The supersingular Endomorphism Ring and One Endomorphism problems are equivalent

Aurel Page, Benjamin Wesolowski

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, F-33400 Talence, France
[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. 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.

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 Section 2, and the proof is the object of Section 7. The reduction from OneEnd to EndRing is obvious, and the other direction is stated more precisely in Theorem 7.2. 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 (Theorem 8.2). 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 Section 1.2). This is the object of Section 8.1 and Section 8.2.
- EndRing reduces to the isogeny problem (Theorem 8.5). 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 Section 8.3.
- There is an algorithm solving EndRing in expected time $\tilde{O}(p^{1/2})$ (Theorem 8.7), 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 Section 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, refering 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

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

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 $e \in [0,\dots,n-1]$. 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 $\textsc{Rich}^{\mathcal{O}}$ is indeed invariant by conjugation. It is a consequence of Theorem 1.3, our general equidistribution result. The proof of Theorem 1.3 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 (Theorem 3.27) 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 eigenvalues of the adjacency operators give the desired fast mixing result.

Acknowledgements

The authors would like to thank Damien Robert for discussions about this project. The authors were supported by the Agence Nationale de la Recherche under grant ANR MELODIA (ANR-20-CE40-0013), ANR CIAO (ANR-19-CE48-0008), and the France 2030 program under grant agreement No. ANR-22-PETQ-0008 PQ-TLS.

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 group law (defined algebraically) with neutral element O. An algebraic morphism between elliptic curves (preserving the base point) is automatically a group morphism, and is either a constant map or has a finite kernel. In the latter case, we say that the morphism is an *isogeny*. The *degree* of an isogeny φ, written $\deg(\varphi)$, is its degree as a rational map. If φ is separable, we have $\deg(\varphi) = \#\ker(\varphi)$. An isogeny of degree d is also 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* of an elliptic curve E is a morphism from 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 $\#\operatorname{Aut}(E) = 6$ and $\#\operatorname{Aut}(E) = 4$.

Assume that K has positive characteristic p and let E be an elliptic curve over K. We say that E is *supersingular* if $\operatorname{End}(E)$ is an order in a quaternion algebra. In this case, $B = \operatorname{End}(E) \otimes \mathbf{Q}$ is a quaternion algebra over $\mathbf{Q}$ with ramification set $\{p, \infty\}$, the ring $\operatorname{End}(E)$ is a maximal order in B, and E is defined over $\mathbf{F}_{p^2}$. When we see a nonzero endomorphism $\alpha \in \operatorname{End}(E)$ as a quaternion $a \in B$, we have $\deg(\alpha) = \operatorname{nrd}(a)$ and $\operatorname{Tr}(\alpha) = \operatorname{trd}(a)$.

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.

We will use the following proposition.

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 $\alpha \in N\operatorname{End}(E)$, 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_λ 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 $\operatorname{Hom}_\mathcal{C}(x, y)$, sometimes denoted $f\colon x \to y$, an associative composition law for morphisms with compatible source and target, and an identity morphism $\mathrm{id}_x \in \operatorname{Hom}_\mathcal{C}(x, x)$ for every object $x \in \mathcal{C}$. An *isomorphism* is a morphism that admits a two-sided inverse. For $x, y \in \mathcal{C}$, we define the set $\operatorname{End}_\mathcal{C}(x) = \operatorname{Hom}_\mathcal{C}(x, x)$ of *endomorphisms* of x, the set $\operatorname{Isom}_\mathcal{C}(x, y)$ of isomorphisms from x to y, the group $\operatorname{Aut}_\mathcal{C}(x) = \operatorname{Isom}_\mathcal{C}(x, x)$ of *automorphisms* of x. 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[3] and identities. Functors can be composed in the obvious way.

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 $\operatorname{El}(\mathcal{F})$ *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 $f \in \operatorname{Hom}_\mathcal{C}(c, c')$ s.t. $\mathcal{F}(f)(x) = x'$.*

This category is equipped with the natural forgetful functor $\operatorname{El}(\mathcal{F}) \to \mathcal{C}$.

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.

Let $\mathcal{F}, \mathcal{F}'\colon \mathcal{C} \to \mathcal{D}$ be functors. A *morphism of functors* $\psi\colon \mathcal{F} \to \mathcal{F}'$ is a collection $\psi = (\psi_x)_{x \in \mathcal{C}}$ of morphisms $\psi_x\colon \mathcal{F}(x) \to \mathcal{F}'(x)$ in $\mathcal{D}$ such that for every morphism $f\colon x \to y$ in $\mathcal{C}$, we have $\mathcal{F}'(f) \circ \psi_x = \psi_y \circ \mathcal{F}(f)$. A morphism of functors is an isomorphism if and only if every ψ_x is an isomorphism in $\mathcal{D}$. A

[3] All our functors are covariant.

functor $\mathcal{F}\colon \mathcal{C} \to \mathcal{D}$ is an *equivalence of categories* if there exists a functor $\mathcal{G}\colon \mathcal{D} \to \mathcal{C}$ and isomorphisms of functors $\mathcal{G}\circ\mathcal{F} \cong \mathrm{id}_{\mathcal{F}}$ and $\mathcal{F}\circ\mathcal{G} \cong \mathrm{id}_{\mathcal{D}}$. The functor $\mathcal{F}$ is *full* if all the corresponding maps $\mathrm{Hom}_{\mathcal{C}}(x, y) \to \mathrm{Hom}_{\mathcal{D}}(\mathcal{F}(x), \mathcal{F}(y))$ are surjective, and *faithful* if they are all injective. The functor $\mathcal{F}$ is *essentially surjective* if every object in $\mathcal{D}$ is isomorphic to the image of some object in $\mathcal{C}$ under $\mathcal{F}$. A functor is an equivalence if and only if it is full, faithful and essentially surjective. If $\mathcal{F}, \mathcal{F}'\colon \mathcal{C} \to \mathrm{Sets}$ are functors, every isomorphism of functors $\mathcal{F} \cong \mathcal{F}'$ induces an equivalence of categories $\mathrm{El}(\mathcal{F}) \cong \mathrm{El}(\mathcal{F}')$.

The categorical formulation of the Deuring correspondence [Deu41] provides the most versatile way of transfering problems from supersingular elliptic curves to quaternions. Let $\mathrm{SS}(p)$ denote the category with

- objects: supersingular elliptic curves over $\overline{\mathbf{F}}_p$;
- morphisms: algebraic group morphisms.

We fix a base curve $E_0 \in \mathrm{SS}(p)$. Let $\mathcal{O} = \mathrm{End}(E_0)$ and $B = \mathcal{O} \otimes \mathbf{Q}$.
Let $\mathrm{Mod}(\mathcal{O})$ denote the category with

- objects: invertible right $\mathcal{O}$-modules;
- morphisms: right $\mathcal{O}$-module homomorphisms.

Then we have the classical Deuring correspondence ([Voi21, Theorem 42.3.2], except we are using the covariant version; see also [Koh96, Theorem 45]).

Theorem 2.10. *The association* $E \mapsto \mathrm{Hom}(E_0, E)$, $(\varphi\colon E \to E') \mapsto (\psi \mapsto \varphi\psi)$ *defines a equivalence of categories*

$$\mathrm{SS}(p) \longrightarrow \mathrm{Mod}(\mathcal{O}).$$

2.8 Quaternionic automorphic forms

The following preliminaries concern the proof of our result on the equidistribution of isogeny random walks. The reader willing to admit Theorem 3.10 without proof does not need background on automorphic forms. A general reference for this section is [JL70]; a more gentle one, for Borel-type level, is [DV13, Section 3].

A *Hilbert space* V is a complex vector space equipped with a Hermitian inner product $\langle \cdot, \cdot \rangle$, and complete for the induced norm $\| \cdot \|$, which is automatic if V has finite dimension. Let V be a finite-dimensional Hilbert space. The *adjoint* of a linear operator $T\colon V \to V$ is the unique operator T^* satisfying $\langle Tv, w \rangle = \langle v, T^*w \rangle$ for all $v, w \in V$. A *normal operator* is an operator that commutes with its adjoint. The *operator norm* of an operator T is

$$\max_{v \neq 0} \frac{\|Tv\|}{\|v\|}.$$

Every normal operator stabilises the orthogonal complement of every stable subspace, is diagonalisable in an orthogonal basis, and has operator norm equal to the maximum absolute value of its eigenvalues.

Let $\widehat{\mathbf{Z}} = \prod_\ell \mathbf{Z}_\ell$ be the profinite completion of $\mathbf{Z}$. For every abelian group A, we write $\widehat{A} = A \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$. In particular $\widehat{\mathbf{Q}} = \prod'_\ell \mathbf{Q}_\ell$ is the ring of finite adèles of $\mathbf{Q}$.

Let $N \geq 1$ be an integer, and let $U(N) = (1+N\widehat{\mathbf{Z}}) \cap \widehat{\mathbf{Z}}^\times$. Let $H \subset (\mathbf{Z}/N\,\mathbf{Z})^\times$ be a subgroup, and let $U \subset \widehat{\mathbf{Z}}^\times$ be the preimage of H under the quotient map $\widehat{\mathbf{Z}}^\times \to \widehat{\mathbf{Z}}^\times/U(N) = (\mathbf{Z}/N\,\mathbf{Z})^\times$. Then we have an isomorphism

$$\mathbf{Q}_{>0}^\times \backslash \widehat{\mathbf{Q}}^\times / U \cong \widehat{\mathbf{Z}}^\times / U \cong (\mathbf{Z}/N\,\mathbf{Z})^\times / H.$$

Indeed, since every ideal of $\mathbf{Z}$ has a positive generator we have $\widehat{\mathbf{Q}}^\times = \mathbf{Q}_{>0}^\times \widehat{\mathbf{Z}}^\times$.

Let B be a quaternion algebra over $\mathbf{Q}$. The group $\widehat{B}^\times$ admits a measure μ that is bi-invariant under group translations, finite on compact subsets and nontrivial on open subsets, called its *Haar measure*.

Let $N \geq 1$ be an integer, let $U(N) = (1 + N\widehat{\mathcal{O}}) \cap \widehat{\mathcal{O}}^\times$, which is a compact open subgroup of $\widehat{B}^\times$, and let U be a subgroup satisfying $U(N) \subseteq U \subseteq \widehat{\mathcal{O}}^\times$. The set $B^\times \backslash \widehat{B}^\times / U$ is finite. The *space of automorphic forms of level* U is the Hilbert space $L^2(B^\times \backslash \widehat{B}^\times / U)$, equipped with the inner product induced by the projection of the Haar measure:

$$\langle F, G \rangle = \int_{b \in B^\times \backslash \widehat{B}^\times / U} F(b)\overline{G(b)}\mu(bU).$$

Fix, for every ℓ that is unramified in B, an isomorphism $\mathcal{O}_\ell \cong M_2(\mathbf{Z}_\ell)$, and let $\delta_\ell \in \widehat{B}^\times$ have component 1 at every $\ell' \neq \ell$ and that corresponds to $\begin{pmatrix} \ell & 0 \\ 0 & 1 \end{pmatrix}$ at ℓ via the chosen isomorphism. The Hecke operator

$$T_\ell \colon L^2(B^\times \backslash \widehat{B}^\times / U) \longrightarrow L^2(B^\times \backslash \widehat{B}^\times / U)$$

is defined by

$$T_\ell F(B^\times x U) = \sum_{uU \in U\delta_\ell U / U} F(B^\times x u U),$$

and its adjoint admits the expression

$$T_\ell^* F(B^\times x U) = \sum_{uU \in U\delta_\ell^{-1} U / U} F(B^\times x u U).$$

For ℓ that do not divide N, the Hecke operators T_ℓ are normal operators that pairwise commute.

3 Equidistribution of elliptic curves with extra data

The goal of this section is to prove Theorem 3.10. We state our results in Subsection 3.1. In Subsection 3.2, we set up a suitable version of the Deuring correspondence. The goal of Subsection 3.3 is to prove a technical result classifying extra data satisfying a simple property. In Subsection 3.4 we apply automorphic methods to prove the equidistribution theorem.

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*.*

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 $\mathcal{F}\colon \mathcal{C} \to \mathrm{Sets}$ is a functor, $\mathrm{El}(\mathcal{F})$ is the category of "objects $c \in \mathcal{C}$ with extra structure taken from $\mathcal{F}(c)$". A related definition can be found in [LM23, Definition 2.1], formulated at the level of graphs. For us, an advantage of the category-theoretic formulation is that we can forget about $\mathrm{SS}(p)$ and work in a quaternionic category, thanks to the Deuring correspondence.

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

- $\mathrm{Cyc}_N(E)$ 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 is the category of supersingular elliptic curves equipped with an endomorphism modulo N, which will play an important role in Section 4.

We now introduce the graphs of interest (more generally see Definition 3.28).

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 $(\mathrm{mod}\ N)$-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 $(\mathrm{mod}\ N)$-congruence property: indeed, endomorphisms divisible by N act as 0 on N-torsion points.

Example 3.9. The functor $\mathrm{End}\,/N$ from Example 3.3 has the $(\mathrm{mod}\ N)$-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 $(\mathrm{mod}\ N)$-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 is a normal operator on $L^2(\mathcal{G}_\mathcal{F})$ which 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 $1 \bmod N$. The operator norm of A_ℓ on $L^2_{\deg}(\mathcal{G}_\mathcal{F})$ is $\ell + 1$.*
- *$L^2_0(\mathcal{G}_\mathcal{F})$, the orthogonal complement of $L^2_{\deg}(\mathcal{G}_\mathcal{F})$. The operator norm of A_ℓ on $L^2_0(\mathcal{G}_\mathcal{F})$ is at most $2\sqrt{\ell}$.*

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 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 $\frac{1}{\#\operatorname{Aut}(E,x)}$. 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 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 $1 \bmod N$;
(3) the morphism $\mathrm{End}_\Sigma(E_0) \to (\mathrm{End}(E_0)/N\,\mathrm{End}(E_0))^\times$ is surjective, inducing an action of the group $G = (\mathrm{End}(E_0)/N\,\mathrm{End}(E_0))^\times$ on $\mathcal{F}(E_0)$.

Let $x_1, \ldots, x_n$ 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.

- Properties (1) and (2) allow us to transfer what happens at E_0 to any other curve.
- Property (3) allows us to define the H_i. When $p \nmid N$, this can be used to relate $\mathcal{F}$ to the setup of [CL23], using an isomorphism $G \cong \mathrm{GL}_2(\mathbf{Z}/N\,\mathbf{Z})$. Note that when $p \mid N$, the group $(\mathcal{O}/N\mathcal{O})^\times$ 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 $\mathcal{G}_{\mathcal{F}}$.

- Property (5) ensures that the space of functions on $\mathcal{G}_{\deg}$ injects into $L^2(\mathcal{G}_{\mathcal{F}})$ via the map Deg.
- 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 isomorphim $G \cong \mathrm{GL}_2(\mathbf{Z}/N\,\mathbf{Z})$ and a compatible bijection $\mathcal{F}(E_0) \cong \{\mathbf{Z}/N\,\mathbf{Z}\text{-lines in } (\mathbf{Z}/N\,\mathbf{Z})^2\}$. 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 Adélic Deuring correspondence

Since automorphic forms are usually defined using adélic language, we will reformulate the Deuring correspondence using adèles (Corollary 3.21).

The following terminology will be convenient.

Definition 3.14. *A* degree map *on a category $\mathcal{C}$ is the data, for every morphism f of $\mathcal{C}$, of an integer $\deg(f) \in \mathbf{Z}_{\geq 0}$ such that $\deg(fg) = \deg(f)\deg(g)$ for every morphisms f, g that can be composed, and such that $\deg(\mathrm{id}_x) = 1$ for all $x \in \mathcal{C}$.*

A functor $\mathcal{F}$ between categories equipped with degree maps is degree-preserving *if $\deg(\mathcal{F}(f)) = \deg(f)$ for every morphism f.*

When $\mathcal{C}$ is a category with a degree map and $\mathcal{F}\colon \mathcal{C} \to \mathrm{Sets}$ is a functor, we equip the category of elements $\mathrm{El}(\mathcal{F})$ with its inherited degree map.

Remark 3.15. In a category with a degree map, every isomorphism has degree 1.

Example 3.16. The category $\mathrm{SS}(p)$ is equipped with a degree map: the usual degree for isogenies, and 0 for the zero morphism.

The following is the basic object underlying the adélic Deuring correspondence.

Definition 3.17. *We define the category* $\mathrm{Cosets}(\widehat{\mathcal{O}}^\times)$ *with*

- *objects: cosets $[x] := x\widehat{\mathcal{O}}^\times \in \widehat{B}^\times/\widehat{\mathcal{O}}^\times$ for $x \in \widehat{B}^\times$;*
- *morphisms:* $\mathrm{Hom}([x],[y]) = B \cap y\widehat{\mathcal{O}}x^{-1}$*, using multiplication in B as composition.*

We equip the category $\mathrm{Cosets}(\widehat{\mathcal{O}}^\times)$ *with a degree map as follows: for every* $b \in \mathrm{Hom}([x],[y])$*, we define the* degree of b *to be the positive integer* $\deg(b)$ *such that* $\deg(b)\widehat{\mathbf{Z}} = \mathrm{nrd}(u)\widehat{\mathbf{Z}}$ *where* $b = yux^{-1}$*.*

Remark 3.18. We warn the reader that a single element $b \in B$ can represent different morphisms, depending on the source $[x]$ and target $[y]$. Moreover, the degree of the morphism is in general not the reduced norm of b.

We reformulate the Deuring correspondence adélically as follows.

Proposition 3.19. *The association* $[x] \mapsto B \cap x\widehat{\mathcal{O}}$, $g \in \mathrm{Hom}([x],[y]) \mapsto (b \mapsto gb)$ *defines a equivalence of categories*

$$\mathrm{Cosets}(\widehat{\mathcal{O}}^\times) \longrightarrow \mathrm{Mod}(\mathcal{O}).$$

Its composition with the equivalence of Theorem 2.10 is a degree-preserving equivalence of categories

$$\mathrm{Cosets}(\widehat{\mathcal{O}}^\times) \longrightarrow \mathrm{SS}(p).$$

Proof. The association described clearly defines a faithful functor.

We claim that the functor is full. Indeed, let $f \in \mathrm{Hom}(B \cap x\widehat{\mathcal{O}}, B \cap y\widehat{\mathcal{O}})$. Since every right B-module endomorphism of B is a left multiplication by an element of B, there exists $g \in B$ such that $f(b) = gb$ for all $b \in B \cap x\widehat{\mathcal{O}}$. Moreover, by weak approximation the closure of $B \cap x\widehat{\mathcal{O}}$ in $\widehat{B}$ is $x\widehat{\mathcal{O}}$, so we must have $gx\widehat{\mathcal{O}} \subset y\widehat{\mathcal{O}}$ and therefore $g \in y\widehat{\mathcal{O}}x^{-1}$, so that $g \in \mathrm{Hom}([x],[y])$ as claimed.

Finally, the functor is essentially surjective since every right invertible $\mathcal{O}$-module is isomorphic to a right invertible $\mathcal{O}$-ideal I, and such an ideal is locally principal and therefore of the form $I = B \cap x\widehat{\mathcal{O}}$.

By examining the determinant of a morphism on the modules, we see that the equivalence preserves the degree.

□

Fix Σ be a set of primes. Let $\widehat{\mathcal{O}}_\Sigma$ denote the ring obtained from $\widehat{\mathcal{O}}$ by inverting all primes in Σ. Then $\widehat{\mathcal{O}} \cap \widehat{\mathcal{O}}_\Sigma^\times$ is the set of elements $u \in \widehat{\mathcal{O}}$ such that $\mathrm{nrd}(u)\widehat{\mathbf{Z}}$ is generated by a product of the primes in Σ.

Definition 3.20. *Let* $\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times)$ *be the category with*

- *objects: cosets* $[x] = x\widehat{\mathcal{O}}^\times \in \widehat{B}^\times/\widehat{\mathcal{O}}^\times$ *for* $x \in \widehat{B}^\times$*;*
- *morphisms:* $\mathrm{Hom}_\Sigma([x],[y]) = B^\times \cap y(\widehat{\mathcal{O}} \cap \widehat{\mathcal{O}}_\Sigma^\times)x^{-1} = \mathrm{Hom}([x],[y]) \cap y\widehat{\mathcal{O}}_\Sigma^\times x^{-1}$*, using multiplication in* B *as composition.*

Equivalently, the morphisms in $\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times)$ are the morphisms in $\mathrm{Cosets}(\widehat{\mathcal{O}}^\times)$ whose degree is a product of the primes in Σ. We obtain the following corollary.

Corollary 3.21 (Adélic Deuring correspondence). *The second equivalence from Proposition 3.19 induces a degree-preserving equivalence of categories*

$$\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times) \longrightarrow \mathrm{SS}_\Sigma(p).$$

3.3 Extra data of congruence type

Fix $N \geq 1$ an integer not divisible by any prime in Σ.

In this subsection, we will study categories of elements of various functors, as in Definition 2.8. We will use functors coming from quaternionic constructions, which will allow us to apply automorphic methods. The main result of this subsection is Theorem 3.27, which classifies functors satisfying a simple property in terms of adélic groups.

Let $U(N) = (1 + N\widehat{\mathcal{O}}) \cap \widehat{\mathcal{O}}^\times$, which is a finite index subgroup of $\widehat{\mathcal{O}}^\times$, and similarly $U_\Sigma(N) = (1+N\widehat{\mathcal{O}}_\Sigma)\cap\widehat{\mathcal{O}}_\Sigma^\times$. Let U be a subgroup of $\widehat{\mathcal{O}}^\times$ containing $U(N)$, and let $U_\Sigma = U \cdot U_\Sigma(N)$, so that $U = U_\Sigma \cap \widehat{\mathcal{O}}^\times$. Note that the natural map $\widehat{\mathcal{O}}^\times/U \to \widehat{\mathcal{O}}_\Sigma^\times/U_\Sigma$ is a bijection, i.e. we have $\widehat{\mathcal{O}}_\Sigma^\times = \widehat{\mathcal{O}}^\times U_\Sigma$.

It is helpful to think about these definition in terms of the product decomposition $\widehat{B}^\times = \prod'_\ell B_\ell^\times$ as follows: we have

$$U = U' \times \prod_{\ell\nmid N} \mathcal{O}_\ell^\times \text{ and } U_\Sigma = U' \times \prod_{\ell\nmid N, \ell\notin\Sigma} \mathcal{O}_\ell^\times \times \prod_{\ell\in\Sigma}{}' B_\ell^\times$$

where U' is the image of U in $\prod_{\ell|N} \mathcal{O}_\ell^\times$.

Definition 3.22. *Let $\mathcal{F}_U$ be the functor*

$$\mathcal{F}_U\colon \operatorname{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times) \to \operatorname{Sets}$$

defined by

- $\mathcal{F}_U([x]) = x\widehat{\mathcal{O}}_\Sigma^\times/U_\Sigma$,
- *for $b \in \operatorname{Hom}_\Sigma([x],[y])$, the map $\mathcal{F}_U(b)\colon \mathcal{F}_U([x]) \to \mathcal{F}_U([y])$ is left multiplication by b.*

Let $\operatorname{Cosets}_\Sigma(U)$ be the category with

- *objects: cosets $xU \in \widehat{B}^\times/U$ for $x \in \widehat{B}^\times$;*
- *morphisms: $\operatorname{Hom}_U(xU, yU) = B^\times \cap y(\widehat{\mathcal{O}}\cap U_\Sigma)x^{-1} = \operatorname{Hom}([x],[y])\cap yU_\Sigma x^{-1}$, using multiplication in B as composition.*

In other words, morphism are required to respect U at the primes dividing N, and to have degree a product of the primes in Σ.

Example 3.23. Assume that p does not divide N. Let $\mathcal{O}_0(N) \subset \mathcal{O}$ be an Eichler order of level N, and let $U = \widehat{\mathcal{O}}_0(N)^\times$. Then for all $x \in \widehat{B}^\times$, the set $\mathcal{F}_U([x])$ is in bijection with $\mathbf{P}^1(\mathbf{Z}/N\,\mathbf{Z})$, or more naturally with the set of Eichler orders contained in the maximal order $x\widehat{\mathcal{O}}x^{-1} \cap B$.

The following proposition is the bridge between functors on the quaternionic side of the Deuring correspondence and automorphic forms.

Proposition 3.24. *The association*

- $([x], xgU_\Sigma) \mapsto xg'U$ *where* $xgU_\Sigma = xg'U_\Sigma$ *and* $g' \in \widehat{\mathcal{O}}^\times$;
- $b \in \mathrm{Hom}_{\mathrm{El}(\mathcal{F}_U)}(([x], xgU_\Sigma), ([y], yhU_\Sigma)) \mapsto b \in \mathrm{Hom}_U(xg'U, yh'U)$

defines a degree-preserving equivalence of categories

$$\mathrm{El}(\mathcal{F}_U) \longrightarrow \mathrm{Cosets}_\Sigma(U).$$

Proof. First, the association is well-defined on objects: if two elements $g', g'' \in \widehat{\mathcal{O}}^\times$ satisfy $xg'U_\Sigma = xg''U_\Sigma$, then $(g')^{-1}g'' \in U_\Sigma \cap \widehat{\mathcal{O}}^\times = U$ so $xg'U = xg''U$. Next, it is well-defined on morphisms: let $b \in \mathrm{Hom}_{\mathrm{El}(\mathcal{F}_U)}(([x], xgU_\Sigma), ([y], yhU_\Sigma))$, and $xgU_\Sigma = xg'U_\Sigma$ and $yhU_\Sigma = yh'U_\Sigma$ with $g', h' \in \widehat{\mathcal{O}}^\times$; then $bxg'U_\Sigma = yh'U_\Sigma$ so $b \in yh'U_\Sigma(xg')^{-1}$ and therefore $b \in \mathrm{Hom}_U(xg'U, yh'U)$. Since the association is clearly multiplicative on morphisms, it defines a functor. Moreover, the functor $\mathrm{Cosets}_\Sigma(U) \longrightarrow \mathrm{El}(\mathcal{F}_U)$ defined by

- $xU \mapsto ([x], xU_\Sigma)$;
- $b \in \mathrm{Hom}_U(xU, yU) \mapsto b$

is clearly an inverse, so we obtain an equivalence as claimed. □

We will need the following consequence of the strong approximation theorem.

Lemma 3.25. *Assume that Σ contains at least one prime different from p and that it generates $(\mathbf{Z}/N\mathbf{Z})^\times$. Then for every $g, x \in \widehat{B}^\times$, we have*

$$\mathrm{Hom}_\Sigma([x], [gx]) \cap gxU_\Sigma(N)x^{-1} \neq \emptyset.$$

Proof. Let $x \in \widehat{B}^\times$, and let $H = xU_\Sigma(N)x^{-1}$. Since Σ contains a prime different from p, strong approximation holds ([Voi21, Theorem 28.5.3, see also 28.5.5]), so the reduced norm induces a bijection

$$B^\times \backslash \widehat{B}^\times / H \longrightarrow \mathbf{Q}_{>0}^\times \backslash \widehat{\mathbf{Q}}^\times / \mathrm{nrd}(H).$$

On the other hand, the group $\mathbf{Q}_{>0}^\times \backslash \widehat{\mathbf{Q}}^\times / \mathrm{nrd}(H)$ is isomorphic to a quotient of $(\mathbf{Z}/N\mathbf{Z})^\times / \langle \Sigma \rangle$. Since Σ generates $(\mathbf{Z}/N\mathbf{Z})^\times$ the latter quotient is trivial, so $\widehat{B}^\times = B^\times H$. Now let $g \in \widehat{B}^\times$. Write $g = b_0 x u_0^{-1} x$ with $b_0 \in B^\times$ and $u_0 \in U_\Sigma(N)$, so that $b_0 = gxu_0x^{-1}$. Let λ be a product of the primes in Σ such that $\lambda u_0 \in \widehat{\mathcal{O}}$ and $\lambda \equiv 1 \bmod N$. Let $b = \lambda b_0$ and $u = \lambda u_0$, which satisfy $gxux^{-1} = b \in B^\times$ and $u \in \widehat{\mathcal{O}} \cap U_\Sigma(N)$. Then $b \in \mathrm{Hom}_\Sigma([x], [gx]) \cap gxU_\Sigma(N)x^{-1}$, which is therefore not empty. □

We are now in position to state and prove the main result of this subsection.

Definition 3.26. *Let $\mathcal{F}\colon \mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times) \to \mathrm{Sets}$ be a functor. We say that $\mathcal{F}$ is* of N-congruence type *if $\mathcal{F}$ is isomorphic to a disjoint union of functors $\mathcal{F}_U$. We say that $\mathcal{F}$* satisfies the (mod N)-congruence property *if for every $x \in \widehat{B}^\times$ and for every $a, b \in \mathrm{End}_\Sigma([x])$ such that $a - b \in N \cdot \mathrm{End}([x])$, we have $\mathcal{F}(a) = \mathcal{F}(b)$.*

Theorem 3.27. *Assume that Σ contains at least one prime different from p and that it generates $(\mathbf{Z}/N\mathbf{Z})^\times$. Let $\mathcal{F}\colon \mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times) \to \mathrm{Sets}$ be a functor. Then $\mathcal{F}$ is of N-congruence type if and only if $\mathcal{F}$ satisfies the $(\mathrm{mod}\ N)$-congruence property. More precisely, assume that $\mathcal{F}$ satisfies the $(\mathrm{mod}\ N)$-congruence property. Then:*

(1) for every morphism f in $\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times)$, the map $\mathcal{F}(f)$ is a bijection; and
(2) the morphism $\mathrm{End}_\Sigma([1]) \to (\mathcal{O}/N\mathcal{O})^\times$ is surjective, inducing an action of the group $G = (\mathcal{O}/N\mathcal{O})^\times$ on $\mathcal{F}([1])$.

Choose a G-equivariant bijection $\mathcal{F}([1]) \cong \bigsqcup_i G/H_i$, and for all i, let U_i be the preimage of H_i under the quotient map $\widehat{\mathcal{O}}^\times \to (\mathcal{O}/N\mathcal{O})^\times$. Then:

(3) there exists an isomorphism of functors

$$\mathcal{F} \cong \bigsqcup_i \mathcal{F}_{U_i}.$$

Proof. First, every functor of N-congruence type clearly satisfies the $(\mathrm{mod}N)$-congruence property.

Assume that $\mathcal{F}$ satisfies the $(\mathrm{mod}\ N)$-congruence property.

Step 1. We claim that all $\mathcal{F}(f)$ are bijections. Indeed, let $f \in \mathrm{End}_\Sigma([x])$ be an endomorphism in $\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times)$. Then its reduction modulo N has finite order, say k, so that $f^k - 1 \in Nx\widehat{\mathcal{O}}x^{-1}$. By the $(\mathrm{mod}\ N)$-congruence property this implies $\mathcal{F}(f)^k = \mathcal{F}(1) = \mathrm{id}$, so that $\mathcal{F}(f)$ is invertible. Now let $f\colon [x] \to [y]$ be an arbitrary morphism in $\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times)$. Then there exists a morphism $g\colon [y] \to [x]$ such that $fg = \deg(f) \in \mathrm{End}_\Sigma([y])$ and $gf = \deg(f) \in \mathrm{End}_\Sigma([x])$. By the endomorphism case, this proves that $\mathcal{F}(fg)$ and $\mathcal{F}(gf)$ are invertible, hence that $\mathcal{F}(f)$ is. This proves (1).

Step 2. We claim that for all $x, y \in \widehat{B}^\times$ and all $a, b \in \mathrm{Hom}_\Sigma([x],[y])$, if $a - b \in Ny\widehat{\mathcal{O}}x^{-1}$ then $\mathcal{F}(a) = \mathcal{F}(b)$. Indeed under these conditions there exists $c \in \mathrm{Hom}_\Sigma([y],[x])$. Then $ca - cb \in \mathrm{End}_\Sigma([x]) \cap Nx\widehat{\mathcal{O}}x^{-1}$. By the $(\mathrm{mod}\ N)$-congruence property we have $\mathcal{F}(ca) = \mathcal{F}(cb)$. Since $\mathcal{F}(c)$ is invertible, this proves $\mathcal{F}(a) = \mathcal{F}(b)$.

We also claim that for all $x, y \in \widehat{B}^\times$ and all $a, b \in \mathrm{Hom}_\Sigma([x],[y])$, if $b \in axU_\Sigma(N)x^{-1}$ then $\mathcal{F}(a) = \mathcal{F}(b)$. Indeed the condition implies that

$$b \in a(1 + Nx\widehat{\mathcal{O}}_\Sigma x^{-1}) = a + Nax\widehat{\mathcal{O}}_\Sigma x^{-1} = a + Ny\widehat{\mathcal{O}}_\Sigma x^{-1}.$$

Since $\widehat{\mathcal{O}} \cap N\widehat{\mathcal{O}}_\Sigma = N\widehat{\mathcal{O}}$, the previous claim applies, and therefore $\mathcal{F}(a) = \mathcal{F}(b)$.

Step 3. Inspired by Proposition 3.24, we are going to define an action of $\widehat{B}^\times$ on $\bigsqcup_{[x]} \mathcal{F}([x])$. Let $g, x \in \widehat{B}^\times$. By Lemma 3.25, there exists $b \in \mathrm{Hom}_\Sigma([x],[gx]) \cap gxU_\Sigma(N)x^{-1}$. For $A \in \mathcal{F}([x])$, we define $g \cdot A = \mathcal{F}(b)(A) \in \mathcal{F}([gx])$. To see that this is well-defined, let $b' \in \mathrm{Hom}_\Sigma([x],[gx]) \cap gxU_\Sigma(N)x^{-1}$ be another element. We have $b' \in bxU_\Sigma(N)x^{-1}$, so that $\mathcal{F}(b) = \mathcal{F}(b')$.

The defined action is multiplicative, because when $d \in \mathrm{Hom}_\Sigma([x],[hx]) \cap hxU_\Sigma(N)x^{-1}$ and $c \in \mathrm{Hom}_\Sigma([hx],[ghx]) \cap ghxU_\Sigma(N)(hx)^{-1}$ we have $cd \in \mathrm{Hom}_\Sigma([x],[ghx]) \cap ghxU_\Sigma(N)x^{-1}$: the action of gh is given by $\mathcal{F}(cd) = \mathcal{F}(c)\mathcal{F}(d)$.

We therefore get an action of $\widehat{B}^\times$ on $\bigsqcup_{[x]} \mathcal{F}([x])$ with the following properties for $g, x \in \widehat{B}^\times$:

- the action of g induces a bijection $\mathcal{F}([x]) \to \mathcal{F}([gx])$;
- $x\widehat{\mathcal{O}}^\times x^{-1}$ stabilises $\mathcal{F}([x])$;
- $xU(N)x^{-1}$ acts trivially on $\mathcal{F}([x])$.

In particular, we obtain an action of $\widehat{\mathcal{O}}^\times$ on $\mathcal{F}([1])$. By decomposing this action into orbits, we obtain an $\widehat{\mathcal{O}}^\times$-equivariant bijection

$$\psi_{[1]} \colon \bigsqcup_{i \in I} \widehat{\mathcal{O}}^\times / U_i \longrightarrow \mathcal{F}([1]),$$

where the U_i are subgroups of $\widehat{\mathcal{O}}^\times$ containing $U(N)$. Recalling the definition of the functors $\mathcal{F}_{U_i}$ (Definition 3.22) we see that this is the same as an $\widehat{\mathcal{O}}^\times$-equivariant bijection

$$\psi_{[1]} \colon \bigsqcup_{i \in I} \mathcal{F}_{U_i}([1]) \longrightarrow \mathcal{F}([1]).$$

In fact, the action of $\widehat{\mathcal{O}}^\times$ factors through $\widehat{\mathcal{O}}^\times \to \widehat{\mathcal{O}}^\times / U(N) \cong (\mathcal{O}/N\mathcal{O})^\times$, and comes from the application of Lemma 3.25 to $x = 1$ and $g \in \widehat{\mathcal{O}}^\times$, from which we see that (2) holds and one can choose the U_i compatibly with the H_i from the statement of the theorem.

Step 4. We extend $\psi_{[1]}$ to an isomorphism of functors $\psi \colon \bigsqcup_{i \in I} \mathcal{F}_{U_i} \to \mathcal{F}$. Let $x \in \widehat{B}^\times$. We define

$$\psi_{[x]} \colon \bigsqcup_{i \in I} \mathcal{F}_{U_i}([x]) \longrightarrow \mathcal{F}([x])$$

by setting for every $U = U_i$ and every $xgU_\Sigma \in \mathcal{F}_U([x])$ with $g \in \widehat{\mathcal{O}}^\times_\Sigma$,

$$\psi_{[x]}(xgU_\Sigma) = x \cdot \psi_{[1]}(gU_\Sigma).$$

The map $\psi_{[x]}$ is well-defined since $xgU_\Sigma = xg'U_\Sigma$ implies $gU_\Sigma = g'U_\Sigma$. In addition, $\psi_{[x]}$ depends only on $[x]$: for all $u \in \widehat{\mathcal{O}}^\times$ we have $(xu) \cdot \psi_{[1]}(u^{-1}gU_\Sigma) = x \cdot \psi_{[1]}(gU_\Sigma)$ by $\widehat{\mathcal{O}}^\times$-equivariance. Since the multiplication by x^{-1} from $\mathcal{F}_U([x])$ to $\mathcal{F}_U([1])$, the map $\psi_{[1]}$ and the action of x from $\mathcal{F}([1])$ to $\mathcal{F}([x])$ are all bijections, the map $\psi_{[x]}$ is a bijection. We now prove that $\psi = (\psi_{[x]})_{[x]}$ is a morphism of functors.

The proof will follow the following diagram:

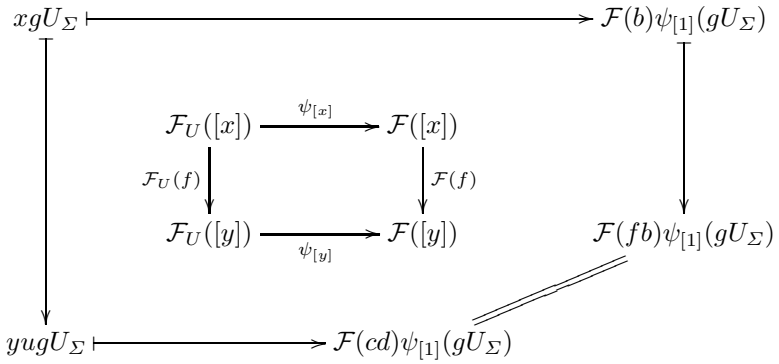

Let $x, y \in \widehat{B}^\times$, $f \in \mathrm{Hom}_\Sigma([x],[y])$ and U be one of the U_i; we will prove that $\mathcal{F}(f) \circ \psi_{[x]} = \psi_{[y]} \circ \mathcal{F}_U(f)$ holds on $\mathcal{F}_U([x])$. Let $xgU_\Sigma \in \mathcal{F}_U([x])$. We have

$$\mathcal{F}(f) \circ \psi_{[x]}(xgU_\Sigma) = \mathcal{F}(f)\big(x \cdot \psi_{[1]}(gU_\Sigma)\big) = \mathcal{F}(fb)\psi_{[1]}(gU_\Sigma),$$

where $b \in \mathrm{Hom}_\Sigma([1],[x]) \cap xU_\Sigma(N)$. Write $f = yux^{-1}$ with $u \in \widehat{\mathcal{O}} \cap \mathcal{O}_\Sigma^\times$, and note that $fxgU_\Sigma = yugU_\Sigma \in \mathcal{F}_U([y])$. We therefore have

$$\psi_{[y]} \circ \mathcal{F}_U(f)(xgU_\Sigma) = \psi_{[y]}(fxgU_\Sigma) = y \cdot \psi_{[1]}(ugU_\Sigma) = \mathcal{F}(c)\psi_{[1]}(ugU_\Sigma)$$

where $c \in \mathrm{Hom}_\Sigma([1],[y]) \cap yU_\Sigma(N)$. Now $u \in vU_\Sigma(N)$ for some $v \in \widehat{\mathcal{O}}^\times$, so that $ugU_\Sigma = vgU_\Sigma$ since $\widehat{\mathcal{O}}_\Sigma^\times$ normalises $U_\Sigma(N)$, and by equivariance of $\psi_{[1]}$ we have

$$\psi_{[1]}(ugU_\Sigma) = \psi_{[1]}(vgU_\Sigma) = v \cdot \psi_{[1]}(gU_\Sigma) = \mathcal{F}(d)\psi_{[1]}(gU_\Sigma)$$

where $d \in \mathrm{End}_\Sigma([1]) \cap vU_\Sigma(N)$. We get

$$\psi_{[y]} \circ \mathcal{F}_U(f)(xgU_\Sigma) = \mathcal{F}(cd)\psi_{[1]}(gU_\Sigma).$$

We finally compare fb and cd. We have

$$fb \in (yux^{-1})xU_\Sigma(N) = yuU_\Sigma(N),$$

and

$$cd \in yU_\Sigma(N)vU_\Sigma(N) = yvU_\Sigma(N) = yuU_\Sigma(N).$$

Since fb and cd both belong to $\mathrm{Hom}_\Sigma([1],[y]) = \mathrm{Hom}_\Sigma([1],[yu])$, this proves that $\mathcal{F}(fb) = \mathcal{F}(cd)$. This proves that ψ is an isomorphism of functors, so that $\mathcal{F}$ is of N-congruence type, proving (3) and concluding the proof. □

3.4 Associated graphs and equidistribution

In this subsection, we study the graphs of interest and prove our main equidistribution theorem: Theorem 3.10 and its companion Proposition 3.11.

We first introduce a categorical construction of graphs generalising Definition 3.4.

Definition 3.28. *Let $\mathcal{C}$ be category with finitely many isomorphism classes of objects, finite automorphism groups, and equipped with a degree map. We define the graph* $\mathrm{Graph}(\mathcal{C})$ *with:*

- *vertices: isomorphism classes of objects in $\mathcal{C}$;*
- *edges: let $x \in \mathcal{C}$; the set of edges from the vertex corresponding to x is the set of classes of morphisms from x modulo the relation $(f\colon x \to y) \sim (g\colon x \to z)$ if and only there exists $u \in \mathrm{Isom}_{\mathcal{C}}(y,z)$ such that $g = uf$; the endpoint of the edge corresponding to $f\colon x \to y$ is the isomorphism class of y. In other words, the set of edges between the classes of $x, y \in \mathcal{C}$ is $\mathrm{Aut}(y)\backslash \mathrm{Hom}(x,y)$.*

The degree *of an edge is the degree of the corresponding morphism.*

We define a measure on the set of vertices of $\mathrm{Graph}(\mathcal{C})$ *by giving each vertex v measure $\frac{1}{\#\mathrm{Aut}(x)}$ where v corresponds to $x \in \mathcal{C}$, and we write $L^2(\mathrm{Graph}(\mathcal{C}))$ the Hilbert space of complex functions on the set of vertices of* $\mathrm{Graph}(\mathcal{C})$.

For every prime ℓ, we define an adjacency operator A_ℓ on $L^2(\mathrm{Graph}(\mathcal{C}))$ given by

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

where the sum runs over edges of degree ℓ leaving x.

Remark 3.29. For every functor $\mathcal{F}$ as in Definition 3.4, we have $\mathcal{G}_{\mathcal{F}} = \mathrm{Graph}(\mathrm{El}(\mathcal{F}))$. Every degree-preserving equivalence of categories $\mathcal{C} \cong \mathcal{D}$ induces an isomorphism of graphs $\mathrm{Graph}(\mathcal{C}) \cong \mathrm{Graph}(\mathcal{D})$ compatible with all the structure from Definition 3.28.

The following lemma relates the graphs obtained from our quaternionic categories to automorphic forms.

Lemma 3.30. *The category* $\mathrm{Cosets}_\Sigma(U)$ *and its associated graph have the following properties:*

(1) Two objects $x, y \in \widehat{B}^\times/U$ are isomorphic if and only if they have the same image in the quotient $B^\times\backslash\widehat{B}^\times/U$.

(2) The projection to $B^\times\backslash\widehat{B}^\times/U$ of a Haar measure on $\widehat{B}^\times$ coincides with the measure on the set of vertices of $\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U))$.

(3) For every $x \in \widehat{B}^\times$, the map

$$\mathrm{edg}\colon u \in \widehat{\mathcal{O}} \cap U_\Sigma \longmapsto 1 \in \mathrm{Hom}_U(xU, xu^{-1}U)$$

induces a bijection between $U\backslash(\widehat{\mathcal{O}} \cap U_\Sigma)$ and the set of edges leaving the vertex $B^\times xU$ in $\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U))$.

(4) For every prime $\ell \in \Sigma$ different from p, the adjacency operator A_ℓ coincides with the adjoint of the Hecke operator T_ℓ on $L^2(B^\times\backslash\widehat{B}^\times/U)$.

Proof.

(1) Isomorphisms in $\mathrm{Cosets}_\Sigma(U)$ are exactly morphisms of degree 1, so that the set of isomorphisms between two cosets xU, yU is $B^\times \cap yUx^{-1}$, i.e. the set of elements $b \in B^\times$ such that $bxU = yU$. This proves the claim.
(2) Since cosets of U are open and form a disjoint union, by translation invariance every element of $\widehat{B}^\times/U$ has the same nonzero measure. We normalise the Haar measure so that each coset of U has measure 1. For every $x \in \widehat{B}^\times$, every fiber of the projection map $xU \mapsto B^\times xU$ has cardinality $\#(B^\times \cap xUx^{-1}) = \#\operatorname{Aut}_U(xU)$, so the projected measure of $B^\times xU$ is the inverse of this cardinality, as claimed.
(3) Let $x \in \widehat{B}^\times$. For every $u \in \widehat{\mathcal{O}} \cap U_\Sigma$, we have $1 = (xu^{-1})ux \in xu^{-1}(\widehat{\mathcal{O}} \cap U_\Sigma)x \cap B^\times = \mathrm{Hom}_U(xU, xu^{-1}U)$, so the map edg is well-defined. Let $f \in \mathrm{Hom}_U(xU, yU)$ represent an edge $x \to y$ in $\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U))$. Then $1 \in \mathrm{Hom}_U(xU, f^{-1}yU)$ represents the same edge, which is therefore $\mathrm{edg}(y^{-1}fx)$. Moreover, two morphisms $f \in \mathrm{Hom}_U(xU, yU)$ and $g \in \mathrm{Hom}_U(xU, zU)$ represent the same edge if and only if there exists $b \in B^\times$ such that $g = bf$. For morphisms in the image of edg, this can only happen with $b = 1$, so the edge $\mathrm{edg}(u)$ is completely determined by its endpoint $xu^{-1}U$, i.e. by the coset Uu.
(4) Consider the cosets $Uu \in U\backslash(\widehat{\mathcal{O}} \cap U_\Sigma)$ such that $\mathrm{nrd}(u)\widehat{\mathbf{Z}} = \ell\widehat{\mathbf{Z}}$. Then for every $\ell' \neq \ell$, the ℓ'-component of u is in the ℓ'-component of U, so we may replace it by 1. Choosing an isomorphism $\mathcal{O} \otimes \mathbf{Z}_\ell \cong M_2(\mathbf{Z}_\ell)$, the possible cosets correspond to the cosets in $\mathrm{GL}_2(\mathbf{Z}_\ell)\backslash \mathrm{GL}_2(\mathbf{Q}_\ell)$ whose determinant has valuation 1: these are exactly the cosets of $\mathrm{GL}_2(\mathbf{Z}_\ell)$ that belong to the double coset

$$\mathrm{GL}_2(\mathbf{Z}_\ell) \begin{pmatrix} \ell & 0 \\ 0 & 1 \end{pmatrix} \mathrm{GL}_2(\mathbf{Z}_\ell).$$

Let $F \in L^2(\mathrm{Graph}(\mathrm{Cosets}_U(\Sigma)))$. From the above we have

$$A_\ell F(B^\times xU) = \sum_{Uu \in U\backslash U\delta_\ell U} F(B^\times xu^{-1}U).$$

This is the adjoint of the Hecke operator T_ℓ, as claimed.

□

The following proposition is the quaternionic version of our equidistribution result.

Proposition 3.31. *Let $L^2_{\mathrm{nrd}}(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U))) \subset L^2(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ denote the subspace of functions that factor through the reduced norm map*

$$B^\times\backslash\widehat{B}^\times/U \longrightarrow \mathbf{Q}^\times_{>0}\backslash\widehat{\mathbf{Q}}^\times/\operatorname{nrd}(U),$$

and let $L^2_0(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ denote the orthogonal complement of the subspace $L^2_{\mathrm{nrd}}(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$. Then, for every $\ell \in \Sigma$ different from p, the adjacency operator A_ℓ is a normal operator, stabilises $L^2_{\mathrm{nrd}}(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ and $L^2_0(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$, and its operator norm on $L^2_0(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ is at most $2\sqrt{\ell}$. Moreover, the A_ℓ for $\ell \in \Sigma$ pairwise commute.

Proof. It is clear that $L^2_{\mathrm{nrd}}(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ is stable under A_ℓ. The operators A_ℓ are normal and pairwise commute by Lemma 3.30 (4), and therefore leave $L^2_0(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ stable and are diagonalisable. We bound the operator norm of A_ℓ by bounding its eigenvalues, equivalently by bounding the eigenvalues of the Hecke operator T_ℓ. The space $L^2_{\mathrm{nrd}}(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ is exactly the subspace of $L^2(B^\times \backslash \widehat{B}^\times / U)$ of automorphic forms that generate a one-dimensional automorphic representation (i.e. of the form $g \mapsto \chi(\mathrm{nrd}\, g)$ for some Dirichlet character χ). Therefore, by the Jacquet–Langlands correspondence [JL70, Theorem 14.4], every system of Hecke eigenvalues appearing in $L^2_0(\mathrm{Graph}(\mathrm{Cosets}_\Sigma(U)))$ is also the one attached to a cuspidal modular newform of weight 2 ramified only at primes dividing pN. Therefore, by Deligne's theorem [Del73, Theorem 8.2], the absolute values of the eigenvalues of T_ℓ are bounded by $2\sqrt{\ell}$. This proves the proposition. □

Remark 3.32. Using the full statement of the Jacquet–Langlands correspondence, one could obtain the exact eigenvalues in terms of classical modular forms. This is not needed in our applications.

We can finally prove Theorem 3.10 and Proposition 3.11.

Proof (Theorem 3.10 and Proposition 3.11). First, we use Corollary 3.21 to transfer the entire situation to the quaternionic category $\mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times)$; in particular $\mathcal{F}$ induces a functor $\mathcal{F}' \colon \mathrm{Cosets}_\Sigma(\widehat{\mathcal{O}}^\times) \to \mathrm{Sets}$. Since the equivalence of Proposition 3.19 is additive, the functor $\mathcal{F}'$ satisfies the (mod N)-congruence property in the sense of Definition 3.26, so that we can apply Theorem 3.27. From (1) and (2) of Theorem 3.27, we obtain (1) and (3) respectively. Moreover, we can choose the H_i of Theorem 3.27 to coincide with those of Proposition 3.11.

Let $E \in \mathrm{SS}(p)$. It is standard that there exists $\psi \in \mathrm{Hom}_\Sigma(E_0, E)$. By (3), there exists $\alpha \in \mathrm{End}_\Sigma(E_0)$ whose degree is the inverse of $\deg(\psi) \bmod N$, so that $\varphi = \psi\alpha \in \mathrm{Hom}_\Sigma(E_0, E)$ has degree 1 mod N. This proves (2) when one of the curves is E_0, and therefore in general by going via E_0.

By (2), every vertex of $\mathcal{G}_\mathcal{F}$ is connected to one above E_0. Moreover, two vertices above E_0 are connected if and only if they are related by an element of $\mathrm{End}_\Sigma(E_0)$, if and only if they are in the same orbit under G. In particular there is exactly one vertex of the form (E_0, x_i) in each connected component of $\mathcal{G}_\mathcal{F}$. Since every vertex of $\mathcal{G}_{\deg}$ has exactly one outgoing edge labelled by each element of $(\mathbf{Z}/N\,\mathbf{Z})^\times$, this proves that there is at most one morphism of graphs satisfying the properties of (4).

We now prove the existence of Deg. Let U_i be as in Theorem 3.27. Applying (3) of that theorem and Proposition 3.24 we obtain a degree-preserving equivalence of categories

$$\mathrm{El}(\mathcal{F}) \cong \bigsqcup_i \mathrm{Cosets}_\Sigma(U_i),$$

inducing an isomorphism of graphs

$$\mathcal{G}_\mathcal{F} \cong \bigsqcup_i \mathrm{Graph}(\mathrm{Cosets}_\Sigma(U_i)).$$

By Lemma 3.30 (1) and (3), the reduced norm map

$$\mathrm{nrd}\colon B^\times \backslash \widehat{B}^\times / U_i \to \mathbf{Q}_{>0}^\times \backslash \widehat{\mathbf{Q}}^\times / \mathrm{nrd}(U_i)$$

combined with the isomorphism

$$\mathbf{Q}_{>0}^\times \backslash \widehat{\mathbf{Q}}^\times / \mathrm{nrd}(U_i) \cong (\mathbf{Z}/N\,\mathbf{Z})^\times / \mathrm{nrd}(H_i)$$

translates into a graph morphism

$$\mathrm{Deg}\colon \mathcal{G}_\mathcal{F} \longrightarrow \mathcal{G}_{\mathrm{deg}}$$

satisfying the properties of (4). Since the adélic reduced norm map $\widehat{B}^\times \to \widehat{\mathbf{Q}}^\times$ is surjective, so is Deg, proving (5).

Let $L^2_{\mathrm{Deg}}(\mathcal{G}_\mathcal{F}) \subset L^2(\mathcal{G}_\mathcal{F})$ be the subspace of functions that factor through Deg. Since vertices connected by an edge of degree 1 mod N clearly have the same image under Deg, we have $L^2_{\mathrm{Deg}}(\mathcal{G}_\mathcal{F}) \subseteq L^2_{\mathrm{deg}}(\mathcal{G}_\mathcal{F})$. Moreover, if two vertices have the same image under Deg, then they are in the same connected component of $\mathcal{G}_\mathcal{F}$ from the above analysis, and are therefore connected by a single edge by composing the morphisms corresponding to a path between them; the degree of this edge must therefore be 1 mod N by the properties of Deg. So we have the reverse inclusion, and $L^2_{\mathrm{Deg}}(\mathcal{G}_\mathcal{F}) = L^2_{\mathrm{deg}}(\mathcal{G}_\mathcal{F})$. This proves (6) and concludes the proof of Proposition 3.11.

Finally, applying Proposition 3.31 to each U_i, and the isomorphisms above, yields Theorem 3.10. □

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 Section 3.

The following lemma relates conjugation-invariance of distributions to the abstract setup of Section 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 $G \in L^2_{\mathrm{deg}}(\mathcal{G}_{\mathrm{End}/N})$, 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 mod 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

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

□

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* $\mathcal{O}$ *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* $\textsc{Rich}_k^{\mathcal{O}}(E)$*. Then for every element* $g \in$

Algorithm 1 $\textsc{Rich}_k^{\mathscr{O}}$: turning an oracle $\mathscr{O}$ for OneEnd into a 'richer' oracle $\textsc{Rich}_k^{\mathscr{O}}$, 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 OneEnd 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** $\hat{\varphi} \circ \alpha \circ \varphi$

$(\mathrm{End}(E)/N\,\mathrm{End}(E))^\times$ *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*

$$\frac{1+\sqrt{3}}{4}\lambda^k N^2\sqrt{p+13} = O(\lambda^k N^2\sqrt{p}),$$

where $\lambda = \frac{2\sqrt{2}}{3} \approx 0.94$.

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

$$F(E', \beta) = \Pr[\mathscr{O}(E') \bmod N = \beta],$$

so that

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

Indeed, $\left(\frac{A_2}{3}\right)^k F(E, \beta)$ is the average, over all random walks $\varphi\colon E \to E'$ that Algorithm 4 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 $\mathrm{End}(E')/N\,\mathrm{End}(E')$,

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

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

$$\left\| \left(\frac{A_2}{3}\right)^k F_0 \right\| \leq \lambda^k \|F_0\| \leq \lambda^k \|F\|.$$

Finally, with β ranging over $\mathrm{End}(E)/N\,\mathrm{End}(E)$, the statistical distance in the statement of the theorem is

$$\begin{aligned}
&\frac{1}{2}\sum_\beta \left|\Pr[\textsc{Rich}_k^{\mathscr{O}}(E) \bmod N = \beta] - \Pr[\textsc{Rich}_k^{\mathscr{O}}(E) \bmod N = g\beta g^{-1}]\right| \\
&= \frac{1}{2}\sum_\beta \left| \left(\frac{A_2}{3}\right)^k F(E, 4^{-k}\beta) - c_g \left(\frac{A_2}{3}\right)^k F(E, 4^{-k}\beta) \right| \\
&= \frac{1}{2}\sum_\beta \left| \left(\frac{A_2}{3}\right)^k F(E, \beta) - c_g \left(\frac{A_2}{3}\right)^k F(E, \beta) \right| \text{ since } \beta \mapsto 4^k\beta \text{ is a bijection} \\
&\leq \frac{1}{2}\left(N^4 \sum_\beta \left| \left(\frac{A_2}{3}\right)^k F(E, \beta) - c_g \left(\frac{A_2}{3}\right)^k F(E, \beta) \right|^2 \right)^{\frac{1}{2}} \text{ by the Cauchy–Schwarz inequality} \\
&\leq \frac{1}{2} N^2 \sqrt{6} \left\| \left(\frac{A_2}{3}\right)^k F - c_g \left(\frac{A_2}{3}\right)^k F \right\| \text{ since } \#\,\mathrm{Aut}(E, \beta) \leq 6 \\
&\leq \frac{1}{2}(1+\sqrt{3}) N^2 \sqrt{6} \left\| \left(\frac{A_2}{3}\right)^k F_0 \right\| \leq \frac{1}{2}(1+\sqrt{3})\lambda^k N^2 \sqrt{6} \|F\| \\
&\leq \frac{1}{2}(1+\sqrt{3})\lambda^k N^2 \sqrt{6 \cdot \frac{p+13}{24}} = \frac{1}{4}(1+\sqrt{3})\lambda^k N^2 \sqrt{p+13}, \text{ as claimed.}
\end{aligned}$$

□

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. 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 $\mathsf{Success}$ 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 present 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. The running time of $\textsc{Saturate}_\ell$ is polynomial in ℓ, not in $\log \ell$, so we only use it for small ℓ.

Proposition 6.1. *Algorithm 2 ($\textsc{Saturate}_\ell$) is correct and runs in time polynomial in ℓ and the size of the input.*

Algorithm 2 SATURATE$_\ell(R_0)$: turns an order into a super-order which is maximal at ℓ.

Require: A supersingular elliptic curve $E/\mathbf{F}_{p^2}$, a prime $\ell \neq p$, and an order $R_0 \subset \mathrm{End}(E)$.
Ensure: An order R such that $R_0 \subseteq R \subseteq \mathrm{End}(E)$ and $R \otimes \mathbf{Z}_\ell$ is maximal.
$L \leftarrow R_0$
while $\gcd(\mathrm{disc}(L), \ell) \neq 1$ **do**
  for lattices L' such that $[L' : L] = \ell$ **do**
    $\alpha \leftarrow$ an element of L such that $\alpha/\ell \in L' \setminus L$
    $\beta \leftarrow$ DIVIDE(α, ℓ) an efficient representation of α/ℓ {Proposition 2.3}
    if $\beta \neq \perp$ **then**
      $L \leftarrow L + \mathbf{Z}\,\beta$
    end if
  end for
end while
return L

Algorithm 3 SATURATERAM(R_0): turns an order into a super-order which is maximal at p.

Require: A supersingular elliptic curve $E/\mathbf{F}_{p^2}$, and an order $R_0 \subset \mathrm{End}(E)$.
Ensure: An order R such that $R_0 \subseteq R \subseteq \mathrm{End}(E)$ and $R \otimes \mathbf{Z}_p$ is maximal.
$(\mathcal{O}_0, \iota) \leftarrow$ an order $\mathcal{O}_0$ given by multiplication table, and an isomorphism $\iota\colon \mathcal{O}_0 \to R_0$
$\mathcal{O} \leftarrow$ an order containing $\mathcal{O}_0$, maximal at p with $[\mathcal{O} : \mathcal{O}_0]$ a power of p {[IR93] or [Voi13]}
$(1, b_1, b_2, b_3) \leftarrow$ a basis of $\mathcal{O}$
for $i \in \{1, 2, 3\}$ **do**
  Write $b_i = a_i/p^{k_i}$ with $a_i \in \mathcal{O}_0$
  $\alpha_i \leftarrow \iota(a_i)$
  $\beta_i \leftarrow$ DIVIDE(α_i, p^{k_i}) an efficient representation of α_i/p^{k_i} {Proposition 2.3}
end for
return $\mathbf{Z} + \mathbf{Z}\,\beta_1 + \mathbf{Z}\,\beta_2 + \mathbf{Z}\,\beta_3$

Proof. Since $\ell \neq p$, the discriminant of the maximal order $\mathrm{End}(E)$ is coprime to ℓ. Therefore, at each iteration, there is at least one L' that is contained in $\mathrm{End}(E)$, so that the division succeeds. Every iteration divides the discriminant by ℓ^2, so the number of iterations is half the valuation of $\mathrm{disc}(R_0)$ at ℓ, which is polynomial. At every iteration, there are $O(\ell^3)$ lattices L' since they correspond exactly to lines $\ell L'/\ell L \subseteq L/\ell L$. Every operation performed in the loops takes polynomial time. This proves that the algorithm terminates within the claimed running time. Consider the following properties for a lattice M in $\mathrm{End}(E)$:

(1) $R_0 \subseteq M \subseteq \mathrm{End}(E)$;
(2) $[M : R_0]$ is a power of ℓ;
(3) $[\mathrm{End}(E) : M]$ is coprime to ℓ.

There exists at most one M satsifying (1)–(3). When the algorithm terminates, the lattice L satisfies (1)–(3). On the other hand, there exists an order R satisfying (1)–(3). Therefore $R = L$, as claimed. □

Proposition 6.2. *Algorithm 3 (*SaturateRam*) is correct and runs in polynomial time.*

Proof. Since $R_0 \otimes \mathbf{Q}$ is ramified at p, there is a unique order R containing R_0, maximal at p with $[R : R_0]$ a power of p, and this order is contained in $\mathrm{End}(E)$. This imples $(\iota \otimes \mathbf{Q})(\mathcal{O}) = R \subset \mathrm{End}(E)$, so all the divisions succeed, the family $(1, \beta_1, \beta_2, \beta_3)$ is a basis of R, and the algorithm is correct. All the operations take polynomial time. □

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

Algorithm 4 $\textsc{Reduce}_N(\alpha)$: 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** $\mathrm{Tr}(\beta) \equiv 0 \mod 4$ **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 4 (*$\textsc{Reduce}_N$*) 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 5 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 OneEnd 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**
  $\alpha \leftarrow \textsc{Rich}_{k_1}^{\mathscr{O}}(E)$, a random endomorphism of E {Algorithm 1}
  $R \leftarrow$ the ring generated by R and α
end while
$R \leftarrow \textsc{SaturateRam}_2(R)$ {Algorithm 2}
$R \leftarrow \textsc{SaturateRam}(R)$ {Algorithm 3}
$[\mathrm{End}(E) : R] \leftarrow \sqrt{\mathrm{disc}(R)/p}$
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 $\textsc{Reduce}_N(\alpha)$
  $\alpha_i \leftarrow \textsc{Rich}_{k_2}^{\mathscr{O}_N}(E)$ 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 $\textsc{Rich}_k^{\mathscr{O}_N}$ *satisfies the conditions of Proposition 5.11.*

Proof. Let $\varphi\colon E \to E'$ of degree a power of 2. For any endomorphism $\beta \in \mathrm{End}(E')$, 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 $\textsc{Rich}_k^{\mathscr{O}_N}$ is of the form $\hat{\varphi} \circ \textsc{Reduce}_N(\alpha) \circ \varphi$, so is N-reduced. So the distribution of $\textsc{Rich}_k^{\mathscr{O}_N}$ satisfies Item 1 of Proposition 5.11.

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

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

With

$$k \geq \frac{\log\left(6000000 \cdot (\log N)^{12} \cdot \frac{1+\sqrt{3}}{4} N^2\sqrt{p+13}\right)}{\log\left(\frac{3}{2\sqrt{2}}\right)},$$

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 5 is a reduction from* EndRing *to* $\textsc{OneEnd}_\lambda$ *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 $\mathrm{End}(E)$, and the success condition $[\mathrm{End}(E) : R] = 1$ implies $R = \mathrm{End}(E)$.

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 $\mathrm{SL}_2(\mathbf{F}_3)$-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 $t = \max_i \mathrm{lev}_3(\alpha_i)$, and $\beta_i = 3^{t-\mathrm{lev}_3(\alpha_i)}\alpha_i$, so all β_i are at the same level t. Like the variables α_i, the variables β_i are ε-close to $\mathrm{SL}_2(\mathbf{F}_3)$-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

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

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 Equation (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

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

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 Section 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 8.7 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 variant SQIsignHD [DLRW23], offer the most compact public keys and signatures of all known post-quantum constructions. Note that the results from this section apply equally to SQIsign and SQIsignHD. This digital signature scheme 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] heuristically argues that the protocol is sound if EndRing is hard, our main theorem allows us to prove it.

Let SQIsign.param be the SQIsign 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 SQIsign.keygen be the SQIsign 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.

Let $\mathscr{V}$ be an honest verifier for the SQIsign identification protocol. For any (malicious) prover $\mathscr{P}^*$ and parameters pp, run the following experiment: first, sample a key pair $(\mathsf{pk}, \mathsf{sk}) \leftarrow \text{SQIsign.keygen}(\mathsf{pp})$, and give pk to $\mathscr{P}^*$. Then, run the SQIsign 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 an 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 SQIsign). *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 $\mathscr{P}^*$. From [DKL+20, Theorem 1], there is an algorithm of expected running time

$$r' = O\left(\frac{r}{\mathrm{Adv}^{\mathscr{P}^*}_{\mathrm{SQIsound}}(\mathsf{pp})}\right)$$

for the *supersingular smooth endomorphism problem*, with solutions of length $O(\log p)$. The supersingular smooth endomorphism problem is defined as OneEnd with the additional contraint that the output has smooth degree. Therefore, the same algorithm solves OneEnd. The result follows from the equivalence between OneEnd and EndRing (Theorem 1.1). □

The same theorem is true 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. In the proof, we would get

$$r' = O\left(\frac{r}{\pi^{\mathscr{P}^*}(\mathsf{pp}) - 1/c}\right).$$

Note that r' is not necessarily polynomial in $\pi^{\mathscr{P}^*}(\mathsf{pp})^{-1}$, as $\pi^{\mathscr{P}^*}(\mathsf{pp})$ could be arbitrarily close to $1/c$. We can consider two cases: first, if $\pi^{\mathscr{P}^*}(\mathsf{pp}) > 2/c$, we have $\pi^{\mathscr{P}^*}(\mathsf{pp}) - 1/c > \pi^{\mathscr{P}^*}(\mathsf{pp})/2$ and we can conclude as above. Second, if $\pi^{\mathscr{P}^*}(\mathsf{pp}) \leq 2/c$, then $\pi^{\mathscr{P}^*}(\mathsf{pp})^{-1} \geq c/2 = \Omega(2^k) = \Omega(p^{1/2})$, so we can conclude from the fact that there exists an algorithm for EndRing in expected time $\tilde{O}(p^{1/2})$ (see Theorem 8.7).

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 8.3 (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 IsogenyPath$_\lambda$ problem denotes the IsogenyPath 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 IsogenyPath reduces to EndRing.

Proposition 8.4 (IsogenyPath reduces to EndRing). *Assuming the generalised Riemann hypothesis, the problem* IsogenyPath$_\lambda$ *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. IsogenyPath 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 IsogenyPath$_\lambda$ reduces to EndRing. □

The converse reduction is trickier. As a solution to IsogenyPath 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 8.5 below is unconditional. In particular, it implies that EndRing reduces to the ℓ-isogeny path problem independently of the generalised Riemann hypothesis.

Theorem 8.5 (EndRing reduces to IsogenyPath). *Given an oracle for* IsogenyPath$_\lambda$, *there is an algorithm for* EndRing *that runs in expected polynomial time in* $\log(p)$ *and* $\lambda(\log p)$.

Algorithm 6 Solving OneEnd given an IsogenyPath oracle.

Require: A supersingular elliptic curve $E/\mathbf{F}_{p^2}$, a parameter $\varepsilon > 0$, an oracle $\mathscr{O}_{\text{IsogenyPath}}$ solving the IsogenyPath$_\lambda$ problem.
Ensure: An endomorphism $\alpha \in \mathrm{End}(E) \setminus \mathbf{Z}$ in efficient representation.

$P \leftarrow$ an arbitrary nonzero point in $E[2]$
$n \leftarrow \lceil 2\log_3(p) - 4\log_3(\varepsilon) \rceil$
while true **do**
  $\varphi \leftarrow$ a non-backtracking random walk $\varphi\colon E \to E'$ of length n in the 3-isogeny graph
  $\nu \leftarrow$ the isogeny $\nu\colon E' \to E''$ of kernel $\langle\varphi(P)\rangle$
  $\psi \leftarrow \mathscr{O}_{\text{IsogenyPath}}(E'', E)$, an isogeny $\psi\colon E'' \to E$
  $\alpha \leftarrow (\psi \circ \nu \circ \varphi)/2^e \in \mathrm{End}(E)$ for the largest possible e
  if $2 \mid \deg(\alpha)$ **then**
    return α
  end if
end while

Proof. Since EndRing is equivalent to OneEnd (Theorem 1.1), let us prove that OneEnd reduces to IsogenyPath. Suppose we have an oracle $\mathscr{O}_{\text{IsogenyPath}}$ for IsogenyPath$_\lambda$. Let E be a supersingular curve for which we want to solve OneEnd. Consider

a parameter ε. The reduction is described in Algorithm 6. 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 $\|P - Q\|_1 < \varepsilon$, since φ has length $\lceil 2\log_3(p) - 4\log_3(\varepsilon) \rceil$.

We now obtain the following bound:

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

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 $2/3 - \varepsilon$. With $\varepsilon < 1/3$, the expected number of iterations is at most $(2/3 - \varepsilon)^{-1} \leq 3 = O(1)$. □

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 8.5). Let us start by recalling the following folklore solution to ℓ-IsogenyPath.

Proposition 8.6. *Algorithm 7 solves the* ℓ-IsogenyPath *problem in expected time* $\mathrm{poly}(\ell, \log p)p^{1/2}$ *and returns paths of length* $O(\log p)$.

Algorithm 7 Solving ℓ-IsogenyPath.

Require: Two supersingular elliptic curves $E_0/\mathbf{F}_{p^2}$ and $E_1/\mathbf{F}_{p^2}$, a parameter n.
Ensure: An ℓ-isogeny path $E_0 \to E_1$.

$T \leftarrow \emptyset$ an empty hash table
while $\#T < p^{1/2}$ **do**
  $\varphi \leftarrow$ a random walk $\varphi\colon E_0 \to E$ of length n in the ℓ-isogeny graph
  if $j(E)$ is not the key of any entry in T **then**
    Record φ in T, with key $j(E)$
  end if
end while
while true **do**
  $\psi \leftarrow$ a random walk $\psi\colon E_1 \to E$ of length n in the ℓ-isogeny graph
  if $j(E)$ is the key of a recorded entry φ in T **then**
    return $\hat{\varphi} \circ \psi\colon E_0 \to E_1$
  end if
end while

Proof. Algorithm 7 is a standard bi-directional pathfinding algorithm. Choose the parameter n as in Proposition 2.7, so that random isogeny paths of length $O(\log p)$ reach a target at statistical distance $O(1/p)$ from uniform. The ℓ-isogeny graph has $O(p)$ vertices, and each sampled curve is close to uniform, so the table is complete after $O(p^{1/2})$ iterations of the first loop. By the same token, thanks to the birthday paradox, the second loop finds a matching entry after an expected $O(p^{1/2})$ number of attempts. The factor $\mathrm{poly}(\ell, \log p)$ accounts for the cost of sampling an isogeny path and checking that a candidate is in the table. □

Theorem 8.7. *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 8.6), and EndRing reduces to polynomially many instances of the ℓ-isogeny path problem (Theorem 8.5). □

References

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

BCC+23. Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. Supersingular curves you can trust. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023*, volume 14005 of *Lecture Notes in Computer Science*, pages 405–437. Springer, 2023.

CD23. Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In *Advances in cryptology—EUROCRYPT 2023. Part V*, volume 14008 of *Lecture Notes in Comput. Sci.*, pages 423–447. Springer, Cham, [2023] ©2023.

CL23. Giulio Codogni and Guido Lido. Spectral theory of isogeny graphs. Preprint arXiv:2308.13913, 2023. https://arxiv.org/abs/2308.13913.

CLG09. Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, Jan 2009.

Del73. Pierre Deligne. La conjecture de Weil. I. *Publ. Math., Inst. Hautes Étud. Sci.*, 43:273–307, 1973.

Deu41. Max Deuring. Die typen der multiplikatorenringe elliptischer funktionenkörper. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 14(1):197–272, 1941.

DKL+20. Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. Sqisign: Compact post-quantum signatures from quaternions and isogenies. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security*, volume 12491 of *Lecture Notes in Computer Science*, pages 64–93. Springer, 2020.

DLRW23. Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQISignHD: New dimensions in cryptography. IACR Cryptology ePrint Archive, Report 2023/436, 2023. https://eprint.iacr.org/2023/436.

DV13. Lassina Dembélé and John Voight. Explicit methods for Hilbert modular forms. In *Elliptic curves, Hilbert modular forms and Galois deformations.*, pages 135–198. Basel: Birkhäuser/Springer, 2013.

EHL+18. Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 329–368, Cham, 2018. Springer International Publishing.

FIK+23. Jenny Fuselier, Annamaria Iezzi, Mark Kozek, Travis Morrison, and Changningphaabi Namoijam. Computing supersingular endomorphism rings using inseparable endomorphisms. Preprint arXiv:2306.03051, 2023. https://arxiv.org/abs/2306.03051.

HLMW23. Arthur Herlédan Le Merdy and Benjamin Wesolowski. The supersingular endomorphism ring problem given one endomorphism. Preprint arXiv:2309.11912, 2023. https://arxiv.org/abs/2309.11912.

IR93. Gábor Ivanyos and Lajos Rónyai. Finding maximal orders in semisimple algebras over $\mathbb{Q}$. *Comput. Complexity*, 3(3):245–261, 1993.

JL70. H. Jacquet and R. P. Langlands. *Automorphic forms on GL(2)*, volume 114 of *Lect. Notes Math.* Springer, Cham, 1970.

KLPT14. David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ-isogeny path problem. *LMS Journal of Computation and Mathematics*, 17(A):418–432, 2014.

Koh96. David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.

LM23. Antonio Lei and Katharina Müller. On towers of isogeny graphs with full level structure. Preprint arXiv:2309.00524, 2023. `https://arxiv.org/abs/2309.00524`.

Mes86. Jean-Francois Mestre. La méthode des graphes. exemples et applications. In *Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata)*, pages 217–242, 1986.

ML98. Saunders Mac Lane. *Categories for the working mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1998.

MMP+23. Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023. Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 448–471. Springer, 2023.

Piz90. Arnold K Pizer. Ramanujan graphs and Hecke operators. *Bulletin of the American Mathematical Society*, 23(1):127–137, 1990.

Rob22. Damien Robert. 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. Damien Robert. Breaking SIDH in polynomial time. In *Advances in cryptology—EUROCRYPT 2023. Part V*, volume 14008 of *Lecture Notes in Comput. Sci.*, pages 472–503. Springer, Cham, [2023] ©2023.

Sil86. Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Gradute Texts in Mathematics*. Springer-Verlag, 1986.

Voi13. John Voight. Identifying the matrix ring: algorithms for quaternion algebras and quadratic forms. In *Quadratic and higher degree forms*, pages 255–298. New York, NY: Springer, 2013.

Voi21. John Voight. *Quaternion Algebras*. Springer International Publishing, 2021. Graduate Texts in Mathematics, No. 288.

Wes22a. Benjamin Wesolowski. Orientations and the supersingular endomorphism ring problem. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, volume 13277 of *Lecture Notes in Computer Science*, pages 345–371. Springer, 2022.

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

A Illustration of our equidistribution theorem

In this appendix, we give some more examples and comments on the use of Theorem 3.10 and Proposition 3.11.

A.1 Distributions and functions

Recall the definition of our adjacency operator on $L^2(\mathcal{G}_{\mathcal{F}})$:

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

where the sum runs over edges of degree ℓ leaving x. This sum always has $\ell + 1$ terms. In some situations it is more natural to use a different operator B_ℓ:

$$B_\ell F(x) = \sum_{x \leftarrow y} F(y)$$

where the sum runs over edges of degree ℓ arriving at x. This sum may have fewer than $\ell + 1$ terms, due to automorphisms.

In order to relate the two operators, it is convenient to compute the adjoint of A_ℓ. Let μ be the measure on $\mathcal{G}_{\mathcal{F}}$ (recall $\mu(x) = \frac{1}{\#\operatorname{Aut}(x)}$). We have

$$\begin{aligned}
\langle A_\ell F, G\rangle &= \sum_x A_\ell F(x)\overline{G(x)}\mu(x) \\
&= \sum_x \sum_{x \to y} F(y)\overline{G(x)}\mu(x) \\
&= \sum_y \sum_{x \leftarrow y} F(y)\overline{G(x)}\mu(x) \\
&= \sum_y F(y) \left(\frac{1}{\mu(y)} \sum_{x \leftarrow y} \overline{G(x)}\mu(x) \right) \mu(y) \\
&= \langle F, A_\ell^* G\rangle,
\end{aligned}$$

where

$$A_\ell^* G(x) = \frac{1}{\mu(x)} \sum_{x \leftarrow y} G(y)\mu(y).$$

We therefore introduce the "diagonal" operator M on $L^2(\mathcal{G}_{\mathcal{F}})$ defined by

$$MF(x) = F(x)\mu(x),$$

so that we have

$$A_\ell^* = M^{-1} B_\ell M \text{ i.e. } B_\ell = M A_\ell^* M^{-1}.$$

Since A_ℓ^* has the same orthogonal eigenvectors as A_ℓ, with complex conjugate eigenvalues, this gives the spectral decomposition of B_ℓ.

For instance, the action of one step of a degree ℓ random walk on a distribution on $\mathcal{G}_\mathcal{F}$ is given by $\frac{B_\ell}{\ell+1}$. Therefore, in the case where $\mathcal{G}_\mathcal{F}^1$ is connected, $L^2_{\deg}(\mathcal{G}_\mathcal{F})$ is one-dimensional generated by the constant function $\mathbf{1}$ equal to 1 everywhere, so we obtain that for every F, the sequence $(\frac{A^*_\ell}{\ell+1})^k F$ quickly converges to a constant function, and therefore for every distribution f, the distribution $(\frac{B_\ell}{\ell+1})^k f$ obtained after k steps of the random walk converges to the distribution $M\mathbf{1} = \mu$, i.e. to the stationary distribution.

A.2 Explicit orthogonal projection onto $L^2_{\deg}$

Another useful tool is the explicit decomposition of functions according to $L^2(\mathcal{G}_\mathcal{F}) = L^2_{\deg}(\mathcal{G}_\mathcal{F}) \oplus L^2_0(\mathcal{G}_\mathcal{F})$. We first introduce some notation: for two vertices x, y of $\mathcal{G}_\mathcal{F}$, write $x \sim y$ if they are in the same connected component of $\mathcal{G}_\mathcal{F}^1$ (this is an equivalence relation). Moreover, for every vertex x, let

$$W(x) = \sum_{y \sim x} \mu(y).$$

If $x \sim y$, then $W(x) = W(y)$. We now define an operator P on $L^2(\mathcal{G}_\mathcal{F})$:

$$PF(x) = \frac{1}{W(x)} \sum_{y \sim x} F(y)\mu(y).$$

The function PF is clearly in $L^2_{\deg}(\mathcal{G}_\mathcal{F})$, and if $F \in L^2_{\deg}(\mathcal{G}_\mathcal{F})$ then $PF = F$; therefore P is a projector onto $L^2_{\deg}(\mathcal{G}_\mathcal{F})$. In order to prove that P is the desired orthogonal projector, it is enough to show that it is self-adjoint. Let $F, G \in L^2(\mathcal{G}_\mathcal{F})$, then

$$\begin{aligned}
\langle PF, G\rangle &= \sum_x PF(x)\overline{G(x)}\mu(x) \\
&= \sum_x \frac{1}{W(x)} \sum_{y\sim x} F(y)\mu(y)\overline{G(x)}\mu(x) \\
&= \sum_y \sum_{x\sim y} \frac{1}{W(x)} F(y)\mu(y)\overline{G(x)}\mu(x) \\
&= \sum_y \sum_{x\sim y} \frac{1}{W(y)} F(y)\mu(y)\overline{G(x)}\mu(x) \\
&= \sum_y F(y) \left(\frac{1}{W(y)} \sum_{x\sim y} \overline{G(x)}\mu(x) \right) \mu(y) \\
&= \langle F, PG\rangle.
\end{aligned}$$

So P, being self-adjoint, is the orthogonal projection onto $L^2_{\deg}(\mathcal{G}_\mathcal{F})$. Note that one can also see this formula as the sum of the orthogonal projections onto the indicator functions of the connected components of $\mathcal{G}_\mathcal{F}^1$.

A.3 Equidistribution over an entire connected component of $\mathcal{G}_{\mathcal{F}}$

Assume that $\mathcal{G}_{\mathcal{F}}$ is connected (one can always reduce to this case by considering connected components). The existence of the map Deg and the resulting action of A_ℓ on $L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$ often force the ℓ-part of $\mathcal{G}_{\mathcal{F}}$ to be disconnected and/or multipartite, preventing degree ℓ random walks from properly equidistributing. One may be tempted to think that there is a fundamental obstruction to the equidistribution of random walks on the whole of $\mathcal{G}_{\mathcal{F}}$. We will see here that this is not the case: one can easily obtain full equidistribution, simply by using several primes for the random walk (in other words, in general the disconnectedness of $\mathcal{G}_{\mathcal{F}}$ is the only fundamental obstruction to equidistribution).

Pick a bound X, and assume that Σ contains every prime $\ell < X$ that does not divide pN. Let $\mathcal{N}(X)$ denote the number of such primes, so that we have $\mathcal{N}(X) \approx \frac{X}{\log X}$. Define the operator Δ on $L^2(\mathcal{G}_{\mathcal{F}})$ by

$$\Delta = \frac{1}{\mathcal{N}(X)} \sum_{\ell<X} \frac{A_\ell}{\ell+1}.$$

The interpretation of Δ is that one step of the corresponding random walk consists in choosing a prime $\ell < X$ uniformly at random, and then using one step of the degree ℓ random walk.

Since the A_ℓ are normal operators that pairwise commute, Δ is also a normal operator, stabilises $L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$ and $L^2_0(\mathcal{G}_{\mathcal{F}})$, and is diagonalisable in the same orthogonal basis as the A_ℓ. We bound its eigenvalues and operator norm.

- On $L^2_0(\mathcal{G}_{\mathcal{F}})$, the operator norm of Δ is bounded by

$$\frac{1}{\mathcal{N}(X)} \sum_{\ell<X} \frac{2\sqrt{\ell}}{\ell+1} \approx \frac{1}{\sqrt{X}}.$$

- On $L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$, there is one eigenvector for each character χ of $(\mathbf{Z}/N\mathbf{Z})^\times$ that vanishes on $\deg(H)$ (with the notations of Proposition 3.11), with eigenvalue 1 if χ is the trivial character, and otherwise

$$\frac{1}{\mathcal{N}(X)} \sum_{\ell<X} \chi(\ell) \approx \frac{1}{\sqrt{X}}.$$

For X large enough, all the eigenvalues are therefore small, except for the eigenvalue 1 corresponding to the constant function. Moreover, these approximation can be turned into good bounds under the generalised Riemann hypothesis. Therefore, for a moderate value of X, the random walk corresponding to Δ will quickly equidistribute over the whole of $\mathcal{G}_{\mathcal{F}}$.

A.4 Example: graphs attached to endomorphisms modulo ℓ

It is often convenient to use a functor slightly different from End $/N$, namely the functor $\mathcal{F}$ (for Σ the set of all primes not dividing N) defined by

- $\mathcal{F}(E) = \mathrm{End}(E)/N\,\mathrm{End}(E)$;
- $\mathcal{F}(\varphi)\colon \alpha \mapsto (\deg\varphi)^{-1}\varphi\alpha\hat{\varphi}$ (which makes sense: $\deg\varphi$ is invertible $\bmod N$).

The main reason to prefer this functor is that for every isogeny $\varphi\colon E \to E'$, the map $\mathcal{F}(\varphi)\colon \mathrm{End}(E)/N\,\mathrm{End}(E) \to \mathrm{End}(E')/N\,\mathrm{End}(E')$ is a ring homomorphism, and therefore preserves the trace, degree, dual, minimal polynomial and level. This functor clearly satisfies the $(\bmod N)$-congruence property. We will describe the connected components of the graphs $\mathcal{G}_{\mathcal{F}}$, and $\mathcal{G}^1_{\mathcal{F}}$, and the graph $\mathcal{G}_{\deg}$.

For simplicity, we will assume $N = \ell$ is an odd prime different from p, but the other cases can be treated similarly. We recall the classification of conjugacy classes in $M_2(\mathbf{F}_\ell)$ and their centraliser in $\mathrm{GL}_2(\mathbf{F}_\ell)$:

(1) Homotheties
$$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix} \text{ for } a \in \mathbf{F}_\ell\,.$$
The centraliser H of such a matrix is $\mathrm{GL}_2(\mathbf{F}_\ell)$, and $\det(H) = \mathbf{F}_\ell^\times$.

(2) Diagonalisable matrices, with representatives
$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} \text{ for } a \neq b \in \mathbf{F}_\ell\,.$$
The centraliser H of such a matrix is
$$\begin{pmatrix} * & 0 \\ 0 & * \end{pmatrix} \cong \mathbf{F}_\ell^\times \times \mathbf{F}_\ell^\times,$$
and $\det(H) = \mathbf{F}_\ell^\times$. Two matrices in such a conjugacy class are also $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugate.

(3) Semisimple, non-diagonalisable matrices, with representatives
$$\Psi(\lambda) \text{ for } \lambda \in \mathbf{F}_{\ell^2},$$
where $\Psi\colon \mathbf{F}_{\ell^2} \to M_2(\mathbf{F}_\ell)$ is the ring homomorphism given by the action on an $\mathbf{F}_\ell$-basis of $\mathbf{F}_{\ell^2}$. The centraliser of such a matrix is $H = \Psi(\mathbf{F}_{\ell^2}^\times) \cong \mathbf{F}_{\ell^2}^\times$ and $\det(H) = \mathbf{F}_\ell^\times$. Two matrices in such a conjugacy class are also $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugate.

(4) Non-semisimple matrices, with representatives
$$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix} \text{ for } a \in \mathbf{F}_\ell\,.$$
The centraliser of such a matrix is
$$H = \left\{ \begin{pmatrix} c & d \\ 0 & c \end{pmatrix} \mid c \in \mathbf{F}_\ell^\times, d \in \mathbf{F}_\ell \right\},$$
and $\det(H) = (\mathbf{F}_\ell^\times)^2$. Each such conjugacy class splits into two $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugacy class, with representatives
$$\begin{pmatrix} a & \varepsilon \\ 0 & a \end{pmatrix} \text{ for } a \in \mathbf{F}_\ell \text{ and } \varepsilon \in \mathbf{F}_\ell^\times/(\mathbf{F}_\ell^\times)^2;$$

under conjugacy by an element $g \in \mathrm{GL}_2(\mathbf{F}_\ell)$, the entry ε of the representative gets multiplied by $\det(g)$ modulo $(\mathbf{F}_\ell^\times)^2$.

In particular, conjugacy classes are completely characterised by the pair (characteristic polynomial, minimal polynomial).

For every $E \in \mathrm{SS}(p)$, choose an isomorphism $\psi_E \colon \mathrm{End}(E)/\ell\,\mathrm{End}(E) \cong M_2(\mathbf{F}_\ell)$. By Proposition 3.11 (3), two vertices (E, α) and (E, β) above the same curve E are connected if and only if α and β are conjugate in $\mathrm{End}(E)/\ell\,\mathrm{End}(E)$, if and only if $\psi_E(\alpha)$ and $\psi_E(\beta)$ are conjugate. Moreover, since any two curves are connected and the $\mathcal{F}(\varphi)$ are ring isomorphisms, the connected components of $\mathcal{G}_\mathcal{F}$ are exactly the

$$\{(E, \alpha) : \psi_E(\alpha) \in C\}$$

where C is a conjugacy class in $M_2(\mathbf{F}_\ell)$. Fix $\eta \in \mathbf{F}_\ell^\times \setminus (\mathbf{F}_\ell^\times)^2$. We have $\deg(H) = \mathbf{F}_\ell^\times$ for all types, except (4) where $\deg(H) = (\mathbf{F}_\ell^\times)^2$. So the graph $\mathcal{G}_{\deg}$ consists in one isolated vertex for each C, except for type (4) which gives a connected component with two vertices $\{1, \eta\}$, connected by edges labelled by the elements of $\mathbf{F}_\ell^\times \setminus (\mathbf{F}_\ell^\times)^2$.

We now consider $\mathcal{G}^1_\mathcal{F}$. By Proposition 3.11 (3), two vertices (E, α) and (E, β) above the same curve E are connected in $\mathcal{G}^1_\mathcal{F}$ if and only if α and β are conjugate in $\mathrm{End}(E)/\ell\,\mathrm{End}(E)$ by an element of degree $1 \bmod \ell$, if and only if $\psi_E(\alpha)$ and $\psi_E(\beta)$ are $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugate.

For any ring R and $g \in R^\times$, let $\mathrm{Ad}(g)$ denote the endomorphism of R given by $r \mapsto grg^{-1}$. Let $E \in \mathrm{SS}(p)$. By Proposition 3.11 (2), there exists $\varphi \colon E_0 \to E$ of degree $1 \bmod \ell$. The composition $\psi_E \circ \mathcal{F}(\varphi) \circ \psi_{E_0}^{-1}$ is an automorphism of $M_2(\mathbf{F}_\ell)$, hence of the form $\mathrm{Ad}(g)$ for some $g \in \mathrm{GL}_2(\mathbf{F}_\ell)$. Changing ψ_E if necessary (by an interior automorphism), we may assume that $\det g \in (\mathbf{F}_\ell^\times)^2$. In this case, this property holds for all $\varphi' \colon E_0 \to E$ of degree $1 \bmod \ell$:

$$\begin{aligned}
\psi_E \circ \mathcal{F}(\varphi') \circ \psi_{E_0}^{-1} &= \psi_E \circ \mathcal{F}(\varphi) \circ \mathcal{F}(\varphi^{-1}\varphi') \circ \psi_{E_0}^{-1} \\
&= \psi_E \circ \mathcal{F}(\varphi) \circ \mathrm{Ad}(\varphi^{-1}\varphi') \circ \psi_{E_0}^{-1} \\
&= \psi_E \circ \mathcal{F}(\varphi) \circ \psi_{E_0}^{-1} \circ \mathrm{Ad}(\psi_{E_0}(\varphi^{-1}\varphi')) \\
&= \mathrm{Ad}(g) \circ \mathrm{Ad}(\psi_{E_0}(\varphi^{-1}\varphi')) \\
&= \mathrm{Ad}(g\psi_{E_0}(\varphi^{-1}\varphi')),
\end{aligned}$$

and $\det \psi_{E_0}(\varphi^{-1}\varphi') = \deg(\varphi^{-1}\varphi') = 1 \bmod \ell$.

Now let (E_0, α) and (E, β) be vertices. Since E and E_0 are connected by an isogeny φ of degree $1 \bmod \ell$, the vertices (E_0, α) and (E, β) are connected if and only if $(E, \mathcal{F}(\varphi)(\alpha))$ and (E, β) are connected, if and only if $\psi_E(\mathcal{F}(\varphi)(\alpha))$ and $\psi_E(\beta)$ are $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugate, if and only if $\psi_{E_0}(\alpha)$ and $\psi_E(\beta)$ are $\mathrm{SL}_2(\mathbf{F}_\ell)$-conjugate by our assumption on ψ_E.

In other words, the connected components of $\mathcal{G}^1_\mathcal{F}$ are the

$$\{(E, \alpha) : \psi_E(\alpha) \in C\}$$

for C a conjugacy class in $M_2(\mathbf{F}_\ell)$ not of type (4), and the

$$\{(E, \alpha) : \psi_E(\alpha) \in C_1\} \text{ and } \{(E, \alpha) : \psi_E(\alpha) \in C_\eta\}$$

for C a conjugacy class in $M_2(\mathbf{F}_\ell)$ of type (4). (The point of our assumption was to make sure that components C_1 and C_η are not swapped by an isogeny φ of degree 1 mod ℓ.)

A.5 Example: endomorphism transported by a random walk

Let us examine the following situation: let $E_0 \in \mathrm{SS}(p)$ and $\alpha_0 \in \mathrm{End}(E_0)$, let $\varphi\colon E_0 \to E$ be the result of a k-steps random walk of 2-isogenies, and let $\alpha = \varphi\alpha_0\hat{\varphi}$. What is the distribution of $(E, \alpha \bmod N)$ as $k \to \infty$? Again, for simplicity we treat the prime case $N = \ell \neq p$.

We first determine the behaviour of the random walk using the functor $\mathcal{F}$ from Section A.4. We choose $\psi_E\colon \mathrm{End}(E)/\ell\,\mathrm{End}(E) \cong M_2(\mathbf{F}_\ell)$ with the same compatibility condition as in that section. Let C be the conjugacy class of the matrix $\psi_{E_0}(\alpha \bmod \ell)$.

Let f be a distribution on the vertices of $\mathcal{G}_\mathcal{F}$. As in Section A.1, since the effect of one step of random walk is given by $f \mapsto \frac{B_2}{3}f$, it is convenient to encode f into a function $F \in L^2(\mathcal{G}_\mathcal{F})$ by the formula $F(E,\beta) = f(E,\beta)\mu(E,\beta)$, so that the action is given by $F \mapsto \frac{A_2^*}{3}F$. The initial distribution f_0 is defined by

$$f_0(E_0, \alpha_0 \bmod \ell) = 1$$

and f_0 is 0 everywhere else, so the corresponding initial function F_0 is defined by

$$F_0(E_0, \alpha_0 \bmod \ell) = \mu(E_0, \alpha_0 \bmod \ell)^{-1}$$

and F_0 is 0 everywhere else. Using the projection formula from Section A.2, we see that PF_0 is the indicator function of the connected component of $(E_0, \alpha_0 \bmod \ell)$, scaled by $W(E_0, \alpha_0 \bmod \ell)$. We use the corresponding function on $\mathcal{G}_\mathrm{deg}$ to determine the action of $\frac{A_2^*}{3}$.

- C is not of type (4): the vertex $\mathrm{Deg}(E_0, \alpha_0 \bmod \ell)$ is an isolated vertex in $\mathcal{G}_\mathrm{deg}$, and the action of $\frac{A_2^*}{3}$ is trivial.
- C is of type (4): the connected component of $\mathrm{Deg}(E_0, \alpha_0 \bmod \ell)$ in $\mathcal{G}_\mathrm{deg}$ consists of two vertices. The action of $\frac{A_2^*}{3}$ is trivial if $2 \in (\mathbf{F}_\ell^\times)^2$, and swaps the two vertices otherwise.

Coming back to functions on $\mathcal{G}_\mathcal{F}$, we obtain the following.

- If C is not of type (4), or if 2 is a square modulo ℓ, or if k is even, then we have $(\frac{A_2^*}{3})^k PF_0 = PF_0$, i.e. $(\frac{B_2}{3})^k f_0$ is close to the distribution supported on the connected component of $(E_0, \alpha_0 \bmod \ell)$ and where the probability of (E,β) is proportional to $\mu(E,\beta)$.
- Otherwise, $(\frac{A_2^*}{3})^k PF_0$ is the scaled indicator function of the connected component of $\mathcal{G}_\mathcal{F}^1$ corresponding to $C_{\eta\,\varepsilon}$, where $\psi_{E_0}(\alpha_0 \bmod \ell) \in C_\varepsilon$, i.e. $(\frac{B_2}{3})^k f_0$ is close to the distribution supported on that connected component and where the probability of (E,β) is proportional to $\mu(E,\beta)$.

Finally, the actual distribution of $\alpha \bmod \ell$ is obtained by taking the distributions described above, and multiplying the corresponding random endomorphism mod ℓ by 2^k (which usually changes the conjugacy class). The statistical distance to the distribution obtained by projection onto $L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$ can be estimated using the eigenvalue bounds of Theorem 3.10.

A.6 Distribution of isogenies produced by random walks

Let ℓ be a prime and $E_0 \in \mathrm{SS}(p)$. A natural question is: what is the distribution of the isogenies $\varphi\colon E_0 \to E$ produced by a long random ℓ-isogeny walk starting from E_0? Our equidistribution theorem gives nontrivial information about this, in the following form. Let $N \geq 2$ be an integer not divisible by ℓ. We are going to look at the distibution of $\varphi \bmod N \in \mathrm{Hom}(E_0, E)/N\,\mathrm{Hom}(E_0, E)$.

Let Σ be the set of all primes not dividing N. We introduce the functor $\mathcal{F} = (\mathrm{Hom}(E_0, -)/N)^\times\colon \mathrm{SS}_\Sigma(p) \to \mathrm{Sets}$:

- $\mathcal{F}(E) = \{\varphi \in \mathrm{Hom}(E_0, E)/N\,\mathrm{Hom}(E_0, E) \mid \deg(\varphi) \in (\mathbf{Z}/N\,\mathbf{Z})^\times\}$;
- for $\psi \in \mathrm{Hom}_\Sigma(E, E')$, the map $\mathcal{F}(\psi)\colon \mathcal{F}(E) \to \mathcal{F}(E')$ is $\varphi \mapsto \psi \circ \varphi$.

The functor $\mathcal{F}$ clearly satisfies the (mod N)-congruence property.

The action of the group $G = (\mathrm{End}(E_0)/N\,\mathrm{End}(E_0))^\times$ on the set $\mathcal{F}(E_0) = (\mathrm{End}(E_0)/N\,\mathrm{End}(E_0))^\times$ is the left regular action, and therefore has a unique orbit with trivial stabiliser $H = 1$ (we choose $x = \mathrm{id}$ as the base point). Therefore $\mathcal{G}_{\deg}$ is the Cayley graph of $(\mathbf{Z}/N\,\mathbf{Z})^\times$, and the map $\mathrm{Deg}\colon \mathcal{G}_{\mathcal{F}} \to \mathcal{G}_{\deg}$ sends (E, φ) to $\deg\varphi \bmod N$.

We obtain the following proposition.

Proposition A.1. *Let $k \geq 0$. Let $\varphi\colon E_0 \to E_\varphi$ be the random isogeny obtained by a k-step ℓ-isogeny random walk from E_0. Let ν be the distribution on pairs (E, ψ) where $\psi \in \mathrm{Hom}(E_0, E)/N\,\mathrm{Hom}(E_0, E)$ has degree $\ell^k \bmod N$, up to isomorphism, given by probability proportional to $\frac{1}{\#\,\mathrm{Aut}(E,\psi)}$. Then the statistical distance between the distribution of $(E_\varphi, \varphi \bmod N)$ and ν is at most*

$$\frac{1}{2\sqrt{6}}\left(\frac{2\sqrt{\ell}}{\ell+1}\right)^k N^2\sqrt{p}.$$

Proof. Let $f \in L^2(\mathcal{G}_{\mathcal{F}})$ be such that $f(E_0, \mathrm{id}) = \#\,\mathrm{Aut}(E_0, \mathrm{id})$ and $f(E, \psi) = 0$ for all other vertices (E, ψ). Then for all $(E, \psi) \in \mathcal{G}_{\mathcal{F}}$ we have

$$\frac{1}{\#\,\mathrm{Aut}(E,\psi)}\left(\frac{A^*_\ell}{\ell+1}\right)^k f(E, \psi) = \Pr[(E_\varphi, \varphi \bmod N) = (E, \psi)]$$

(see Section A.1). Let $f = f_0 + f_1$ with $f \in L^2_0(\mathcal{G}_{\mathcal{F}})$ and $f_1 \in L^2_{\deg}(\mathcal{G}_{\mathcal{F}})$ be the orthogonal decomposition of f. Then f_1 is proportional to the function that takes the value 1 on all (E, ψ) with $\deg\psi = 1 \in (\mathbf{Z}/N\,\mathbf{Z})^\times$ and 0 elsewhere, and

$$\frac{1}{\#\,\mathrm{Aut}(E,\psi)}\left(\frac{A^*_\ell}{\ell+1}\right)^k f_1(E, \psi) = \Pr{}_\nu[(E, \psi)].$$

The statistical distance in the statement is therefore

$$\begin{aligned}
&\sum_{(E,\psi)} \frac{1}{\#\operatorname{Aut}(E,\psi)} \left| \left(\frac{A_\ell^*}{\ell+1}\right)^k f_0(E,\psi) \right| \\
&\leq \left(\sum_{(E,\psi)} \frac{1}{\#\operatorname{Aut}(E,\psi)} \right)^{1/2} \left\| \left(\frac{A_\ell^*}{\ell+1}\right)^k f_0 \right\| \quad \text{(Cauchy–Schwarz)} \\
&\leq \left(N^4 \frac{p-1}{24} \right)^{1/2} \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^k \|f\| \\
&\leq N^2 \sqrt{\frac{p}{24}} \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^k,
\end{aligned}$$

as claimed. □

Corollary A.2. *Keep the notations of Proposition A.1. There exists a bound* $n = O(\log_\ell(pN) - \log_\ell(\varepsilon))$ *such that for all* E *and all* $k \geq n$*, conditional on* $E_\varphi = E$*, the distribution of* $\varphi \bmod N$ *is* ε*-close to uniform among isogenies of degree* $\ell^k \bmod N$*.*

A.7 Computation of End(E): the obvious $\tilde{O}(p^{1/2})$ algorithm works

Theorem 8.7 states that one can compute $\operatorname{End}(E)$ in expected time $\tilde{O}(p^{1/2})$ unconditionally. However, if we unravel the reductions leading to this theorem, the resulting algorithm seems needlessly complicated. Here, we unconditionally show that the obvious collision-based algorithm to compute endomorphisms also yields $\operatorname{End}(E)$ in expected time $\tilde{O}(p^{1/2})$.

Algorithm 8 Finding an endomorphism by collisions

Require: A supersingular elliptic curve $E_0/\mathbf{F}_{p^2}$, a parameter k.
Ensure: An endomorphism of E_0.

```
T ← ∅ an empty hash table
while true do
    ψ ← a random walk ψ: E_0 → E of length k in the ℓ-isogeny graph
    if j(E) is the key of a recorded entry φ in T then
        return φ̂ ∘ ψ ∈ End(E_0)
    else
        Record φ in T, with key j(E)
    end if
end while
```

An easy consequence of Corollary A.2 is the following.

Proposition A.3. *There exists a bound* $n = O(\log_\ell(pN))$ *such that for all choices of parameter* $k \geq n$*, Algorithm 8 runs in expected time* $\mathrm{poly}(\ell, \log p, k)p^{1/2}$*, and the endomorphisms produced are* $O(1/p)$*-close to uniform modulo* N *among endomorphisms of degree* $\ell^{2k} \in (\mathbf{Z}/N\,\mathbf{Z})^\times$*.*

From Proposition A.3, it is easy to see (by an analysis similar to Section 5 and Section 7 but easier) that for polynomial choices of k with $\ell \in \{2, 3\}$, the endomorphisms output by Algorithm 8 generate $\mathrm{End}(E_0)$ after polynomially many calls. A variant also provides a reduction from EndRing to Isogeny alternative to the proof of Theorem 8.5.