander (Eds.): EUROCRYPT 2024, LNCS 14656, pp. 388–417, 2024.
https://doi.org/10.1007/978-3-031-58751-1_14

Fix a supersingular elliptic curve E. An endomorphism of E is an isogeny from E to itself (or the zero morphism). The collection of all endomorphisms of E forms the endomorphism ring $\mathrm{End}(E)$. The *supersingular endomorphism ring problem*, or EndRing, consists in computing $\mathrm{End}(E)$, when given E. Assuming the generalised Riemann hypothesis, this problem is equivalent to the ℓ-isogeny path problem (see [Wes22b], and the earlier heuristic equivalence [EHL+18]), cementing its importance in the field.

The endomorphism ring contains scalars $\mathbf{Z} \subseteq \mathrm{End}(E)$, simple elements which are always easy to compute. While EndRing asks to find all endomorphisms, it has long been believed that finding even a single non-scalar endomorphism is hard. We call this the *one endomorphism problem*, or OneEnd. Unfortunately, former heuristic arguments suggesting that OneEnd should be as hard as EndRing do not withstand close scrutiny, and actually fail in simple cases. Yet, the connection between these two problems bears important consequences on the hardness of EndRing, on its connection with variants of the isogeny problem, and on the security of cryptosystems such as the CGL hash function [CLG09] or the SQIsign digital signature scheme [DKL+20].

1.1 Contributions

In this article, we prove the following theorem.

Theorem 1.1. *The* EndRing *and* OneEnd *problems are equivalent, under probabilistic polynomial time reductions.*

Formal definitions are provided in Sect. 2, and the proof is the object of Sect. 7. The reduction from OneEnd to EndRing is obvious, and the other direction is stated more precisely in Theorem 7.2. This reduction transforms one instance of EndRing into polynomially many instances of OneEnd.

As a consequence of the main theorem, we prove the following:

- If EndRing is hard, then the CGL hash function is collision resistant (Theorem 8.1), and the SQIsign identification scheme is sound when the keys are indistinguishable from uniform (Theorem 8.2, expressed for RigorousSQIsignHD [DLRW23], a variant for which keys are proved to be statistically indistinguishable from uniform). Previous security proofs relied on the hardness of OneEnd (see [DKL+20, Theorem 1]), or on flawed heuristic reductions (see [EHL+18, Algorithm 8], and the flaws discussed Sect. 1.2). This is the object of Sect. 8.1 and Sect. 8.2.
- EndRing reduces to the isogeny problem (Theorem 8.6). Here, the isogeny problem refers to the problem of finding *any* isogeny between two elliptic curves. Previous results [EHL+18, Wes22b] only applied to isogenies of smooth degree (like the ℓ-isogeny path problem), and were conditional on the generalised Riemann hypothesis. This is the object of Sect. 8.3.
- There is an algorithm solving EndRing in expected time $\tilde{O}(p^{1/2})$ (Theorem 8.8), where $p > 0$ is the characteristic. Previous algorithms were conditional on the generalised Riemann hypothesis (via the conditional equivalence

with the ℓ-isogeny path problem [Wes22b]; see also [FIK+23, Theorem 5.7] for a more direct approach). Previous unconditional algorithms ran in time $\tilde{O}(p)$ and only returned a full-rank subring [Koh96, Theorem 75]. This is the object of Sect. 8.4.

Our main technical tool is an equidistribution result for isogeny walks in the graph of supersingular elliptic curves equipped with an endomorphism modulo N. In fact, we prove a more general equidistribution result generalising the classical one (see [Mes86, Piz90] and Proposition 2.7), which we think is of independent interest. We state this result informally here, referring the reader to the body of the paper for a formal statement.

Definition 1.2. *Equipping the set of supersingular elliptic curves with extra data consists in defining for each such curve E a finite set $\mathcal{F}(E)$, and for every isogeny $\varphi\colon E \to E'$ a map $\mathcal{F}(\varphi)\colon \mathcal{F}(E) \to \mathcal{F}(E')$, compatible under composition of isogenies (see Definitions 2.8 and 3.1). We obtain the isogeny graph $\mathcal{G}_{\mathcal{F}}$ of pairs (E, x) where $x \in \mathcal{F}(E)$ (see Definition 3.4).*

Let $N \geq 1$ be an integer. The extra data satisfies the (mod N)-congruence property *if for every curve E, pairs of endomorphisms of E that are congruent modulo N act identically on $\mathcal{F}(E)$ (see Definition 3.7).*

Our equidistribution result, stated informally, reads as follows.

Theorem 1.3. *Let $N \geq 1$ be an integer. Random walks in the isogeny graph of supersingular elliptic curves equipped with extra data satisfying the* (mod N)*-congruence property equidistribute optimally.*

We refer to Theorem 3.10 for a formal statement. The optimality refers to the fact that the graphs can be disconnected or multipartite, resulting in the adjacency matrix having several forced eigenvalues (see Proposition 3.11 and Remark 3.12), but all the remaining eigenvalues are as small as possible. A similar general result was recently proved by Codogni and Lido [CL23], so we point out some similarities and differences. In [CL23], the extra data needs to be expressed in terms of N-torsion points (a *level structure*), whereas we allow for extra data of arbitrary nature, only requiring it to satisfy a simple property (the (mod N)-congruence property). We hope that this makes our theorem flexible, and easy to use in a variety of situations. In particular, the extra data used in our main application trivially fits within our framework; in contrast, this data is not a level structure, so does not directly fit the framework of [CL23]. Moreover, we allow p to divide N, contrary to the results in [CL23]. Both proofs use Deligne's bounds, but the proof in [CL23] is purely algebro-geometric, whereas ours proceeds via the Deuring correspondence and the Jacquet–Langlands correspondence; as a result, the two proofs could have different interesting generalisations.

1.2 Technical Overview

The ideas behind our reduction are as follows. Assume we have an oracle $\mathscr{O}$ for OneEnd and we want to compute $\mathrm{End}(E)$ for a given E.

The ring $\mathrm{End}(E)$ is a lattice of dimension 4 and volume $p/4$. Computing $\mathrm{End}(E)$ consists in finding a basis: four endomorphisms that generate all the others. Given a collection of endomorphisms, one can compute the volume of the lattice they generate, and easily check whether they generate $\mathrm{End}(E)$.

A First Flawed Attempt. We thus need a way to generate several endomorphisms of E. Naively, one could repeatedly call $\mathscr{O}(E)$, hoping to eventually obtain a generating set. This can fail, for instance if the oracle is deterministic and $\mathscr{O}(E)$ always returns the same endomorphism.

To circumvent this issue, it was proposed in [EHL+18] to randomise the curve. More precisely, one constructs a richer, randomised oracle $\textsc{Rich}^{\mathscr{O}}$ from $\mathscr{O}$ as follows. On input E, walk randomly on the 2-isogeny graph, resulting in an isogeny $\varphi\colon E \to E'$. This graph has rapid mixing properties, so E' is close to uniformly distributed among supersingular curves. Now, call the oracle $\mathscr{O}$ on E', to get an endomorphism $\beta \in \mathrm{End}(E')$. The composition $\alpha = \hat{\varphi} \circ \beta \circ \varphi$ is an endomorphism of E, the output of $\textsc{Rich}^{\mathscr{O}}$.

With this randomisation, there is hope that calling $\textsc{Rich}^{\mathscr{O}}$ repeatedly on E could yield several independent endomorphisms that would eventually generate $\mathrm{End}(E)$. This method is essentially [EHL+18, Algorithm 8]. In that article, it is heuristically assumed that endomorphisms produced by $\textsc{Rich}^{\mathscr{O}}$ are very nicely distributed, and they deduce that a generating set for $\mathrm{End}(E)$ is rapidly obtained. This heuristic has a critical flaw: one can construct oracles that contradict it. Consider an integer $M > 1$, and suppose that for any input E, the oracle $\mathscr{O}$ returns an endomorphism from the strict subring $\mathbf{Z} + M\,\mathrm{End}(E)$. Then, the above algorithm would fail, because the randomisation $\textsc{Rich}^{\mathscr{O}}$ would still be stuck within the subring $\mathbf{Z} + M\,\mathrm{End}(E)$. Worse, juggling with several related integers M, we will see that there are oracles for which this algorithm only stabilises after an exponential time.

Identifying and Resolving Obstructions. The core of our method rests on the idea that this issue is, in essence, the only possible obstruction. The key is *invariance by conjugation*. If $\varphi, \varphi' : E \to E'$ are two random walks of the same length, and β is an endomorphism of $\mathrm{End}(E')$, the elements $\alpha = \hat{\varphi} \circ \beta \circ \varphi$ and $\alpha' = \hat{\varphi}' \circ \beta \circ \varphi'$ are equally likely outputs of $\textsc{Rich}^{\mathscr{O}}$. These two elements are conjugates of each other in $\mathrm{End}(E)/N\,\mathrm{End}(E)$ for any odd integer N, as

$$\alpha = \frac{\hat{\varphi} \circ \varphi'}{[\deg(\varphi')]} \circ \alpha' \circ \frac{\hat{\varphi}' \circ \varphi}{[\deg(\varphi')]} \bmod N.$$

From there, one can prove that the output of $\textsc{Rich}^{\mathscr{O}}$ follows a distribution that is invariant by conjugation: each output is as likely as any of its conjugates, modulo odd integers N (up to some bound). Intuitively, for the outputs of $\textsc{Rich}^{\mathscr{O}}$ to be "stuck" in a subring (such as $\mathbf{Z} + M\,\mathrm{End}(E)$ above), that subring must itself be stable by conjugation (modulo odd integers N). There comes the next key: every subring of $\mathrm{End}(E)$ (of finite index not divisible by p) stable by conjugation

modulo all integers is of the form $\mathbf{Z} + M\,\mathrm{End}(E)$. From a basis of $\mathbf{Z} + M\,\mathrm{End}(E)$, it is easy to recover a basis of $\mathrm{End}(E)$ essentially by dividing by M (using a method due to Robert [Rob22] that stems from the attacks on SIDH).

This intuition does not immediately translate into an algorithm, as an oracle could be "bad" without really being stuck in a subring. Imagine an oracle that outputs an element of $\mathbf{Z} + 2^e\,\mathrm{End}(E)$ (and not in $\mathbf{Z} + 2^{e+1}\,\mathrm{End}(E)$) with probability 2^{e-n} for each $e \in [0, \ldots, n-1]$, for some integer n. A sequence of samples $(\alpha_i)_i$ could eventually generate $\mathrm{End}(E)$, but only after an amount of time exponential in n. This particular case could be resolved as follows: for each sample α, identify the largest e such that $\beta = (2\alpha - \mathrm{Tr}(\alpha))/2^e$ is an endomorphism. A sequence of samples $(\beta_i)_i$ could rapidly generate $\mathbf{Z} + 2\,\mathrm{End}(E)$, from which one easily recovers $\mathrm{End}(E)$. This resolution first identifies the prime 2 as the source of the obstruction, then "reduces" each sample "at 2". In general, such obstructive primes would appear as factors of $\mathrm{disc}(\alpha)$. Identifying these primes, and ensuring that each sample is "reduced" at each of them, one gets, in principle, a complete algorithm. However, factoring $\mathrm{disc}(\alpha)$ could be hard. Instead, we implement an optimistic approach: we identify obstructive pseudo-primes using a polynomial time partial-factoring algorithm. The factors may still be composite, but it is fine: the algorithm will either behave as if they were prime, or reveal a new factor.

Equidistribution in Isogeny Graphs. The technical core of our result is the proof that the distribution of $\textsc{Rich}^{\mathscr{O}}$ is indeed invariant by conjugation. It is a consequence of Theorem 1.3, our general equidistribution result, whose proof proceeds as follows. We use a categorical version of the Deuring correspondence to bring everything to the quaternion world. We then use a technical result to show that extra data satisfying the congruence property yield graphs isomorphic to special ones constructed from quaternionic groups. Finally, these special graphs are directly related to automorphic forms, so we can apply the Jacquet–Langlands correspondence and Deligne's bounds on coefficients of modular forms. The resulting bounds on the adjacency operators is the desired result.

2 Preliminaries

2.1 Notation

We write $\mathbf{Z}$, $\mathbf{Q}$, $\mathbf{R}$ and $\mathbf{C}$ for the ring of integers, the fields of rational, real, and complex numbers. For any prime ℓ, we write $\mathbf{Z}_\ell$ and $\mathbf{Q}_\ell$ for the ring of ℓ-adic integers and the field of ℓ-adic numbers. For any prime power q, we write $\mathbf{F}_q$ for the finite field with q elements. For any field K, we write $\overline{K}$ for its algebraic closure. For any set S, we write $\#S$ for its cardinality. We write $f = O(g)$ for the classic big O notation, and equivalently $g = \Omega(f)$ for the classic Ω notation. We also write $f = \Theta(g)$ if we have both $f = O(g)$ and $f = \Omega(g)$. We use the soft O notation $\tilde{O}(g) = \log(g)^{O(1)} \cdot O(g)$. We also write $\mathrm{poly}(f_1, \ldots, f_n) =$

$(f_1 + \cdots + f_n)^{O(1)}$. The logarithm function log is in base 2. For any ring R, we write $R^\times$ the multiplicative group of invertible elements, and $M_2(R)$ the ring of 2×2 matrices with coefficients in R.

2.2 Quaternion Algebras

A general reference for this section is [Voi21]. A *quaternion algebra over* $\mathbf{Q}$ is a ring B having a $\mathbf{Q}$-basis $1, i, j, k$ satisfying the multiplication rules $i^2 = a$, $j^2 = b$ and $k = ij = -ji$, for some $a, b \in \mathbf{Q}^\times$. Let $w = x + yi + zj + tk \in B$. The *reduced trace* of w is $\mathrm{trd}(w) = 2x$. The *reduced norm* of w is $\mathrm{nrd}(w) = x^2 - ay^2 - bz^2 - abt^2$. The reduced norm map is multiplicative. A *lattice* in a $\mathbf{Q}$-vector space V of finite dimension d is a subgroup $L \subset V$ of rank d over $\mathbf{Z}$ and such that $V = L\,\mathbf{Q}$. The *discriminant* of a lattice L in B is $\mathrm{disc}(L) = \det(\mathrm{trd}(b_i b_j)) \neq 0$ where (b_i) is a $\mathbf{Z}$-basis of L. When $L' \subset L$ is a sublattice, we have $\mathrm{disc}(L') = [L : L']^2 \,\mathrm{disc}(L)$. An *order* in B is a subring $\mathcal{O} \subset B$ that is also a lattice. A *maximal order* is an order that is not properly contained in another order. The algebra B is *ramified at* ∞ if $B \otimes \mathbf{R} \ncong M_2(\mathbf{R})$. Let ℓ be a prime number. The algebra B is *ramified at* ℓ if $B_\ell := B \otimes \mathbf{Q}_\ell \ncong M_2(\mathbf{Q}_\ell)$. If ℓ is unramified and $\mathcal{O}$ a maximal order, then $\mathcal{O}_\ell := \mathcal{O} \otimes \mathbf{Z}_\ell \cong M_2(\mathbf{Z}_\ell)$. The discriminant of a maximal order in B is the square of the product of the ramified primes of B. When B is ramified at ∞, the quadratic form nrd is positive definite, and for every lattice L in B, the volume $\mathrm{Vol}(L)$ satisfies $\mathrm{disc}(L) = 16\,\mathrm{Vol}(B)^2$.

2.3 Elliptic Curves

A general reference for this section is [Sil86]. An *elliptic curve* over a field K is a genus 1 projective curve with a specified base point O. An elliptic curve has a unique algebraic group law with neutral element O. An algebraic morphism between elliptic curves (preserving the base point) is automatically a group morphism, and is either constant or surjective. In the latter case, we say that it is an *isogeny*. The *degree* $\deg(\varphi)$ of an isogeny φ is its degree as a rational map. An isogeny of degree d is called a *d-isogeny*. For every integer $n \neq 0$, the multiplication-by-n map $[n] \colon E \to E$ is an isogeny of degree n^2. Every isogeny $\varphi \colon E \to E'$ has a *dual isogeny* $\hat{\varphi} \colon E' \to E$ such that $\varphi\hat{\varphi} = [\deg \varphi]$ and $\hat{\varphi}\varphi = [\deg \varphi]$. An *endomorphism* is a morphism $E \to E$. We denote $\mathrm{End}(E)$ the ring of endomorphisms of E defined over $\overline{K}$. The degree map is a positive definite quadratic form on $\mathrm{End}(E)$. For $\alpha \in \mathrm{End}(E)$, the endomorphism $\alpha + \hat{\alpha}$ equals the multiplication map by an integer, the *trace* $\mathrm{Tr}(\alpha)$ of α, and we have $\mathrm{Tr}(\alpha)^2 \leq 4\deg(\alpha)$; we also define the *discriminant* $\mathrm{disc}(\alpha) = \mathrm{Tr}(\alpha)^2 - 4\deg(\alpha)$, which satisfies $|\mathrm{disc}(\alpha)| \leq 4\deg(\alpha)$ and $\mathrm{disc}(\alpha + [n]) = \mathrm{disc}(\alpha)$ for all $n \in \mathbf{Z}$. If the characteristic of K is not 2 or 3, we have $\mathrm{Aut}(E) = \{\pm 1\}$ for all E, except two isomorphism classes over $\overline{K}$ having respectively $\#\,\mathrm{Aut}(E) = 6$ and $\#\,\mathrm{Aut}(E) = 4$. Assume that K has positive characteristic p and let E be an elliptic curve over K. We say that E is *supersingular* if $\mathrm{End}(E)$ is an order in a quaternion algebra. In this case, $B = \mathrm{End}(E) \otimes \mathbf{Q}$ is a quaternion algebra over $\mathbf{Q}$ with ramification

set $\{p, \infty\}$, the ring $\mathrm{End}(E)$ is a maximal order in B, and E is defined over $\mathbf{F}_{p^2}$. When we see a nonzero endomorphism $\alpha \in \mathrm{End}(E)$ as a quaternion $a \in B$, we have $\deg(\alpha) = \mathrm{nrd}(a)$ and $\mathrm{Tr}(\alpha) = \mathrm{trd}(a)$.

2.4 Computing with Isogenies

Let us formalise how one can computationally encode isogenies. All we need is a notion of *efficient representation*: some data efficiently represents an isogeny if it allows to evaluate it efficiently on arbitrary inputs.

Definition 2.1 (Efficient representation). *Let $\mathscr{A}$ be an algorithm, and let $\varphi : E \to E'$ be an isogeny over a finite field $\mathbf{F}_q$. An* efficient representation *of φ (with respect to $\mathscr{A}$) is some data $D_\varphi \in \{0,1\}^*$ such that*

- *$D_\varphi \in \{0,1\}^*$ has size polynomial in $\log(\deg(\varphi))$ and $\log q$, and*
- *on input D_φ and $P \in E(\mathbf{F}_{q^k})$, the algorithm $\mathscr{A}$ returns $\varphi(P)$, and runs in polynomial time in $\log(\deg(\varphi))$, $\log q$, and k.*

Remark 2.2. When we say that an isogeny is in efficient representation, the algorithm $\mathscr{A}$ is often left implicit. There are only a handful of known algorithms to evaluate isogenies, so one can think of $\mathscr{A}$ as an algorithm that implements each of these, and D_φ would start with an indicator of which algorithm to use.

Proposition 2.3. *There is an algorithm* Divide *which takes as input*

- *a supersingular elliptic curve $E/\mathbf{F}_{p^2}$,*
- *an endomorphism α of E in efficient representation, and*
- *an integer N,*

and returns an efficient representation of α/N if $\alpha \in N\,\mathrm{End}(E)$, and $\perp$ otherwise, and runs in time polynomial in the length of the input.

Proof. This is the division algorithm introduced by Robert [Rob22] that was derived from the attacks on SIDH [CD23, MMP+23, Rob23]. Note that in [Rob22], the algorithm is only presented for particular endomorphisms (translates of the Frobenius), but it works, mostly unchanged, in all generality. The general statement and detailed proof can be found in [HLMW23]. □

2.5 Computational Problems

The endomorphism ring problem is the following.

Problem 2.4 (EndRing). *Given a prime p and a supersingular elliptic curve E over $\mathbf{F}_{p^2}$, find four endomorphisms in efficient representation that form a basis of $\mathrm{End}(E)$ as a lattice.*

As the endomorphism ring problem asks to find, in a sense, all the endomorphisms, it is natural to study the problem of finding even a single one. Scalar multiplications $[m]$ for $m \in \mathbf{Z}$ are trivial to find, so we exclude them.

Problem 2.5 (OneEnd). *Given a prime p and a supersingular elliptic curve E over $\mathbf{F}_{p^2}$, find an endomorphism in $\mathrm{End}(E) \setminus \mathbf{Z}$ in efficient representation.*

There exists arbitrarily large endomorphisms, so it is convenient to introduce a bounded version of this problem. Given a function $\lambda\colon \mathbf{Z}_{>0} \to \mathbf{Z}_{>0}$, the OneEnd$_\lambda$ problem denotes the OneEnd problem where the solution α is required to satisfy $\log(\deg \alpha) \leq \lambda(\log p)$ (in other words, the length of the output is bounded by a function of the length of the input).

The ℓ-isogeny path problem is a standard problem in isogeny-based cryptography. Fix a prime ℓ. An ℓ-isogeny path is a sequence of isogenies of degree ℓ such that the target of each isogeny is the source of the next.

Problem 2.6 (ℓ-IsogenyPath). *Given a prime p and two supersingular elliptic curves E and E' over $\mathbf{F}_{p^2}$, find an ℓ-isogeny path from E to E'.*

2.6 Probabilities

Given a random variable X with values in a discrete set $\mathscr{X}$, we say it has distribution f if $f(x) = \Pr[X = x]$ for every $x \in \mathscr{X}$. We also write $f(A) = \sum_{x\in A} f(x)$ for any $A \subseteq \mathscr{X}$. For two distributions f_1 and f_2 over the same set $\mathscr{X}$, their *statistical distance* (or *total variation distance*) is

$$\frac{1}{2}\|f_1 - f_2\|_1 = \frac{1}{2}\sum_{x\in\mathscr{X}} |f_1(x) - f_2(x)| = \sup_{A\subseteq\mathscr{X}} |f_1(A) - f_2(A)|.$$

Random walks play a key role in isogeny-based cryptography. Fix a field $\mathbf{F}_{p^2}$ and a prime number $\ell \neq p$. The supersingular ℓ-isogeny graph has vertices the (finitely many) isomorphism classes of supersingular elliptic curves over $\mathbf{F}_{p^2}$, and edges are the ℓ-isogenies between them (up to isomorphism of the target). At the heart of the Charles–Goren–Lauter hash function [CLG09], one of the first isogeny-based constructions, lies the fact that random walks in supersingular ℓ-isogeny graphs have rapid-mixing properties: they are Ramanujan graphs. This is the following well-known proposition. It is a particular case of our more general Theorem 3.10.

Proposition 2.7. *Let E be a supersingular elliptic curve over $\mathbf{F}_{p^2}$, and $\ell \neq p$ a prime number. Let $\varepsilon > 0$. There is a bound $n = O(\log_\ell(p) - \log_\ell(\varepsilon))$ such that the endpoint of a uniform random walk of length at least n from E in the ℓ-isogeny graph is at statistical distance at most ε from the stationary distribution f, which satisfies $f(E) = \frac{24}{(p-1)\#\operatorname{Aut}(E)}$.*

Proof. This is a standard consequence of Pizer's proof that the supersingular ℓ-isogeny graph is Ramanujan [Piz90]. Details can be found, for instance, in [BCC+23, Theorem 11] for the length of the walk, and in [BCC+23, Theorem 7, Item 2] for the description of the stationary distribution. □

The stationary distribution is at statistical distance $O(1/p)$ of the uniform distribution. Note that rejection sampling allows to efficiently transform a sampler for the stationary distribution into a sampler for the uniform distribution.

2.7 Categories

A general reference for this section is [ML98]. A *category* $\mathcal{C}$ consists of objects, for every objects $x, y \in \mathcal{C}$, a set of morphisms $\mathrm{Hom}_{\mathcal{C}}(x, y)$, sometimes denoted $f\colon x \to y$, an associative composition law for morphisms with compatible source and target, and an identity morphism $\mathrm{id}_x \in \mathrm{Hom}_{\mathcal{C}}(x, x)$ for every object $x \in \mathcal{C}$. Let $\mathcal{C}, \mathcal{D}$ be categories. A *functor* $\mathcal{F}\colon \mathcal{C} \to \mathcal{D}$ is an association of an object $\mathcal{F}(x) \in \mathcal{D}$ for every object $x \in \mathcal{C}$, and of a morphism $\mathcal{F}(f)\colon \mathcal{F}(x) \to \mathcal{F}(y)$ for every morphism $f\colon x \to y$ in $\mathcal{C}$, that respects composition[1] and identities.

Let Sets be the category of sets. The following is a standard construction.

Definition 2.8. *Let $\mathcal{C}$ be a category and $\mathcal{F}\colon \mathcal{C} \to$ Sets be a functor. The* category of elements $\mathrm{El}(\mathcal{F})$ *is the category with*

- *objects: pairs (c, x) where $c \in \mathcal{C}$ and $x \in \mathcal{F}(c)$;*
- *morphisms $(c, x) \to (c', x')$: morphisms $f \in \mathrm{Hom}_{\mathcal{C}}(c, c')$ s.t. $\mathcal{F}(f)(x) = x'$.*

Remark 2.9. One could also use the contravariant version of this definition. All our results would hold in this setting, as one can compose $\mathcal{F}$ with the isogeny duality to reverse the direction of all morphisms.

3 Equidistribution of Elliptic Curves with Extra Data

The goal of this section is to state Theorem 3.10, whose proof is available in the full version [PW23, Section 3].

3.1 Statement of the Equidistribution Theorem

In order to avoid bad primes, we will need to restrict the possible degrees of isogenies under consideration. Let Σ be a set of primes, and let $N \geq 1$ be an integer not divisible by any prime in Σ.

Definition 3.1. *Let $\mathrm{SS}_\Sigma(p)$ denote the category with*

- *objects: supersingular elliptic curves over $\overline{\mathbf{F}}_p$;*
- *morphisms $\mathrm{Hom}_\Sigma(E, E')$: isogenies with degree a product of the primes in Σ.*

When Σ is the set of all primes, we simply write $\mathrm{SS}(p)$.

Our results are expressed in terms of categories of elements of various functors, as in Definition 2.8. For us, this is going to play the role of "equipping with extra structure": when $\mathcal{F}\colon \mathcal{C} \to$ Sets is a functor, $\mathrm{El}(\mathcal{F})$ is the category of "objects $c \in \mathcal{C}$ with extra structure taken from $\mathcal{F}(c)$".

Example 3.2. Assume $p \nmid N$. Let Σ be the set of primes not dividing N. Define the functor $\mathrm{Cyc}_N\colon \mathrm{SS}_\Sigma(p) \to$ Sets by:

[1] All our functors are covariant.

– $\mathrm{Cyc}_N(E)$ is the set of cyclic subgroups of order N of E;
– for every isogeny $\varphi \in \mathrm{Hom}_\Sigma(E, E')$, the map $\mathrm{Cyc}_N(\varphi)$ is $C \mapsto \varphi(C)$.

Then $\mathrm{El}(\mathrm{Cyc}_N)$ is the category of supersingular elliptic curves equipped with a cyclic subgroup of order N.

Example 3.3. Let Σ be the set of primes not dividing N. Let $\mathrm{End}\,/N$ denote the functor $\mathrm{SS}_\Sigma(p) \to \mathrm{Sets}$ defined by

– $(\mathrm{End}\,/N)(E) = \mathrm{End}(E)/N\,\mathrm{End}(E)$;
– for $\varphi\colon E \to E'$, the map $(\mathrm{End}\,/N)(\varphi)$ is $\alpha \mapsto \varphi\alpha\hat{\varphi}$.

Then $\mathrm{El}(\mathrm{End}\,/N)$ is the category of supersingular elliptic curves equipped with an endomorphism modulo N, which will play an important role in Sect. 4.

We now introduce the graphs of interest (more generally see Definition ??).

Definition 3.4. *Let $\mathcal{F}\colon \mathrm{SS}_\Sigma(p) \to \mathrm{Sets}$ be a functor with $\mathcal{F}(E)$ finite for all E. We define the graph $\mathcal{G}_\mathcal{F}$ with:*

– *vertices: isomorphism classes of objects in $\mathrm{El}(\mathcal{F})$;*
– *edges: let $(E, x) \in \mathrm{El}(\mathcal{F})$; edges from (E, x) are isogenies $\varphi \in \mathrm{Hom}_\Sigma(E, E')$ modulo automorphisms of $(E', \mathcal{F}(\varphi)(x))$.*

Let $L^2(\mathcal{G}_\mathcal{F})$ be the space of complex functions on vertices of $\mathcal{G}_\mathcal{F}$, and define

$$\langle F, G\rangle = \sum_{(E,x)\in\mathcal{G}_\mathcal{F}} \frac{F(E,x)\overline{G(E,x)}}{\#\,\mathrm{Aut}(E,x)} \text{ for } F, G \in L^2(\mathcal{G}_\mathcal{F}).$$

For every prime ℓ, we define the adjacency operator A_ℓ on $L^2(\mathcal{G}_\mathcal{F})$ by

$$A_\ell F(E,x) = \sum_{(E,x)\to(E',x')} F(E',x'),$$

where the sum runs over edges of degree ℓ leaving (E, x).

Remark 3.5. The graphs $\mathcal{G}_\mathcal{F}$ have finitely many vertices, but infinitely many edges.

Example 3.6. Assume $p \nmid N$, and let ℓ a prime not dividing Np. The graph obtained from $\mathcal{G}_{\mathrm{Cyc}_N}$ by keeping only the edges of degree ℓ is the ℓ-isogeny graph of supersingular elliptic curves with Borel structure studied in [Arp23] and [BCC+23]. When $N = 1$ this is the classical supersingular ℓ-isogeny graph.

We are now in position to state our equidistribution theorem.

Definition 3.7. *Let $\mathcal{F}\colon \mathrm{SS}_\Sigma(p) \to \mathrm{Sets}$ be a functor and $N \geq 1$ an integer. We say that $\mathcal{F}$* satisfies the (mod N)-congruence property *if for every $E \in \mathrm{SS}(p)$ and every $\varphi, \psi \in \mathrm{End}_\Sigma(E)$ such that $\varphi - \psi \in N\,\mathrm{End}(E)$, we have $\mathcal{F}(\varphi) = \mathcal{F}(\psi)$.*

Example 3.8. Assume that p does not divide N. The functor Cyc_N from Example 3.2 satisfies the (mod N)-congruence property: indeed, endomorphisms divisible by N act as 0 on N-torsion points.

Example 3.9. The functor End/N from Example 3.3 has the (mod N)-congruence property: if $\varphi, \psi \in \mathrm{End}_\Sigma(E)$ and $\alpha, \beta \in \mathrm{End}(E)$ satisfy $\psi = \varphi + N\beta$, then $\psi\alpha\hat{\psi} = (\varphi + N\beta)\alpha(\hat{\varphi} + N\hat{\beta}) \in \varphi\alpha\hat{\varphi} + N\,\mathrm{End}(E)$, so that $(\mathrm{End}/N)(\varphi) = (\mathrm{End}/N)(\psi)$.

Theorem 3.10. *Let p be a prime and $N \geq 1$ an integer. Let Σ be a set of primes that do not divide N, such that Σ generates $(\mathbf{Z}/N\,\mathbf{Z})^\times$. Let $\mathcal{F}\colon \mathrm{SS}_\Sigma(p) \to \mathrm{Sets}$ be a functor satisfying the* (mod N)*-congruence property and such that all sets $\mathcal{F}(E)$ are finite.*

Then, for every $\ell \in \Sigma$ different from p, the adjacency operator A_ℓ on $L^2(\mathcal{G}_\mathcal{F})$ commutes with its adjoint and stabilises the following subspaces:

- *$L^2_{\deg}(\mathcal{G}_\mathcal{F})$, the subspace of functions that are constant on every connected component of the graph $\mathcal{G}^1_\mathcal{F}$ obtained from $\mathcal{G}_\mathcal{F}$ by keeping only the edges of degree* 1 mod N*. For all $f \in L^2_{\deg}(\mathcal{G}_\mathcal{F})$ we have $\|A_\ell f\| \leq (\ell+1)\cdot\|f\|$.*
- *$L^2_0(\mathcal{G}_\mathcal{F})$, the orthogonal complement of $L^2_{\deg}(\mathcal{G}_\mathcal{F})$. For all $f \in L^2_0(\mathcal{G}_\mathcal{F})$ we have $\|A_\ell f\| \leq 2\sqrt{\ell}\cdot\|f\|$.*

Moreover, the A_ℓ for $\ell \in \Sigma$ pairwise commute.

In other words, the normalised operator $A'_\ell = \frac{1}{\ell+1}A_\ell$ makes functions rapidly converge to the subspace $L^2_{\deg}(\mathcal{G}_\mathcal{F})$. This operator A'_ℓ preserves the subset of probability distributions, and closely relates to the effect of a random walk of ℓ-isogenies (see the full version [PW23, Appendix A.1]). In simple cases (such as $N = 1$), the space $L^2_{\deg}(\mathcal{G}_\mathcal{F})$ has dimension 1, is generated by the constant function 1 and the theorem says that random walks in ℓ-isogeny graphs rapidly converge to the unique stationary distribution f with $f(E,x)$ proportional to $\frac{1}{\#\,\mathrm{Aut}(E,x)}$. One thus sees that the classical rapid-mixing property for isogeny graphs (Proposition 2.7) is a particular case of Theorem 3.10. More details and other illustrations of Theorem 3.10 are available in the full version [PW23, Appendix A].

In general $L^2_{\deg}(\mathcal{G}_\mathcal{F})$ could have higher dimension. This reflects the fact that the graph may be disconnected or multipartite, two obstructions for random walks to converge to a unique limit. To ease the application of Theorem 3.10 in such cases, we provide the following companion proposition that gives extra information on the graph $\mathcal{G}_\mathcal{F}$ and an explicit description of the space $L^2_{\deg}(\mathcal{G}_\mathcal{F})$.

Proposition 3.11. *With the same hypotheses and notations as in Theorem 3.10:*

(1) for every isogeny φ in $\mathrm{SS}_\Sigma(p)$, the map $\mathcal{F}(\varphi)$ is a bijection;
(2) for every $E, E' \in \mathrm{SS}(p)$, there exists $\varphi \in \mathrm{Hom}_\Sigma(E,E')$ of degree 1 mod N*;*
(3) for every $E \in \mathrm{SS}(p)$, the morphism $\mathrm{End}_\Sigma(E) \to (\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$ is surjective, inducing an action of $G = (\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$ on $\mathcal{F}(E)$.

Let $x_1, \dots, x_n$ denote representatives of the orbits of the action of G on $\mathcal{F}(E_0)$ and for each i, let H_i denote the stabiliser of x_i in G. Let $\mathcal{G}_{\deg}$ denote the graph with edges labelled by elements of $(\mathbf{Z}/N\mathbf{Z})^\times$ and with

- *vertex set $\bigsqcup_i (\mathbf{Z}/N\mathbf{Z})^\times/\deg(H_i)$;*
- *for every i, every $a \in (\mathbf{Z}/N\mathbf{Z})^\times/\deg(H_i)$ and every $d \in (\mathbf{Z}/N\mathbf{Z})^\times$, an edge $a \to b$ labelled by d, where $b = ad \in (\mathbf{Z}/N\mathbf{Z})^\times/\deg(H_i)$.*

Then:

(4) there exists a unique morphism of graphs $\mathrm{Deg}\colon \mathcal{G}_{\mathcal{F}} \longrightarrow \mathcal{G}_{\deg}$ *such that for all i we have* $\mathrm{Deg}(E_0, x_i) = 1 \in (\mathbf{Z}/N\mathbf{Z})^\times/\deg(H_i)$ *and for every edge φ of $\mathcal{G}_{\mathcal{F}}$, the edge* $\mathrm{Deg}(\varphi)$ *is labelled by* $\deg(\varphi) \bmod N$*;*
(5) the map Deg *is surjective; and*
(6) $L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$ is the space of functions that factor through Deg*.*

Remark 3.12.

- When $p \nmid N$, Property (3) can be used to relate $\mathcal{F}$ to the setup of [CL23], using an isomorphism $G \cong \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$. When $p \mid N$, for any $E \in \mathrm{SS}(p)$, the group $(\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$ is not isomorphic to $\mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$.
- The graph $\mathcal{G}_{\deg}$ is the Cayley graph of the set $\bigsqcup_i (\mathbf{Z}/N\mathbf{Z})^\times/\deg(H_i)$ equipped with its natural action of $(\mathbf{Z}/N\mathbf{Z})^\times$.
- Property (4) amounts to stating the existence of a disconnectedness and a multipartition of $\mathcal{G}_{\mathcal{F}}$.
- Using Properties (5) and (6), one easily obtains the spectra of the adjacency operators A_ℓ on $L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$: for every complex character χ of $(\mathbf{Z}/N\mathbf{Z})^\times$, one obtains the eigenvalue $\chi(\ell)(\ell+1)$ with multiplicity equal to the number of i such that $\chi(\deg(H_i)) = 1$.
- From Proposition 3.11 and Theorem 3.10, since $2\sqrt{\ell} < \ell + 1$, one can simply deduce connectedness and multipartition properties of $\mathcal{G}_{\mathcal{F}}$, its degree ℓ subgraphs, etc. For instance, the graph $\mathcal{G}_{\mathcal{F}}$ has exactly n connected components: the preimages of the $(\mathbf{Z}/N\mathbf{Z})^\times/\deg(H_i)$ via the map Deg.

Example 3.13. Assume $p \nmid N$, let Σ denote the set of all primes that do not divide pN, and apply Theorem 3.10 and Proposition 3.11 to $\mathcal{F} = \mathrm{Cyc}_N$. Then we have an isomorphism $G \cong \mathrm{GL}_2(\mathbf{Z}/N\mathbf{Z})$ and a compatible bijection $\mathcal{F}(E) \cong \{\mathbf{Z}/N\mathbf{Z}\text{-lines in } (\mathbf{Z}/N\mathbf{Z})^2\}$. In particular, there is a single orbit ($n = 1$) and, choosing x_1 corresponding to the line generated by $\binom{1}{0}$, the stabiliser $H = H_1$ corresponds to the subgroup of upper-triangular matrices, so that $\deg(H) = (\mathbf{Z}/N\mathbf{Z})^\times$. The space $L^2_{\deg}(\mathcal{G}_{\mathrm{Cyc}_N})$ is therefore one-dimensional, generated by the constant function 1. Hence Theorem 3.10 recovers [BCC+23, Theorem 8].

3.2 Proof of Theorem 3.10 and Proposition 3.11

The proof of these results is available in the full version [PW23, Section 3].

4 Enriching a OneEnd Oracle

In this section, we show how to turn an oracle for the OneEnd problem into a richer oracle with better distributed output. The quality of this enrichment is quantified in Theorem 4.2. The proof is an application of the equidistribution results of Sect. 3. The following lemma relates conjugation-invariance of distributions to the abstract setup of Sect. 3.

Lemma 4.1. *Let $p > 3$ be a prime, let $N \geq 1$ and let $E \in \mathrm{SS}(p)$. Let $g \in (\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$ be an element of degree $1 \in (\mathbf{Z}/N\,\mathbf{Z})^\times$. Define the linear operator $c_g \colon L^2(\mathcal{G}_{\mathrm{End}/N}) \to L^2(\mathcal{G}_{\mathrm{End}/N})$ by*

$$c_g F(E,\alpha) = F(E, g\alpha g^{-1}) \text{ and } c_g F(E',\alpha') = F(E',\alpha') \text{ for all } E' \neq E.$$

Then:

(1) for all $F \in L^2(\mathcal{G}_{\mathrm{End}/N})$, we have $\|c_g F\|^2 \leq 3\|F\|^2$;
(2) for all $G \in L^2_{\deg}(\mathcal{G}_{\mathrm{End}/N})$, we have $c_g G = G$; and
(3) for every $F = F_0 + F_1 \in L^2(\mathcal{G}_{\mathrm{End}/N})$ with $F_0 \in L^2_0(\mathcal{G}_{\mathrm{End}/N})$ and $F_1 \in L^2_{\deg}(\mathcal{G}_{\mathrm{End}/N})$, we have $\|F - c_g F\| \leq (1+\sqrt{3})\|F_0\|$.

Proof.

(1) Let $F \in L^2(\mathcal{G}_{\mathrm{End}/N})$. We have, where E' ranges over the set of supersingular curves up to isomorphism except E, the elements α and β range over $\mathrm{End}(E)/N\,\mathrm{End}(E)$ and α' over $\mathrm{End}(E')/N\,\mathrm{End}(E')$,

$$\|F\|^2 = \sum_{(E,\alpha)} \frac{1}{\#\,\mathrm{Aut}(E,\alpha)} |F(E,\alpha)|^2 + \sum_{(E',\alpha')} \frac{1}{\#\,\mathrm{Aut}(E',\alpha')} |F(E',\alpha')|^2,$$

and

$$\begin{aligned}
\|c_g F\|^2 &= \sum_{(E,\alpha)} \frac{|F(E,g\alpha g^{-1})|^2}{\#\,\mathrm{Aut}(E,\alpha)} + \sum_{(E',\alpha')} \frac{|F(E',\alpha')|^2}{\#\,\mathrm{Aut}(E',\alpha')} \\
&= \sum_{(E,\beta)} \frac{|F(E,\beta)|^2}{\#\,\mathrm{Aut}(E,g^{-1}\beta g)} + \sum_{(E',\alpha')} \frac{|F(E',\alpha')|^2}{\#\,\mathrm{Aut}(E',\alpha')} \\
&= \sum_{(E,\beta)} \frac{\#\,\mathrm{Aut}(E,\beta)}{\#\,\mathrm{Aut}(E,g^{-1}\beta g)} \frac{|F(E,\beta)|^2}{\#\,\mathrm{Aut}(E,\beta)} + \sum_{(E',\alpha')} \frac{|F(E',\alpha')|^2}{\#\,\mathrm{Aut}(E',\alpha')} \\
&\leq 3 \sum_{(E,\beta)} \frac{|F(E,\beta)|^2}{\#\,\mathrm{Aut}(E,\beta)} + \sum_{(E',\alpha')} \frac{|F(E',\alpha')|^2}{\#\,\mathrm{Aut}(E',\alpha')} \leq 3\|F\|^2,
\end{aligned}$$

where the inequality comes from $\#\,\mathrm{Aut}(E,\beta) \leq 6$ and $\#\,\mathrm{Aut}(E,g^{-1}\beta g) \geq 2$.

(2) Let $h \in \mathrm{End}_\Sigma(E)$ be a lift of g, which exists by Proposition 3.11 (3). Let $G \in L^2_{\deg}(\mathcal{G}_{\mathrm{End}/N})$. For every $E' \neq E$ we have $c_g G(E',\alpha) = G(E',\alpha)$. Moreover, h defines an edge $(E,\alpha) \to (E, h\alpha\hat{h}) = (E, h\alpha h^{-1})$ in $\mathcal{G}_{\mathrm{End}/N}$ of degree $1 \bmod N$, so $G(E,\alpha) = G(E, h\alpha h^{-1})$. Since $c_g G(E,\alpha) = G(E, g\alpha g^{-1}) = G(E, h\alpha h^{-1})$, this proves that $c_g G = G$.

(3) Let $F = F_0 + F_1 \in L^2(\mathcal{G}_{\mathrm{End}/N})$ with $F_0 \in L^2_0(\mathcal{G}_{\mathrm{End}/N})$ and $F_1 \in L^2_{\deg}(\mathcal{G}_{\mathrm{End}/N})$. Then

$$\begin{aligned}\|F - c_g F\| &\le \|F_0 - c_g F_0\| + \|F_1 - c_g F_1\| \\ &= \|F_0 - c_g F_0\| \text{ since } c_g F_1 = F_1 \text{ by (2)} \\ &\le (1+\sqrt{3})\|F_0\| \text{ by (1).}\end{aligned}$$

□

Algorithm 1. $\textsc{Rich}^{\mathscr{O}}_k$: turning an oracle $\mathscr{O}$ for OneEnd into a 'richer' oracle $\textsc{Rich}^{\mathscr{O}}_k$, with guarantees on the distribution of the output.

Require: A supersingular elliptic curve $E/\mathbf{F}_{p^2}$, and a parameter $k \in \mathbf{Z}_{>0}$. We suppose access to an oracle $\mathscr{O}$ that solves the *ProblemOneEnd* problem.
Ensure: An endomorphism $\alpha \in \mathrm{End}(E)$.
1: $\varphi \leftarrow$ a 2-isogenies random walk of length k from E
2: $E' \leftarrow$ endpoint of φ
3: $\alpha \leftarrow \mathscr{O}(E')$, a non-scalar endomorphism of E'
4: **return** $\hat{\varphi} \circ \alpha \circ \varphi$

We can now prove the main result of this section.

Theorem 4.2. *Let $p > 3$ be a prime and N an odd integer. Let $\mathscr{O}$ be an oracle for* OneEnd. *Let E be a supersingular elliptic curve defined over $\mathbf{F}_{p^2}$ and let α be the random endomorphism produced by $\textsc{Rich}^{\mathscr{O}}_k(E)$. Then for every element $g \in (\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$ of degree $1 \in (\mathbf{Z}/N\,\mathbf{Z})^\times$, the statistical distance between the distribution of $\alpha \bmod N$ and the distribution of $g^{-1}(\alpha \bmod N)g$ is at most $\frac{1+\sqrt{3}}{4}(\frac{2\sqrt{2}}{3})^k N^2\sqrt{p+13} = O\big((\frac{2\sqrt{2}}{3})^k N^2\sqrt{p}\big)$.*

Proof. Define $F \in L^2(\mathcal{G}_{\mathrm{End}/N})$ by the following formula for every vertex (E', β):

$$F(E', \beta) = \Pr[\mathscr{O}(E') \bmod N = \beta], \text{ so that}$$
$$\Big(\frac{A_2}{3}\Big)^k F(E, \beta) = \Pr[\textsc{Rich}^{\mathscr{O}}_k(E) \bmod N = 4^k\beta].$$

Indeed, $\big(\frac{A_2}{3}\big)^k F(E,\beta)$ is the average, over all random walks $\varphi\colon E \to E'$ that Algorithm 1 could follow from E, of $\Pr[\mathscr{O}(E') \bmod N = \varphi\beta\hat{\varphi}]$, and the equality $\mathscr{O}(E') \bmod N = \varphi\beta\hat{\varphi}$ is equivalent to $\hat{\varphi}\mathscr{O}(E')\varphi \bmod N = \deg(\varphi)\beta\deg(\hat{\varphi}) = 4^k\beta$ since 2 is invertible mod N.

We have, where E' ranges over isomorphism classes in $\mathrm{SS}(p)$ and β over the set $\mathrm{End}(E')/N\,\mathrm{End}(E')$,

$$\begin{aligned}\|F\|^2 &= \sum_{(E',\beta)} \frac{1}{\#\operatorname{Aut}(E',\beta)} \Pr[\mathscr{O}(E') \bmod N = \beta]^2 \\ &\le \frac{1}{2} \sum_{(E',\beta)} \Pr[\mathscr{O}(E') \bmod N = \beta] \\ &= \frac{1}{2} \sum_{E'} 1 \le \frac{p+13}{24} \text{ by [Sil86, Theorem 4.1 (c)].}\end{aligned}$$

Write $F = F_0 + F_1$ with $F_0 \in L^2_0(\mathcal{G}_{\mathrm{End}/N})$ and $F_1 \in L^2_{\deg}(\mathcal{G}_{\mathrm{End}/N})$. Since A_2 preserves the orthogonal decomposition $L^2_0(\mathcal{G}_{\mathrm{End}/N}) \oplus L^2_{\deg}(\mathcal{G}_{\mathrm{End}/N})$, we may apply Lemma 4.1 (3) to $A_2^k F = A_2^k F_0 + A_2^k F_1$, giving $\|A_2^k F - c_g A_2^k F\| \le (1+\sqrt{3})\|A_2^k F_0\|$. On the other hand, by Theorem 3.10 we have $\left\|(\frac{A_2}{3})^k F_0\right\| \le (\frac{2\sqrt{2}}{3})^k \|F_0\| \le (\frac{2\sqrt{2}}{3})^k \|F\|$. Finally, with β ranging over $\mathrm{End}(E)/N\,\mathrm{End}(E)$, the statistical distance in the statement of the theorem is

$$\begin{aligned}&\frac{1}{2}\sum_{\beta}\left|\Pr[\textsc{Rich}_k^{\mathscr{O}}(E) \bmod N = \beta] - \Pr[\textsc{Rich}_k^{\mathscr{O}}(E) \bmod N = g\beta g^{-1}]\right| \\ &= \frac{1}{2}\sum_{\beta}\left|\left(\frac{A_2}{3}\right)^k F(E, 4^{-k}\beta) - c_g\left(\frac{A_2}{3}\right)^k F(E, 4^{-k}\beta)\right| \\ &= \frac{1}{2}\sum_{\beta}\left|\left(\frac{A_2}{3}\right)^k F(E, \beta) - c_g\left(\frac{A_2}{3}\right)^k F(E, \beta)\right| \text{ since } \beta \mapsto 4^k\beta \text{ is a bijection.}\end{aligned}$$

By the Cauchy–Schwarz inequality, this is bounded by

$$\begin{aligned}&\frac{1}{2}\left(N^4 \sum_{\beta}\left|\left(\frac{A_2}{3}\right)^k F(E,\beta) - c_g\left(\frac{A_2}{3}\right)^k F(E,\beta)\right|^2\right)^{\frac{1}{2}} \\ &\le \frac{1}{2} N^2\sqrt{6}\left\|\left(\frac{A_2}{3}\right)^k F - c_g\left(\frac{A_2}{3}\right)^k F\right\| \text{ since } \#\operatorname{Aut}(E,\beta) \le 6 \\ &\le \frac{1}{2}(1+\sqrt{3}) N^2\sqrt{6}\left\|\left(\frac{A_2}{3}\right)^k F_0\right\| \le \frac{1}{2}(1+\sqrt{3})\left(\frac{2\sqrt{2}}{3}\right)^k N^2\sqrt{6}\|F\| \\ &\le \frac{1}{2}(1+\sqrt{3})\left(\frac{2\sqrt{2}}{3}\right)^k N^2\sqrt{6\cdot\frac{p+13}{24}} = \frac{1}{4}(1+\sqrt{3})\left(\frac{2\sqrt{2}}{3}\right)^k N^2\sqrt{p+13},\end{aligned}$$

as claimed. □

5 On Conjugacy-Invariant Distributions

Theorem 4.2 proves that given a OneEnd oracle, the randomization method allows one to sample endomorphisms from a distribution which is (locally) invariant under conjugation by $(\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$. In this section, we study such

conjugacy-invariant distributions, and show that with good probability, such endomorphisms generate interesting suborders. In the whole section, fix B a quaternion algebra over $\mathbf{Q}$ and $\mathcal{O} \subset B$ a maximal order.

5.1 The Local Case

We start by studying the local case. Let ℓ be a prime unramified in B. In this subsection, we study distributions on $M_2(\mathbf{F}_\ell) \cong \mathcal{O}/\ell\mathcal{O}$ and $M_2(\mathbf{Z}_\ell) \cong \mathcal{O}_\ell$.

Definition 5.1. *The distribution of a random $\alpha \in M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ is* ε-close to $\mathrm{SL}_2(\mathbf{F}_\ell)$-invariant *if, for every $g \in \mathrm{SL}_2(\mathbf{F}_\ell)$, the statistical distance between the distributions of α and of $g^{-1}\alpha g$ is at most ε. When the distributions are the same (i.e., $\varepsilon = 0$), we say that the distribution of α is* $\mathrm{SL}_2(\mathbf{F}_\ell)$-invariant.

A key observation is that a conjugacy class cannot be stuck in a subspace.

Lemma 5.2. *Suppose $\ell > 2$. Let $\alpha \in M_2(\mathbf{F}_\ell) \setminus \mathbf{F}_\ell$. Let $V \subsetneq M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ be an $\mathbf{F}_\ell$-linear subspace. Let $\beta \in M_2(\mathbf{F}_\ell)$ be a random element uniformly distributed in the $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugacy class of α. Then, $\beta \in V$ with probability at most $1/2$.*

Proof. The size of the orbit X of α is $\#\mathrm{SL}_2(\mathbf{F}_\ell)/\#C$, where C is the centraliser of α in $\mathrm{SL}_2(\mathbf{F}_\ell)$. The size of this centraliser can be $\ell+1, \ell-1$ or 2ℓ, so $\#X \geq \frac{\ell^2-1}{2}$.

We now bound $\#(X \cap V)$ by noting that every element v of this intersection satisfies the quadratic equation $\mathrm{disc}(v) = \mathrm{disc}(\alpha)$. The discriminant quadratic form on $M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ is isomorphic to $x^2 - yz$, so the maximal dimension of a totally isotropic subspace is 1. If $\dim V = 1$, the number of solutions is at most ℓ. If $\dim V = 2$, either the equation is degenerate and has at most 2ℓ solutions, or it represents a conic and has at most $\ell + 1$ solutions.

So the probability of $\beta \in V$ is at most $2\ell/\frac{\ell^2-1}{2} = \frac{4\ell}{\ell^2-1}$, which is less than $1/2$ for $\ell \geq 11$. We check the bound by bruteforce enumeration for $\ell \in \{3, 5, 7\}$. □

Lemma 5.3. *Suppose $\ell > 2$. Let $\alpha_1, \alpha_2, \alpha_3 \in M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ be independent non-zero $\mathrm{SL}_2(\mathbf{F}_\ell)$-invariant elements. Then, $(\alpha_1, \alpha_2, \alpha_3)$ is a basis of $M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ with probability at least $1/8$.*

Proof. Let $V_1 = \{0\}$ and $V_i = V_{i-1} + \mathbf{F}_\ell \cdot \alpha_i$. By dimensionality, we have $V_i \neq M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ for every $i < 3$. Lemma 5.2 implies that with probability at least $1/8$, we have $\alpha_i \notin V_{i-1}$ for each $i \leq 3$. When this occurs, each V_i is an $\mathbf{F}_\ell$-vector space of dimension i, hence, $V_3 = M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$. □

In our application, we will only approach $\mathrm{SL}_2(\mathbf{F}_\ell)$-invariance, so we now derive the corresponding result for distributions that are close to $\mathrm{SL}_2(\mathbf{F}_\ell)$-invariant.

Proposition 5.4. *Suppose $\ell > 2$. Let $\alpha_1, \alpha_2, \alpha_3 \in M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ be independent non-zero random elements which are ε-close to $\mathrm{SL}_2(\mathbf{F}_\ell)$-invariant. Then, $(\alpha_1, \alpha_2, \alpha_3)$ is a basis of $M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ with probability at least $1/8 - 3\varepsilon$.*

Proof. Let $g_i \in \mathrm{SL}_2(\mathbf{F}_\ell)$ be uniformly distributed and independent. Let $\beta_i = g_i^{-1}\alpha_i g_i$, three independent variables. For each i, the statistical distance between α_i and β_i is at most ε. By the triangle inequality, the statistical distance between $(\alpha_1, \alpha_2, \alpha_3)$ and $(\beta_1, \beta_2, \beta_3)$ is at most 3ε. From Lemma 5.3, $(\beta_1, \beta_2, \beta_3)$ is a basis of $M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ with probability at least $1/8$. Therefore, $(\alpha_1, \alpha_2, \alpha_3)$ is a basis of $M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$ with probability at least $1/8 - 3\varepsilon$. □

We now show that these results about $M_2(\mathbf{F}_\ell)$ have consequences in $M_2(\mathbf{Z}_\ell)$.

Definition 5.5. *The* level *of $\alpha \in M_2(\mathbf{Z}_\ell) \setminus \mathbf{Z}_\ell$ at ℓ is the largest integer* $\mathrm{lev}_\ell(\alpha)$ *such that $\alpha \in \mathbf{Z}_\ell + \ell^{\mathrm{lev}_\ell(\alpha)} M_2(\mathbf{Z}_\ell)$.*

Proposition 5.6. *Suppose $\ell > 2$. Let $\alpha_1, \alpha_2, \alpha_3 \in M_2(\mathbf{Z}_\ell) \setminus \mathbf{Z}_\ell$ be three elements of level a. Then $(1, \alpha_1, \alpha_2, \alpha_3)$ is a $\mathbf{Z}_\ell$-basis of $\mathbf{Z}_\ell + \ell^a M_2(\mathbf{Z}_\ell)$ if and only if $(\alpha_1, \alpha_2, \alpha_3)$ is an $\mathbf{F}_\ell$-basis of $(\mathbf{Z}_\ell + \ell^a M_2(\mathbf{Z}_\ell))/(\mathbf{Z}_\ell + \ell^{a+1} M_2(\mathbf{Z}_\ell)) \cong M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$.*

Proof. The forward implication is clear. The converse is Nakayama's lemma. □

5.2 Dealing with Hard-to-factor Numbers

In the previous section, we have studied the properties of conjugacy-invariant distributions locally at a prime ℓ. However, in our application, we may be confronted to local obstructions at an integer N which is hard to factor; it is then not possible to isolate the primes ℓ to apply the results of the previous section.

In this section, fix a positive integer N. We imagine that N is hard to factor, and rework the previous results "*locally at N*". We suppose that B does not ramify at any prime factor of N. Recall that $\mathcal{O} \subset B$ is a maximal order.

Definition 5.7. *An element $\alpha \in \mathcal{O}$ is N-reduced if $\alpha \notin \mathbf{Z} + N\mathcal{O}$.*

Lemma 5.8. *Let $\alpha \in \mathcal{O}$ be a random variable supported on N-reduced elements. Then, there exist a prime factor ℓ of N and an integer a such that ℓ^{a+1} divides N and $\Pr[\mathrm{lev}_\ell(\alpha) = a] \geq (\log N)^{-1}$.*

Proof. Write the prime factorisation $N = \prod_{i=1}^{t} \ell_i^{e_i}$. Let i and $a < e_i$ which maximise the probability $q = \Pr[\mathrm{lev}_{\ell_i}(\alpha) = a]$. We have

$$\sum_{j=1}^{t} \Pr[\mathrm{lev}_{\ell_j}(\alpha) < e_j] = \sum_{\beta} \Pr[\alpha = \beta] \cdot \#\{j \mid \mathrm{lev}_{\ell_j}(\beta) < e_j\} \geq 1,$$

where the last inequality follows from the fact that the distribution is supported on N-reduced elements, so for every β, there exists j such that $\mathrm{lev}_{\ell_j}(\beta) < e_j$. We get

$$1 \leq \sum_{j=1}^{t} \Pr[\mathrm{lev}_{\ell_j}(\alpha) < e_j] = \sum_{j=1}^{t} \sum_{x < e_j} \Pr[\mathrm{lev}_{\ell_j}(\alpha) = x] \leq q \sum_{j=1}^{t} e_j \leq q \log(N).$$

We deduce $q \geq (\log N)^{-1}$. □

Definition 5.9. *Let M be a ring with an isomorphism $\iota\colon M_2(\mathbf{Z}/N\mathbf{Z}) \to M$. The distribution of a random $\alpha \in M/\iota(\mathbf{Z}/N\mathbf{Z})$ is* ε-close to $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})$-invariant *if, for every $g \in \iota(\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z}))$, the statistical distance between the distributions of α and of $g^{-1}\alpha g$ is at most ε.*

Lemma 5.10. *Let $R = \mathbf{Z}/N\mathbf{Z}$, $M = \mathcal{O}/N\mathcal{O} \cong M_2(R)$ and $\overline{M} = M/R$. Let ℓ be a prime factor of N, and a an integer such that $\ell^{a+1} \mid N$. Consider a distribution ν on $\overline{M}$ that is ε-close to $\mathrm{SL}_2(R)$-invariant. For α sampled from ν, let q be the probability that $\alpha \neq 0$ and that a is the largest integer such that $\alpha \in \ell^a\overline{M}$.*

1. *Let $\alpha_1, \alpha_2, \alpha_3 \in \overline{M}$ independent random elements with distribution ν. Let Λ be the subgroup generated by $(\alpha_1, \alpha_2, \alpha_3)$. We have $\Lambda/\ell^{a+1}\overline{M} = \ell^a\overline{M}/\ell^{a+1}\overline{M}$ with probability at least $q^3/8 - 3\varepsilon$.*
2. *Let $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{O}$ be independent random elements such that $\alpha_i \bmod \mathbf{Z}+N\mathcal{O}$ follows the distribution ν. Let Λ be the lattice generated by $(1, \alpha_1, \alpha_2, \alpha_3)$. Then $\Lambda \otimes \mathbf{Z}_\ell = (\mathbf{Z}+\ell^a\mathcal{O}) \otimes \mathbf{Z}_\ell$ with probability at least $q^3/8 - 3\varepsilon$.*

Proof. **Item 1, with $\varepsilon = 0$.** For any $\alpha \in M$, let $\mathrm{lev}_\ell(\alpha)$ be the largest integer such that $\alpha \in R + \ell^a M$ when it exists, and $\mathrm{lev}_\ell(\alpha) = \infty$ otherwise. Let L be the event that $\mathrm{lev}_\ell(\alpha_i) = a$ for all $i \in \{1, 2, 3\}$. Note that the level is constant over any $\mathrm{SL}_2(R)$-conjugacy class, so conditional on L, the variables α_i are still $\mathrm{SL}_2(R)$-invariant. If L occurs, the random variables $\alpha_i \bmod \ell^{a+1}\overline{M}$ are non-zero and $\mathrm{SL}_2(R)$-invariant in $\ell^a\overline{M}/\ell^{a+1}\overline{M} \cong M_2(\mathbf{F}_\ell)/\mathbf{F}_\ell$. The result follows from Lemma 5.3 and the fact that $\Pr[L] = q^3$.

Item 1, with $\varepsilon > 0$. By the triangular inequality, the triple $(\alpha_1, \alpha_2, \alpha_3)$ is 3ε-close to a triple of $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})$-invariant elements. The result thus follows from the case $\varepsilon = 0$ and the defining property of the statistical distance.

Item 2. This is the combination of Item 1 with Proposition 5.6. □

Proposition 5.11. *Assume that N is not a cube. Let $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{O}$ be three independent random elements from a distribution α that satisfies the following properties:*

(1) α is supported on N-reduced elements;
(2) $\alpha \bmod \mathbf{Z}+N\mathcal{O}$ is ε-close to $\mathrm{SL}_2(\mathbf{Z}/N\mathbf{Z})$-invariant for $\varepsilon < \frac{1}{6000000\cdot(\log N)^{12}}$.

Let Λ be the lattice generated by $(1, \alpha_1, \alpha_2, \alpha_3)$. With probability $\Omega\left((\log N)^{-12}\right)$, either $\gcd(N, [\mathcal{O} : \Lambda]) = 1$, or $[\mathcal{O} : \Lambda] = N^n K$ with $\gcd(N, K) \notin \{1, N\}$.

Remark 5.12. The exhibited event either produces a lattice Λ that is saturated at every prime factor of N (when $\gcd(N, [\mathcal{O} : \Lambda]) = 1$), or reveals a non-trivial factor of N.

Proof. Let Success be the event that either $\gcd(N, [\mathcal{O} : \Lambda]) = 1$, or $[\mathcal{O} : \Lambda] = N^n K$ where $\gcd(N, K) \notin \{1, N\}$. Write the prime factorisation $N = \prod_{i=1}^{t} \ell_i^{e_i}$. Since N is not a cube, we may assume without loss of generality that $\gcd(e_1, 3) = 1$. Write $\mathcal{O}_i = \mathcal{O} \otimes \mathbf{Z}_{\ell_i}$ and $\Lambda_i = \Lambda \otimes \mathbf{Z}_{\ell_i}$.

We now split the proof in two cases, depending on the value of

$$q_+ = \Pr[\mathrm{lev}_{\ell_1}(\alpha) \geq e_1].$$

Case 1: suppose $q_+^3 > 1 - \frac{1}{2}\left(\frac{1}{8\cdot(\log N)^3} - 3\varepsilon\right)$. Let ℓ_i and a_i be the ℓ and a from Lemma 5.8. Let $q = \Pr[\mathrm{lev}_{\ell_i}(\alpha) = a_i] > (\log N)^{-1}$. Let E be the event that $\Lambda_i = \mathbf{Z}_{\ell_i} + \ell_i^{a_i}\mathcal{O}_i$, and let F be the event that $\Lambda_1 \subseteq \mathbf{Z}_{\ell_1} + \ell_1^{e_1}\mathcal{O}_1$. Suppose E and F both happen. In that situation, $[\mathcal{O}_i : \Lambda_i] = \ell_i^{3a_i} < \ell_i^{3e_i}$, and $[\mathcal{O}_1 : \Lambda_1] \geq [\mathcal{O}_1 : \mathbf{Z}_{\ell_1} + \ell^{e_1}\mathcal{O}_1] \geq \ell_1^{3e_1}$, hence if $[\mathcal{O} : \Lambda] = N^n K$, then $\gcd(K, N) \notin \{1, N\}$. So if E and F both happen, then Success happens. We have

$$\Pr[F] = \Pr\left[\bigwedge_{j=1}^{3}(\mathrm{lev}_{\ell_1}(\alpha_j) \geq e_1)\right] = \prod_{j=1}^{3}\Pr\left[\mathrm{lev}_{\ell_1}(\alpha_i) \geq e_1\right] = q_+^3.$$

From Lemma 5.10, $\Pr[E] = q^3/8 - 3\varepsilon$. We deduce

$$\begin{aligned}\Pr[\mathsf{Success}] &\geq \Pr[E \wedge F] \geq \Pr[E] + \Pr[F] - 1 \\ &= \frac{q^3}{8} - 3\varepsilon + q_+^3 - 1 \geq \frac{3}{2}\left(\frac{1}{24\cdot(\log N)^3} - \varepsilon\right) \geq \frac{1}{32\cdot(\log N)^3}.\end{aligned}$$

Case 2: suppose $q_+^3 \leq 1 - \frac{1}{2}\left(\frac{1}{8\cdot(\log N)^3} - 3\varepsilon\right)$. Let $a_1 < e_1$ which maximises the probability $q = \Pr[\mathrm{lev}_{\ell_1}(\alpha) = a_1]$. We have $1 - q_+ = \sum_{x<e_1}\Pr[\mathrm{lev}_{\ell_1}(\alpha) = x] \leq e_1 q$. Then

$$q \geq \frac{1 - q_+}{\log(N)} \geq \frac{1 - q_+^3}{3\log(N)} \geq \frac{1}{48\cdot(\log N)^4} - \frac{\varepsilon}{2\log(N)}.$$

Let G be the event that $\Lambda_1 = \mathbf{Z}_{\ell_1} + \ell_1^{a_1} M_1$. If G happens and $[\mathcal{O} : \Lambda]$ is of the form $N^n K$ with $\gcd(N, K) = 1$, then $3a_1 = ne_1$. In that situation, $\gcd(e_1, 3) = 1$ implies that 3 divides n, and $a_1 < e_1$ implies that $n < 3$; together, these imply $n = 0$, so $\gcd(N, [\mathcal{O} : \Lambda]) = 1$. This proves that when G happens, then Success happens. We deduce

$$\begin{aligned}\Pr[\mathsf{Success}] \geq \Pr[G] \geq \frac{q^3}{8} - 3\varepsilon &= \left(\frac{1}{96\cdot(\log N)^4} - \frac{\varepsilon}{4\log(N)}\right)^3 - 3\varepsilon \\ &\geq 3\left(\frac{1}{3\cdot 100^3\cdot(\log N)^{12}} - \varepsilon\right) \geq \frac{1}{2000000\cdot(\log N)^{12}},\end{aligned}$$

which concludes the proof. □

6 Saturation and Reduction

In this section, we introduce three algorithms to saturate a known order of endomorphisms of a supersingular curve, and to reduce an endomorphism (in

the sense of Definition 5.7). The overall strategy is folklore, but only a crucial new ingredient allows it to work in polynomial time: the division algorithm due to Robert [Rob22] (see Proposition 2.3).

Let us start with saturation, which is used to deal with problematic primes in the main reduction.

Proposition 6.1. *There exists an algorithm* $\textsc{Saturate}_\ell$ *that takes as input a suborder* R_0 *of* $\mathrm{End}(E)$ *and a prime* ℓ *and returns an order* $R \subset \mathrm{End}(E)$ *containing* R_0 *such that* $[\mathrm{End}(E) : R]$ *is coprime to* ℓ*, and that runs in time polynomial in* ℓ *and the size of the input.*

Proof. The proof is available in the full version [PW23, Proposition 6.1].

Proposition 6.2. *There exists an algorithm* SaturateRam *that takes as input a suborder* R_0 *of* $\mathrm{End}(E)$ *and returns an order* $R \subset \mathrm{End}(E)$ *containing* R_0 *such that* $[\mathrm{End}(E) : R]$ *is coprime to* p*, and that runs in polynomial time.*

Proof. The proof is available in the full version [PW23, Proposition 6.2].

We now present an algorithm to reduce endomorphisms at odd integers.

Algorithm 2. $\textsc{Reduce}_N(\alpha)$: reduces an endomorphism α at N.

Require: An endomorphism $\alpha \in \mathrm{End}(E) \setminus \mathbf{Z}$ in efficient representation, and an odd integer N.
Ensure: An N-reduced endomorphism (Definition 5.7) $\beta = \frac{\alpha - t}{N^e}$ with $t, e \in \mathbf{Z}$.
1: $\gamma \leftarrow 2\alpha - \mathrm{Tr}(\alpha)$
2: **repeat**
3: $\quad \beta \leftarrow \gamma$
4: $\quad \gamma \leftarrow \textsc{Divide}(\beta, N)$ an efficient representation of β/N {Proposition 2.3}
5: **until** $\gamma = \perp$
6: **if** $\mathrm{Tr}(\alpha) \equiv 0 \mod 2$ **then**
7: $\quad$ **return** $\textsc{Divide}(\beta, 2)$
8: **else**
9: $\quad$ **return** $\textsc{Divide}(\beta + 1, 2)$
10: **end if**

Proposition 6.3. *Algorithm 2 (*$\textsc{Reduce}_N$*) is correct and runs in polynomial time.*

Proof. Let e be the largest integer such that $\alpha \in \mathbf{Z} + N^e \mathrm{End}(E)$. At Step 1, we have that $\gamma \in N^e \mathrm{End}(E)$ and $\gamma \notin \mathbf{Z} + N^{e+1} \mathrm{End}(E)$. Therefore, at the end of the loop, $\beta \in \mathrm{End}(E)$ and $\beta \notin \mathbf{Z} + N \mathrm{End}(E)$, i.e., β is N-reduced. The last division removes the extra factor 2 introduced in Step 1, to ensure the result is of the form $\beta = \frac{\alpha - t}{N^e}$ with $t \in \mathbf{Z}$.

Let us prove that it runs in polynomial time. We have $N^{2e} \mid \mathrm{disc}(\alpha)$, and at each iteration of the loop, $\mathrm{disc}(\beta)$ gets divided by N^2. So the number of iterations is bounded by $e \leq \log(\mathrm{disc}(\alpha)) = O(\log \deg(\alpha))$, which concludes the proof. □

7 The Reduction

In this section, we prove the main result of the paper (Theorem 1.1). We start with a lemma putting together results from the previous sections.

Algorithm 3. Turning an oracle $\mathscr{O}$ for OneEnd into an EndRing algorithm

Require: A supersingular elliptic curve $E/\mathbf{F}_{p^2}$, and a parameter $k > 0$. We suppose access to an oracle $\mathscr{O}$ that solves the *ProblemOneEnd* problem.
Ensure: The endomorphism ring $\mathrm{End}(E)$.

$k_1 \leftarrow \left\lceil \frac{\log\left(12 \cdot 9 \cdot (1+\sqrt{3}) \cdot \sqrt{p+13}\right)}{\log\left(\frac{3}{2\sqrt{2}}\right)} \right\rceil$
$R \leftarrow \mathbf{Z}$
while $\mathrm{rank}_{\mathbf{Z}}(R) \neq 4$ **do**
  $\alpha \leftarrow \textsc{Rich}^{\mathscr{O}}_{k_1}(E)$, a random endomorphism of E {Algorithm 1}
  $R \leftarrow$ the ring generated by R and α
end while
$R \leftarrow \textsc{Saturate}_2(R)$ {Proposition 6.1}
$R \leftarrow \textsc{SaturateRam}(R)$ {Proposition 6.2}
$[\mathrm{End}(E) : R] \leftarrow \sqrt{\mathrm{disc}(R)/p}$
Factor $[\mathrm{End}(E) : R] = \prod_{i=1}^{t} N_i^{e_i}$ where no N_i is a cube {a complete prime factorisation is not required; the somewhat trivial factorisation $[\mathrm{End}(E) : R] = N_1^{3^n}$ where $N_1^{1/3} \notin \mathbf{Z}$ and $n \geq 0$ is sufficient as a starting point, and the subsequent steps of the algorithm may refine it}
while $[\mathrm{End}(E) : R] \neq 1$ **do**
  $N \leftarrow N_t$
  $k_2 \leftarrow \left\lceil 12 \cdot \log\left(4100000 \cdot (\log N)^{12} N^2 \sqrt{p+13}\right) \right\rceil$
  Let $\mathscr{O}_N$ the oracle which given E, runs $\alpha \leftarrow \mathscr{O}(E)$ and returns $\textsc{Reduce}_N(\alpha)$
  $\alpha_i \leftarrow \textsc{Rich}^{\mathscr{O}_N}_{k_2}(E)$ for $i \in \{1, 2, 3\}$, random endomorphisms of E {Algorithm 1}
  $\Lambda \leftarrow$ the lattice generated by $(1, \alpha_1, \alpha_2, \alpha_3)$
  if $\mathrm{rank}_{\mathbf{Z}}(\Lambda) = 4$ **then**
    $n \leftarrow$ the largest integer such that N^n divides $[\mathrm{End}(E) : \Lambda]$
    $d \leftarrow \gcd([\mathrm{End}(E) : \Lambda]/N^n, N)$
    if $d \neq 1$ **then**
      Update the factorisation of $[\mathrm{End}(E) : R]$ with $N = d \cdot (N/d)$
    end if
    if $\Lambda \not\subset R$ **then**
      $R \leftarrow$ the order generated by R and Λ
      Recompute $[\mathrm{End}(E) : R] = \sqrt{\mathrm{disc}(R)/p}$, and update its factorisation
    end if
  end if
end while
return R

Lemma 7.1. *Let $\mathscr{O}$ be an oracle for* OneEnd*, and N an odd integer. Let $\mathscr{O}_N$ be the oracle which on input E, samples $\alpha \leftarrow \mathscr{O}(E)$, and returns* $\textsc{Reduce}_N(\alpha)$.

For any

$$k \geq 12 \cdot \log\left(4100000 \cdot (\log N)^{12} N^2 \sqrt{p+13}\right),$$

the output of $\text{RICH}_k^{\mathcal{O}_N}$ *satisfies the conditions of Proposition 5.11.*

Proof. Let $\varphi\colon E \to E'$ of degree a power of 2. For any endomorphism $\beta \in \text{End}(E')$, since N is odd, we have that β is N-reduced if and only if $\hat{\varphi} \circ \beta \circ \varphi$ is N-reduced. The output of $\text{RICH}_k^{\mathcal{O}_N}$ is of the form $\hat{\varphi} \circ \text{REDUCE}_N(\alpha) \circ \varphi$, so is N-reduced. So the distribution of $\text{RICH}_k^{\mathcal{O}_N}$ satisfies Item 1 of Proposition 5.11.

From Theorem 4.2, $\text{RICH}_k^{\mathcal{O}_N}$ mod N is ε-close to $\text{SL}_2(\mathbf{Z}/N\,\mathbf{Z})$-invariant for

$$\varepsilon = \frac{1+\sqrt{3}}{4}\left(\frac{2\sqrt{2}}{3}\right)^k N^2\sqrt{p+13}.$$

With $k \geq \log\bigl(6000000 \cdot (\log N)^{12} \cdot \frac{1+\sqrt{3}}{4} N^2\sqrt{p+13}\bigr)/\log\bigl(\frac{3}{2\sqrt{2}}\bigr)$, we have $\varepsilon \leq (6000000 \cdot (\log N)^{12})^{-1}$, satisfying Item 2 of Proposition 5.11. □

We now have all the ingredients to prove our main result.

Theorem 7.2 (ENDRING reduces to ONEEND). *Algorithm 3 is a reduction from* ENDRING *to* ONEEND_λ *of expected polynomial time in* $\log(p)$ *and* $\lambda(\log p)$.

Proof. The correctness is clear as at any time, R is a subring of $\text{End}(E)$, and the success condition $[\text{End}(E) : R] = 1$ implies $R = \text{End}(E)$.

We now analyse the expected running time.

First loop (Step 3 to Step 6). First, let us analyse the expected number of iterations of the first loop. From Theorem 4.2, each α generated during this loop is ε-close to $\text{SL}_2(\mathbf{F}_3)$-invariant with $\varepsilon = \frac{1+\sqrt{3}}{4}\left(\frac{2\sqrt{2}}{3}\right)^{k_1} 3^2\sqrt{p+13}$. Choosing $k_1 = O(\log p)$ as in Step 1, we have $\varepsilon \leq 1/48$.

Consider any three consecutively generated elements $\alpha_1, \alpha_2, \alpha_3$. Let $t = \max_i \text{lev}_3(\alpha_i)$, and $\beta_i = 3^{t-\text{lev}_3(\alpha_i)}\alpha_i$, so all β_i are at the same level t. Like the variables α_i, the variables β_i are ε-close to $\text{SL}_2(\mathbf{F}_3)$-invariant. Combining Proposition 5.4 and Proposition 5.6, the tuple $(1, \beta_1, \beta_2, \beta_3)$ generates a full-rank lattice with probability at least $1/8 - 3\varepsilon$, and so does $(1, \alpha_1, \alpha_2, \alpha_3)$. Choosing k_1 as above, this probability is at least $1/16$. We deduce that the loop terminates after an expected $O(1)$ number of iterations.

Let us now analyse the output of this loop. Let R_1 be the order R obtained at the end of the first loop. Let α_i be any three elements generated during the loop such that $(1, \alpha_1, \alpha_2, \alpha_3)$ are independent. Combining the bound $\deg(\alpha_i) \leq 2^{2k_1\lambda(\log p)}$ and Hadamard's inequality, we get

$$\text{disc}(R_1) = 16 \cdot \text{Vol}(R_1)^2 \leq 16 \cdot \prod_{i=1}^{3} \sqrt{\deg(\alpha_i)} \leq 16 \cdot 2^{6k_1\lambda(\log p)}.$$

We deduce that

$$[\mathrm{End}(E) : R_1] \leq 2^{3k_1\lambda(\log p)+2}/p = 2^{O(\log(p)\cdot\lambda(\log p))}. \tag{1}$$

Second loop (Step 11 to Step 28). It remains to analyse the second loop. An iteration of this loop is a *success* if either Step 21 or Step 24 is reached. In case of success, either a new factor of $[\mathrm{End}(E) : R]$ is found (Step 21), or $[\mathrm{End}(E) : R]$ gets divided by an integer at least 2 (Step 24). The number of successes is thus polynomially bounded in $\log([\mathrm{End}(E) : R_1])$, hence in $\mathrm{poly}(\log p, \lambda(\log p))$ (thanks to Eq. (1)). Therefore, we only have to prove that as long as $R \neq \mathrm{End}(E)$, each iteration has a good probability of success.

The event analysed in Proposition 5.11 corresponds precisely to a success. By Lemma 7.1, the distribution of α_i satisfies the conditions of Proposition 5.11. Therefore, Proposition 5.11 implies that each iteration has a probability of success $\Omega((\log N)^{-12})$, which concludes the proof. □

8 Applications

In this section we describe four applications of our main result.

8.1 Collision Resistance of the Charles–Goren–Lauter Hash Function

The first cryptographic construction based on the supersingular isogeny problem is a hash function proposed by Charles, Goren and Lauter [CLG09], the *CGL hash function*. Fix a (small) prime number ℓ, typically $\ell = 2$. For any elliptic curve E, there are $\ell + 1$ outgoing ℓ-isogenies $E \to E'$ (up to isomorphism of the target), so given a curve and an incoming $E'' \to E$, there remain ℓ non-backtracking ℓ-isogenies from E, which can be arbitrarily labelled by the set $\{0, \ldots, \ell - 1\}$. Then, fixing an initial curve E_0 and an arbitrary isogeny $E_{-1} \to E_0$, the set $\{0, \ldots, \ell - 1\}^*$ encodes non-backtracking paths from E_0 in the ℓ-isogeny graph. The CGL hash function

$$\mathrm{CGL}_{E_0} : \{0, \ldots, \ell - 1\}^* \longrightarrow \mathbf{F}_{p^2}$$

associates to any sequence $(x_i)_i$ the j-invariant of the endpoint of the walk from E_0 it encodes. Clearly, this function is pre-image resistant if and only if ℓ-IsogenyPath is hard. However, if $\mathrm{End}(E_0)$ is known, one can find collisions in polynomial time [KLPT14, EHL+18]. Therefore, it was proposed to sample the starting curve randomly. Let SampleSS(p) be an algorithm sampling a uniformly random supersingular elliptic curve over $\mathbf{F}_{p^2}$. We define the advantage of a collision-finding algorithm $\mathscr{A}$ for the CGL family of hash functions as

$$\mathrm{Adv}^{\mathscr{A}}_{\mathrm{CGL}}(p) = \Pr\left[\begin{array}{c} m \neq m' \text{ and} \\ \mathrm{CGL}_E(m) = \mathrm{CGL}_E(m') \end{array} \middle| \begin{array}{c} E \leftarrow \text{SampleSS}(p) \\ (m, m') \leftarrow \mathscr{A}(E) \end{array}\right].$$

It was heuristically argued in [EHL+18] that the collision resistance of this construction is equivalent to EndRing. The flaws of the heuristics are discussed in Sect. 1.2. With our main theorem, we can now prove this resistance.

Theorem 8.1 (Collision resistance of the CGL hash function). *For any algorithm $\mathscr{A}$, there is an algorithm to solve* EndRing *in expected polynomial time in* $\log(p)$*, in* $\mathrm{Adv}^{\mathscr{A}}_{\mathrm{CGL}}(p)^{-1}$ *and in the expected running time of* $\mathscr{A}$*.*

Proof. Since EndRing is equivalent to OneEnd (Theorem 1.1), it is sufficient to prove that $\mathscr{A}$ can be used to solve OneEnd. First, let us prove that a successful collision for CGL_E gives a non-scalar endomorphism of E. Let $\varphi, \psi\colon E \to E'$ be two distinct non-backtracking walks, i.e., isogenies of cyclic kernel of order ℓ^a and ℓ^b respectively. If $\hat{\varphi} \circ \psi$ is scalar, the degrees imply that $a + b$ is even and $\hat{\varphi} \circ \psi = [\ell^{\frac{a+b}{2}}]$. Without loss of generality, suppose $b \geq a$. From the defining property of the dual isogeny, we deduce that $\hat{\psi} = [\ell^{\frac{b-a}{2}}]\hat{\varphi}$. Taking the dual again, we get $\psi = [\ell^{\frac{b-a}{2}}]\varphi$. If $b > a$, then $\{0_E\} \neq E[\ell^{\frac{b-a}{2}}] \subseteq \ker\psi$, contradicting the cyclicity of $\ker\psi$. Therefore $b = a$, and we conclude that $\psi = \varphi$, a contradiction. So $\hat{\varphi} \circ \psi$ is non-scalar.

Now, given a curve E, we can solve OneEnd as follows:

1. First take a random walk $\eta\colon E \to E'$, so that E' has statistical distance $\varepsilon = O(1/p)$ from uniform (Proposition 2.7);
2. Then call $\mathscr{A}(E')$, which gives a non-scalar endomorphism α of E' with probability at least $\mathrm{Adv}^{\mathscr{A}}_{\mathrm{CGL}}(p) - \varepsilon$,
3. Return $\hat{\eta} \circ \alpha \circ \eta$.

The algorithm is successful after an expected $(\mathrm{Adv}^{\mathscr{A}}_{\mathrm{CGL}}(p) - \varepsilon)^{-1}$ number of attempts. This works within the claimed running time if $\mathrm{Adv}^{\mathscr{A}}_{\mathrm{CGL}}(p) > 2\varepsilon$. Otherwise, we have $(\mathrm{Adv}^{\mathscr{A}}_{\mathrm{CGL}}(p))^{-1} = \Omega(p)$, and one can indeed solve EndRing in time polynomial in p (see [Koh96, Theorem 75] for the first such algorithm, in time $\tilde{O}(p)$, or Theorem 8.8 below for time $\tilde{O}(p^{1/2})$). □

8.2 Soundness of the SQIsign Identification Scheme

SQIsign is a digital signature scheme proposed in [DKL+20]. SQIsign and its variants offer the most compact public keys and signatures of all known post-quantum constructions.

Each digital signature scheme in this family is constructed as an identification protocol, turned into a signature by the Fiat–Shamir transform. The protocol proves knowledge of a witness for a problem that closely resembles OneEnd. While [DKL+20] or [DLRW23] heuristically argue that the protocol is sound if EndRing is hard, our main theorem provides a way to prove it.

Among the several variants of SQIsign, we illustrate our result on RigorousSQIsignHD [DLRW23], which currently benefits from the cleanest security arguments. Note that the method could be applied to each variant, but the resulting statements may vary. In particular, a version of Theorem 8.2 below holds for the original SQIsign design if we assume that public keys are computationally indistinguishable from the uniform distribution on supersingular elliptic curves.

Let SQIsignHD.param be the RigorousSQIsignHD public parameter generation procedure, which on input a security level k, outputs data pp

which encodes, among other things, a prime number $p = \Theta(2^{2k})$. Let SQIsignHD.KeyGen be the RigorousSQIsignHD key generation procedure, which on input pp, outputs a pair $(\mathsf{pk}, \mathsf{sk})$. The public key pk is a supersingular elliptic curve over $\mathbf{F}_{p^2}$, and sk is its endomorphism ring. The random elliptic curve pk is at statistical distance $\tilde{O}(p^{-1/2})$ from uniform [DLRW23, Appendix B.2].

Let $\mathscr{V}$ be a honest verifier for the RigorousSQIsignHD identification protocol. For any (malicious) prover $\mathscr{P}^*$ and parameters pp, run the following experiment: first, sample a key pair $(\mathsf{pk}, \mathsf{sk}) \leftarrow \text{SQIsignHD.KeyGen}(\mathsf{pp})$, and give pk to $\mathscr{P}^*$. Then, run the RigorousSQIsignHD identification protocol between $\mathscr{P}^*$ and $\mathscr{V}$ with input pk. Let $\pi^{\mathscr{P}^*}(\mathsf{pp})$ be the probability that $\mathscr{V}$ outputs $\top$ at the end of the protocol. We define the *soundness advantage* $\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp}) = \pi^{\mathscr{P}^*}(\mathsf{pp}) - 1/c$, where $c = \Theta(2^k)$ is the size of the challenge space.

In other words, $\pi^{\mathscr{P}^*}(\mathsf{pp})$ is the probability that $\mathscr{P}^*$ successfully fools a honest verifier, for a random key. Since there is a simple malicious prover achieving $\pi^{\mathscr{P}^*}(\mathsf{pp}) = 1/c$ (by guessing the challenge at the start of the protocol), the advantage $\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp})$ measures how much better $\mathscr{P}^*$ performs.

Theorem 8.2 (Soundness of RigorousSQIsignHD). *Let $\mathscr{P}^*$ be a malicious prover. Consider public parameters pp, encoding the prime p. There is an algorithm to solve* EndRing *for curves over $\mathbf{F}_{p^2}$ in expected polynomial time in $\log(p)$, in $\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp})^{-1}$ and in the expected running time of $\mathscr{P}^*$.*

Proof. Let r denote the expected running time of $\mathscr{P}^*$. Let $\pi^{\mathscr{P}^*}(\mathsf{pp}, E)$ be the probability that $\mathscr{V}$ outputs $\top$ at the end of the protocol given that $\mathsf{pk} = E$. For simplicity, we write $\pi = \pi^{\mathscr{P}^*}(\mathsf{pp})$ and $\pi_E = \pi^{\mathscr{P}^*}(\mathsf{pp}, E)$. Let $\varepsilon_{\mathsf{pk}} = \tilde{O}(p^{-1/2})$ be the statistical distance of pk from uniform.

As mentionned in the proof of Theorem 8.1, one can solve EndRing in time polynomial in p (see [Koh96, Theorem 75], or Theorem 8.8 below), so we may assume that $\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp}) > 2/c + 4\varepsilon_{\mathsf{pk}} = \tilde{O}(p^{-1/2})$. In particular, we have $\pi > 3/c$, and $\pi > 4\varepsilon_{\mathsf{pk}}$.

From [DLRW23, Proposition D.1.7], RigorousSQIsignHD is a proof of knowledge with soundness error $1/c$ for the relation $\{(E, \alpha) \mid \alpha \in \mathrm{End}(E) \setminus \mathbf{Z}\}$ (and witnesses of size $O(\log p)$), hence, whenever $\pi_E > 1/c$, there is an algorithm which solves OneEnd on E with expected running time $R(E) = O\left(\frac{r}{\pi_E - 1/c}\right)$. We get an algorithm for OneEnd on arbitrary input E as follows:

1. First take a random walk $\eta\colon E \to E'$, so that E' has statistical distance $\varepsilon = O(1/p^2)$ from uniform (Proposition 2.7);
2. Run the above algorithm for OneEnd on input E', which has expected time $R(E')$ (if $\pi_{E'} > 1/c$). Upon termination, go to the step 3. While this runs, repeat from step (1) in a new concurrent session.
3. When one of the concurrent sessions terminates, say with $\eta\colon E \to E'$ and $\alpha \in \mathrm{End}(E')$, stop all other sessions, and return $\hat{\eta} \circ \alpha \circ \eta$.

The purpose of running concurrent sessions is that any single E' is not guaranteed to terminate. Let $X = \{E \mid \pi_E \geq \pi/2\} \subseteq \{E \mid \pi_E > 1/c\}$. We have

$$\pi \leq \Pr[\mathsf{pk} \in X] + \Pr[\mathsf{pk} \notin X]\frac{\pi}{2} = \left(1 - \frac{\pi}{2}\right)\Pr[\mathsf{pk} \in X] + \frac{\pi}{2},$$

hence $\Pr[\mathsf{pk} \in X] \geq \frac{\pi - \pi/2}{1 - \pi/2} \geq \pi/2$. Therefore, the curve E' generated in step 1 falls in X with probability at least $\pi/2 - \varepsilon - \varepsilon_{\mathsf{pk}} > \pi/4 - \varepsilon$. Whenever E' is in X, the expected running time $R(E')$ for that session is about

$$\frac{r}{\pi_{E'} - 1/c} \leq \frac{r}{\pi/2 - 1/c} = \frac{4r}{\pi + (\pi - 4/c)} = \frac{4r}{\pi + 1/c} = \frac{4r}{\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp})}.$$

We conclude that this algorithm for OneEnd runs in expected polynomial time in $\log(p)$, in $\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp})^{-1}$ and in r. The result follows from the equivalence between OneEnd and EndRing (Theorem 1.1). □

Remark 8.3. Note that the proof can be adapted to prove the theorem with the quantity $\pi^{\mathscr{P}^*}(\mathsf{pp})^{-1}$ in place of $\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp})^{-1}$, which may be more natural.

8.3 The Endomorphism Ring Problem is Equivalent to the Isogeny Problem

It is known that the problem EndRing is equivalent to the ℓ-isogeny path problem (assuming the generalised Riemann hypothesis [Wes22b]). The same technique shows that EndRing is equivalent to the problem of finding isogenies of *smooth* degree. Lifting this restriction yields the more general Isogeny problem.

Problem 8.4 (Isogeny). *Given a prime p and two supersingular elliptic curves E and E' over $\mathbf{F}_{p^2}$, find an isogeny from E to E' in efficient representation.*

Given a function $\lambda\colon \mathbf{Z}_{>0} \to \mathbf{Z}_{>0}$, the Isogeny_λ problem denotes the Isogeny problem where the solution φ is required to satisfy $\log(\deg\varphi) \leq \lambda(\log p)$ (the length of the output is bounded by a function of the length of the input).

From previous literature, it is easy to see that Isogeny reduces to EndRing.

Proposition 8.5 (Isogeny reduces to EndRing). *Assuming the generalised Riemann hypothesis, the problem Isogeny_λ reduces to EndRing in probabilistic polynomial time (with respect to the length of the instance), for some function $\lambda(\log p) = O(\log p)$.*

Proof. Isogeny immediately reduces to ℓ-IsogenyPath. It is already known that the ℓ-isogeny path problem (with paths of length $O(\log p)$) is equivalent to EndRing [Wes22b], so Isogeny_λ reduces to EndRing. □

The converse reduction is trickier. As a solution to Isogeny is not guaranteed to have smooth degree, previous techniques have failed to prove that it is equivalent to EndRing. Theorem 1.1 unlocks this equivalence. Better yet, contrary to previous results of this form, Theorem 8.6 below is unconditional. In particular, it implies that EndRing reduces to the ℓ-isogeny path problem independently of the generalised Riemann hypothesis.

Theorem 8.6 (EndRing reduces to Isogeny). *Given an oracle for* Isogeny_λ*, there is an algorithm for* EndRing *that runs in expected polynomial time in* $\log(p)$ *and* $\lambda(\log p)$*.*

Algorithm 4. Solving OneEnd given an Isogeny oracle.

Require: A supersingular elliptic curve $E/\mathbf{F}_{p^2}$, a parameter $\varepsilon > 0$, an oracle $\mathscr{O}_{ProblemIsogeny}$ solving the $ProblemIsogeny_\lambda$ problem.
Ensure: An endomorphism $\alpha \in \text{End}(E) \setminus \mathbf{Z}$ in efficient representation.
1: $S \leftarrow$ an arbitrary nonzero point in $E[2]$
2: $n \leftarrow \lceil 2\log_3(p) - 4\log_3(\varepsilon) \rceil$
3: **while** true **do**
4: $\quad \varphi \leftarrow$ a non-backtracking random walk $\varphi\colon E \to E'$ of length n in the 3-isogeny graph
5: $\quad \nu \leftarrow$ the isogeny $\nu\colon E' \to E''$ of kernel $\langle \varphi(S) \rangle$
6: $\quad \psi \leftarrow \mathscr{O}_{ProblemIsogeny}(E'', E)$, an isogeny $\psi\colon E'' \to E$
7: $\quad \alpha \leftarrow (\psi \circ \nu \circ \varphi)/2^e \in \text{End}(E)$ for the largest possible e
8: $\quad$ **if** $2 \mid \deg(\alpha)$ **then**
9: $\quad\quad$ **return** α
10: $\quad$ **end if**
11: **end while**

Proof. Since EndRing is equivalent to OneEnd (Theorem 1.1), let us prove that OneEnd reduces to Isogeny. Suppose we have an oracle $\mathscr{O}_{\text{Isogeny}}$ for Isogeny_λ. Let E be a supersingular curve for which we want to solve OneEnd. Consider a parameter ε. The reduction is described in Algorithm 4. Step 7 and Step 8 ensure that α is not a scalar (indeed, they ensure that upon return, at Step 9, we have $2 \nmid \alpha$ yet $2 \mid \deg(\alpha)$), so is a valid solution to OneEnd.

Let us show that the expected number of iterations of the while-loop is $O(1)$. Let $f \in \mathbf{Z}$ maximal such that $E''[2^f] \subseteq \ker(\psi)$, and let $\psi' = \psi/2^f$. If $\deg(\psi')$ is odd, then α is non-scalar (its degree is divisible by 2 but not by 4) and the loop terminates at this iteration. Now, suppose $\deg(\psi')$ is even and write $\ker(\psi') \cap E''[2] = G_\psi$, a group of order 2. The loop in the reduction terminates in the event that $\ker \hat{\nu} \neq G_\psi$. In the rest of the proof, we bound the probability of this event at each iteration.

Let P be the probability distribution of the pair $(E'', \hat{\nu})$, and Q the probability distribution of the pair (E'', η) where η is uniformly random (among the three 2-isogenies from E''). Note that by construction, the value $Q(E'', \eta)$ does

not depend on η, and we also write it $\tilde{Q}(E'')$. Consider the function τ defined in [BCC+23, Lemma 14]. We have

$$\tau(p,2,3,k) = \frac{1}{4}(p-1)^{1/2}\left(1+\sqrt{3}\right)\left(k+\frac{1}{2}\right)3^{-k/2} \leq p^{1/2}3^{-k/4}.$$

From [BCC+23, Lemma 14], if $\tau(p,2,3,k) \leq \varepsilon$, then the statistical distance $\|P-Q\|_1/2$ is at most ε. This condition is satisfied if the 3-walk φ has length at least

$$\begin{aligned} n(p,2,3,\varepsilon) &= \min\{k \mid \tau(p,2,3,k) \leq \varepsilon\} \\ &\leq \min\{k \mid p^{1/2}3^{-k/4} \leq \varepsilon\} = 2\log_3(p) - 4\log_3(\varepsilon). \end{aligned}$$

We deduce that indeed $\|P-Q\|_1 < 2\varepsilon$, since φ has length $\lceil 2\log_3(p) - 4\log_3(\varepsilon)\rceil$.

We now obtain the following bound:

$$\begin{aligned} \Pr[\ker\hat{\nu} = G_\psi] &= \sum_{(E'',\hat{\nu})} P(E'',\hat{\nu})\Pr[\ker\hat{\nu} = G_\psi \mid (E'',\hat{\nu})] \\ &\leq \sum_{(E'',\hat{\nu})} (Q(E'',\hat{\nu}) + \max_\eta |P(E'',\eta) - Q(E'',\eta)|)\Pr[\ker\hat{\nu} = G_\psi \mid E''] \\ &\leq \sum_{E''} \tilde{Q}(E'') + \sum_{E''} \max_\eta |P(E'',\eta) - Q(E'',\eta)| \leq \frac{1}{3} + 2\varepsilon. \end{aligned}$$

The second line uses that for any fixed E'', the distribution of ψ is independent of ν. In conclusion, at each iteration, the event $\ker\hat{\nu} \neq G_\psi$ (leading to termination) happens with probability at least $2/3 - 2\varepsilon$. With $2\varepsilon < 1/3$, the expected number of iterations is at most $(2/3-2\varepsilon)^{-1} \leq 3 = O(1)$. □

8.4 An Unconditional Algorithm for EndRing in Time $\tilde{O}(p^{1/2})$

As the foundational problem of isogeny-based cryptography, understanding the hardness of EndRing is critical. The fastest known algorithms have complexity in $\tilde{O}(p^{1/2})$, but rely on unproven assumptions such as the generalised Riemann hypothesis. With our new results, we can now prove that EndRing can be solved in time $\tilde{O}(p^{1/2})$ *unconditionally*. In contrast, the previous fastest unconditional algorithm had complexity $\tilde{O}(p)$ and only returned a full-rank subring of the endomorphism ring [Koh96, Theorem 75].

The first method to reach complexity $\tilde{O}(p^{1/2})$ under the generalised Riemann hypothesis consists in reducing EndRing to ℓ-IsogenyPath (via [Wes22b]), and solving ℓ-IsogenyPath by a generic graph path-finding algorithm. Unconditionally, we can follow the same strategy, but using our new reduction from EndRing to ℓ-IsogenyPath (Theorem 8.6). Let us start by recalling the following folklore solution to ℓ-IsogenyPath.

Proposition 8.7. *There exists an algorithm that solves the* ℓ-IsogenyPath *problem in expected time* $\mathrm{poly}(\ell, \log p)p^{1/2}$ *and returns paths of length* $O(\log p)$.

Proof. The proof is available in the full version [PW23, Proposition 8.7].

Theorem 8.8. *There is an algorithm solving* EndRing *in expected time* $\tilde{O}\left(p^{1/2}\right)$.

Proof. This follows from the fact that there is an algorithm of complexity $\tilde{O}\left(p^{1/2}\right)$ for the 2-isogeny path problem (Proposition 8.7), and EndRing reduces to polynomially many instances of the ℓ-isogeny path problem (Theorem 8.6). □

Acknowledgements. The authors would like to thank Damien Robert for fruitful discussions. The authors were supported by the Agence Nationale de la Recherche under grant ANR-20-CE40-0013 (MELODIA) and ANR-19-CE48-0008 (CIAO), the France 2030 program under grant ANR-22-PETQ-0008 (PQ-TLS), and the European Research Council under grant No. 101116169 (AGATHA CRYPTY).

References

Arp23. Arpin, S.: Adding level structure to supersingular elliptic curve isogeny graphs. Preprint arXiv:2203.03531 (2023). https://arxiv.org/abs/2203.03531

BCC+23. Basso, A., et al.: Supersingular curves you can trust. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023. LNCS, vol. 14005, pp. 405–437. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30617-4_14

CD23. Castryck, W., Decru, T.: An efficient key recovery attack on SIDH. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 423–447. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30589-4_15

CL23. Codogni, G., Lido, G.: Spectral theory of isogeny graphs. Preprint arXiv:2308.13913 (2023). https://arxiv.org/abs/2308.13913

CLG09. Charles, D.X., Lauter, K.E., Goren, E.Z.: Cryptographic hash functions from expander graphs. J. Cryptol. **22**(1), 93–113 (2009)

DKL+20. De Feo, L., Kohel, D., Leroux, A., Petit, C., Wesolowski, B.: SQISign: compact post-quantum signatures from quaternions and isogenies. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12491, pp. 64–93. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64837-4_3

DLRW23. Dartois, P., Leroux, A., Robert, D., Wesolowski, B.: SQISignHD: new dimensions in cryptography. IACR Cryptology ePrint Archive, Report 2023/436 (2023). https://eprint.iacr.org/2023/436

EHL+18. Eisenträger, K., Hallgren, S., Lauter, K., Morrison, T., Petit, C.: Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10822, pp. 329–368. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78372-7_11

FIK+23. Fuselier, J., Iezzi, A., Kozek, M., Morrison, T., Namoijam, C.: Computing supersingular endomorphism rings using inseparable endomorphisms. Preprint arXiv:2306.03051 (2023). https://arxiv.org/abs/2306.03051

HLMW23. Le Merdy, A.H., Wesolowski, B.: The supersingular endomorphism ring problem given one endomorphism. Preprint arXiv:2309.11912 (2023). https://arxiv.org/abs/2309.11912

KLPT14. Kohel, D., Lauter, K., Petit, C., Tignol, J.-P.: On the quaternion ℓ-isogeny path problem. LMS J. Comput. Math. **17**(A), 418–432 (2014)

Koh96. Kohel, D.: Endomorphism rings of elliptic curves over finite fields. Ph.D. thesis, University of California, Berkeley (1996)

Mes86. Mestre, J.-F.: La méthode des graphes. Exemples et applications. In: Proceedings of the International Conference on Class Numbers and Fundamental Units of Algebraic Number Fields (Katata), pp. 217–242 (1986)

ML98. Lane, S.M.: Categories for the Working Mathematician. Graduate Texts in Mathematics, 2nd edn, vol. 5. Springer, New York (1998). https://doi.org/10.1007/978-1-4757-4721-8

MMP+23. Maino, L., Martindale, C., Panny, L., Pope, G., Wesolowski, B.: A direct key recovery attack on SIDH. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023. LNCS, vol. 14008, pp. 448–471. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30589-4_16

Piz90. Pizer, A.K.: Ramanujan Graphs and Hecke Operators. Bull. Am. Math. Soc. **23**(1), 127–137 (1990)

PW23. Page, A., Wesolowski, B.: The supersingular endomorphism ring and one endomorphism problems are equivalent. IACR Cryptology ePrint Archive, Report 2023/1399 (2023). https://eprint.iacr.org/2023/1399

Rob22. Robert, D.: Some applications of higher dimensional isogenies to elliptic curves (overview of results). Cryptology ePrint Archive, Paper 2022/1704 (2022). https://eprint.iacr.org/2022/1704

Rob23. Robert, D.: Breaking SIDH in polynomial time. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 472–503. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30589-4_17

Sil86. Silverman, J.H.: The Arithmetic of Elliptic Curves. Gradute Texts in Mathematics, vol. 106. Springer, Heidelberg (1986). https://doi.org/10.1007/978-0-387-09494-6

Voi21. Voight, J.: Quaternion Algebras. Graduate Texts in Mathematics, vol. 288. Springer, Heidelberg (2021). https://doi.org/10.1007/978-3-030-56694-4

Wes22a. Wesolowski, B.: Orientations and the supersingular endomorphism ring problem. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022. LNCS, vol. 13277, pp. 345–371. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07082-2_13

Wes22b. Wesolowski, B.: The supersingular isogeny path and endomorphism ring problems are equivalent. In: FOCS 2021-62nd Annual IEEE Symposium on Foundations of Computer Science (2022)