

Rapid mixing in superspecial isogeny graphs

Maria Corte-Real Santos, Harun Kir, Olivier Taïbi, and Benjamin Wesolowski

ENS de Lyon, CNRS, UMPA, UMR 5669, Lyon, France

Abstract. We prove expansion properties of superspecial isogeny graphs in dimension $g \geq 2$, including arbitrary level structure. In particular, we provide explicit bounds for the spectral gap, and a fine characterization of obstructions created by level structures (disconnectedness and periodicity). Without the level structure (the only case which was previously studied), our bounds dramatically improve the state of the art. These isogenies in dimension greater than one have become increasingly important in isogeny-based cryptography, and previous works have already conjectured consequences of their expansion properties. We present several applications of our result: we transform conjectural predecessors into rigorous algorithms, and we give the first algorithm for computing endomorphism rings of superspecial principally polarized abelian varieties in arbitrary dimension, with complexity $\tilde{O}(p^{g-1})$.

Keywords: isogeny-based cryptography, superspecial principally polarised abelian varieties, expander graphs.

Contents

1	Introduction	1
2	Preliminaries	7
3	Superspecial isogeny graphs	15
4	Statement of main theorem	18
5	Proof of Results	22
6	Applications	48
A	A brief introduction to categories	64

1 Introduction

In 1990, Pizer [Piz90] proved that the graphs formed by isogenies connecting supersingular elliptic curves have the Ramanujan property: they constitute an optimal family of expander graphs, in which random walks rapidly converge to a unique stationary distribution. This motivated their introduction to cryptography when Charles, Goren and Lauter proposed to exploit this rapid mixing to build a hash function in 2006 [CLG09] (henceforth, the *CGL hash function*).

This *supersingular isogeny-based cryptography* has since evolved into one of the main branches of post-quantum cryptography, and the rapid mixing has repeatedly proved to be a critical property. Perhaps most strikingly, it makes

it quite immediate to obtain worst-case to average-case reductions for the core computational problems of the field (in stark contrast with other branches, where such reductions are typically much more complicated, or even unknown). This underlying mechanism is implicitly what motivated the design of the CGL hash function, as preimage-resistance is an average-case property, and the idea has since evolved into a well-understood network of worst-case to average-case reductions [HW25]. Progress in isogeny-based cryptography is intertwined with refinements and generalizations of these walks. Most notably, the addition of various kinds of “level structures” (where vertices of the graph carry extra information, analyzed in [BCC⁺23,CL26,PW24a,Arp24]) has led to new protocols [BCC⁺23,BF23,BFLS26,BFP⁺26] and cryptanalysis [PW24a,DFP24].

The field has long focused primarily on elliptic curves, one-dimensional objects. Early constructions using their higher-dimensional generalisation (principally polarised abelian varieties, henceforth *PPAVs*) remained anecdotal, until a sudden paradigm shift: in 2022 [CD23,MMP⁺23,Rob23], a new algorithm to interpolate isogenies was discovered which made essential use of higher-dimensional abelian varieties. It was first developed as an attack against one of the most well-established isogeny-based protocols, SIDH/SIKE [FJP14], and is now a key component of almost all new isogeny-based constructions.

While these abelian varieties are becoming ubiquitous, they remain much less understood than elliptic curves. The structure and mixing properties of the associated isogeny graphs, which have proved crucial in dimension one, have remained elusive. Yet, heuristic algorithms rely on such properties, conjectures have been formulated, and several applications await progress in this direction.

Previous work. While [CLG09] already hinted at a potential generalisation to higher dimensions, Takashima gives the first concrete proposal for dimension 2 in 2017 [Tak17]. While acknowledging that expansion properties are a decisive factor for the design of this hash function, this early work discusses no evidence for this property to hold in dimension 2. Subsequent work in this direction [FT19,CDS20] remained silent on the expansion issue, until Costello and Smith [CS20] explicitly posit the hypothesis that the two-dimensional isogeny graphs have the Ramanujan property (i.e., they are optimal expanders). Under this hypothesis, they derive new algorithms for the computational problems supporting the cryptosystems.

However, Jordan and Zaytman [JZ23] proved that the hypothesis is false. However, these graphs are still suspected to have excellent expansion properties. These properties are essentially measured by the so-called *spectral gap* of the graph (see Definition 2.8): a quantity between 0 and 1 such that when the gap is large, random walks converge rapidly to a stationary distribution. Florit and Smith give in [FS22] experimental evidence for a large spectral gap (albeit not optimally large); they propose a formal conjecture, including explicit bounds for $\ell = g = 2$. The first rigorous results are obtained by Aikawa, Tanaka, and Yamauchi [ATY24], who prove lower bounds on the spectral gap for the superspecial isogeny graph in dimension $g \geq 2$ (with no level structure). Their bounds

appear to be far from tight: they do not match the prediction of Florit and Smith for $\ell = g = 2$, and become vacuous as $g \rightarrow \infty$.

The case with “level structure” has received even less attention. Motivated by the security analysis of the aforementioned constructions, Decru and Reijnders [DR26] initiate a study of the particular case of “ $(\ell)^g$ -level structures”, and propose a conjecture on the corresponding spectral gap, with heuristic and experimental evidence.

Our graphs of interest. Let us briefly describe our graphs — formal definitions are deferred to Section 3. Fix a finite field $\mathbb{F}_p$, with algebraic closure $\bar{\mathbb{F}}_p$. The main objects are superspecial principally polarized abelian varieties (superspecial PPAVs) over $\bar{\mathbb{F}}_p$, of dimension $g \geq 2$. They are a generalisation to higher dimensions of supersingular elliptic curves (the case $g = 1$). In particular, there is a finite number of superspecial PPAVs over $\bar{\mathbb{F}}_p$ up to isomorphism. They are the vertices of our graphs, at least in the simplest cases.

Just like elliptic curves, PPAVs are connected by isogenies, and isogenies which “behave well” are d -isogenies (Definition 2.1), with d an integer; in the case of elliptic curves, a d -isogeny is an isogeny of degree d . We fix a prime number $\ell \neq p$, and consider the ℓ -isogeny graph $\mathcal{G}_\ell(p, g)$ where

- vertices are the superspecial PPAVs over $\bar{\mathbb{F}}_p$ (up to isomorphism), and
- edges from X to X' are ℓ -isogenies $X \rightarrow X'$ (up to equivalence).

This is the superspecial ℓ -isogeny graph of dimension g over $\bar{\mathbb{F}}_p$. Note that this is actually a *directed multigraph* (see Section 2.3). From any X , there are $k = \prod_{i=1}^g (1 + \ell^i)$ outgoing ℓ -isogenies, so the graph is k -(out-)regular.

Graphs with level structure. We actually consider more general isogeny graphs by adding level structure: each vertex is a pair (X, x) , where X is a superspecial PPAV, and x is some additional information attached to X . The type of information attached is determined by some functor $\mathcal{F}$, and we write $\mathcal{G}_\ell^\mathcal{F}$ for the ℓ -isogeny graph with $\mathcal{F}$ -level structure (see Definition 3.5). This approach is highly flexible, and generalises [PW24a] to higher dimensions.

Let us now illustrate what this “ $\mathcal{F}$ ” means with an example. Given any integer N , and any prime $\ell \nmid pN$, one can define a graph $\mathcal{G}_\ell^{\text{TorPoints}_N}$ whose vertices are all pairs of the form (X, P) where $P \in X[N]$ is an N -torsion point of X . Edges connecting these pairs are ℓ -isogenies which preserve this additional information: there is an edge from (X, P) to (X', P') if and only if there is an ℓ -isogeny $\varphi : X \rightarrow X'$ such that $\varphi(P) = P'$. By walking in these graphs, one follows isogenies while keeping track of a torsion point.

Formally, TorPoints_N is a functor which to any PPAV associates the set $\text{TorPoints}_N(X) = X[N]$ of torsion points: this set $\text{TorPoints}_N(X)$ is the set of possible choices for the additional information x to attach to X , to form a vertex (X, x) . One can choose other functors, like End/N (then, x is an endomorphism of X modulo N), or MaxIso_N (then, x is the kernel of an N -isogeny from X).

These functors are all “ N -well-behaved” (for some integer N) in the sense of Definition 3.8, and any such functor $\mathcal{F}$ gives rise to a graph $\mathcal{G}_\ell^\mathcal{F}$ for any prime $\ell \nmid pN$. The functor $\text{Zero}(X) = \{0\}$ recovers the standard graph $\mathcal{G}_\ell^{\text{Zero}} = \mathcal{G}_\ell(p, g)$.

For any vertex $v = (X, x)$ in the graph, we write $\text{Aut}(v)$ its group of automorphisms. These automorphisms are isomorphisms $\alpha : X \rightarrow X$ of PPAVs which “preserve” the level structure; in the case of $\text{TorPoints}_N(X)$, this means that $\alpha \in \text{Aut}(X, P)$ satisfies $\alpha(P) = P$.

Main contribution: rapid mixing. Our main contribution is a proof that isogeny graphs of superspecial principally polarized abelian varieties, with or without level structure, have good expansion properties. In the simplest cases (i.e., when the level structure does not introduce certain obstructions), we obtain the following result which is proven in Section 4.1.

Theorem 1.1 (Rapid mixing in the aperiodic, connected case). *Fix any primes $p \neq \ell$, any integer $g \geq 2$, and any $\epsilon > 0$. Let $\mathcal{F}$ be any level structure such that $\mathcal{G}_\ell^\mathcal{F}$ is aperiodic and connected (for instance, this is the case for $\mathcal{G}_\ell^{\text{Zero}} = \mathcal{G}_\ell(p, g)$). Write V the set of vertices of the graph. For any vertex $v \in V$, let $\mu(v) = \#\text{Aut}(v)^{-1}$, and $\mu(V) = \sum_{v \in V} \mu(v)$. Then, for any starting vertex v_0 , the endpoint of a uniform random walk in $\mathcal{G}_\ell^\mathcal{F}$ of length at least*

$$n \geq \frac{\frac{1}{2} \log(\mu(V)/\mu(v_0)) - \log(2\epsilon)}{\log(\ell^{g/2} + \ell^{-g/2}) - \log(2)} = O\left(\frac{g^2 \log(p) + g^{O(1)} + \log(\#\mathcal{F}(v_0)) - \log(\epsilon)}{g \log(\ell)}\right).$$

is at statistical distance at most ϵ from the unique stationary distribution $\frac{\mu}{\mu(V)}$. The hidden constant in the O is absolute. More precisely,

$$\log(\mu(V)/\mu(v_0)) \leq \frac{g(g+1)}{2} \log(p) + \log \Phi(g) + \log(\#\mathcal{F}(v_0)),$$

with $\Phi(g) = (15.46g)^{2g} \prod_{i=1}^g \frac{|B_{2i}|}{4i} \leq e^{\text{poly}(g)}$, where B_{2i} is the $2i$ -th Bernoulli number.

We obtain similar bounds for any of the isogeny graphs $\mathcal{G}_\ell^\mathcal{F}$ with $\mathcal{F}$ -level structure, but the general statement is more delicate because the level structure can make the graph disconnected or periodic. Our most precise and general statements are expressed as a spectral analysis of the graph.

The *spectral gap* of the graph is defined as follows. Let $A_\mathcal{G}$ be the adjacency matrix of the graph $\mathcal{G}$ (see Equation (3)). Isogeny graphs are directed (out-)regular graphs, and the degree of regularity, say d , is the largest eigenvalue of $A_\mathcal{G}$ in absolute value. There may be other eigenvalues of absolute value d (although this only happens when level structures are involved), and we call them all the *trivial eigenvalues*. We write $\lambda(\mathcal{G}) = \frac{\max |\lambda|}{d}$, where the maximum is over non-trivial eigenvalues of $A_\mathcal{G}$, and normalize it as $\lambda^{\text{nr}}(\mathcal{G}) = \frac{\lambda(\mathcal{G})}{d}$. The spectral gap $\text{gap}(\mathcal{G}) = 1 - \lambda^{\text{nr}}(\mathcal{G})$ measures the (normalised) distance between the trivial and non-trivial eigenvalues. A large spectral gap (equivalently, a small $\lambda^{\text{nr}}(\mathcal{G})$)

means that the graph has good expansion properties, and random walks in the graph converge rapidly to a stationary distribution (or to a cycle of stationary distributions).

Our main contribution is the proof of the following theorem.

Theorem 1.2 (Informal summary of Theorem 4.1). *Let $\mathcal{G}$ be the ℓ -isogeny graph of superspecial principally polarized abelian varieties of dimension $g \geq 2$ over $\overline{\mathbb{F}}_p$ (possibly with some level structure $\mathcal{F}$). Then, we have $\lambda^{\text{nr}}(\mathcal{G}) \leq \left(\frac{2\sqrt{\ell}}{\ell+1}\right)^g$.*

Conditional on the “weighted fundamental lemma” in the theory of automorphic forms (see Remark 5.35) the bound improves to $\lambda^{\text{nr}}(\mathcal{G}) \leq \left(\frac{2}{\ell g/2 + \ell - g/2}\right)$.

These bounds do not depend on the level structure, but the number of trivial eigenvalues and their arguments do. Proposition 4.2 describes the relation between the level structure and the trivial eigenvalues, and Section 2.4 details their graph-theoretic consequences. These trivial eigenvalues govern the number of connected components, and the periodicity of the graph; our results characterize these graph properties from algebraic properties of the level structures.

Our strategy to prove this theorem is to relate the adjacency operator of the isogeny graph to certain Hecke operators acting on automorphic forms for the group GSp_{2g} . Using this correspondence, we apply deep results from the theory of automorphic representations, in particular, we use Oh’s theorem and Arthur’s endoscopic classification of automorphic representations of symplectic groups, to obtain an upper bound on the second largest (in absolute value) eigenvalue of the Hecke operator (and hence of the adjacency operator).

Comparison with previous work. Theorem 4.1 is the first result of this form which allows for level structures. In the case where there is no level structure, we can compare our bounds with the result of [ATY24] (also obtained via the theory of automorphic representations):

$$\lambda^{\text{nr}}(\mathcal{G}_\ell(p, g)) \leq \lambda_{\text{ATY}}^{\text{nr}}(g, \ell) := 1 - \frac{1}{4(g+2)} \cdot \left(\frac{\ell-1}{2(\ell-1) + 3\sqrt{2\ell(\ell+1)}} \right)^2.$$

Our bounds are significantly better than $\lambda_{\text{ATY}}^{\text{nr}}$ for all g and ℓ (see Table 1). Indeed, for any ℓ and g , we have

$$\lambda_{\text{ATY}}^{\text{nr}}(g, \ell) \geq \lim_{\ell \rightarrow \infty} \lambda_{\text{ATY}}^{\text{nr}}(2, \ell) \geq 0.998.$$

In particular, this bound always stays close to the trivial bound $\lambda^{\text{nr}}(\mathcal{G}) \leq 1$. Then, it degrades as the dimension grows, since $\lim_{g \rightarrow \infty} \lambda_{\text{ATY}}^{\text{nr}}(g, \ell) = 1$ for all ℓ .

In contrast, the largest value of our bounds is 0.888, occurring for $\ell = g = 2$ (or 0.8 conditionally), and they converge to the optimal limit in both g and ℓ :

$$\lim_{g \rightarrow \infty} \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^g = 0 \text{ for all } \ell, \text{ and } \lim_{\ell \rightarrow \infty} \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^g = 0 \text{ for all } g.$$

In terms of random walks, our convergence to 0 as $\ell \rightarrow \infty$ means that a single random ℓ -isogeny already gets us arbitrarily close to the stationary distribution, if ℓ is large enough (such single-step walks matter in applications: for instance, in dimension 1, they are used for the key generation of SQIsign [DKL⁺20]). In contrast, with the bound of $\lambda_{\text{ATY}}^{\text{nr}}$, a walk needs many steps to converge to the stationary distribution even if ℓ is extremely large.

A practical consequence: a 500× improvement for cryptographic constructions. The article [KMM⁺25] presents a generalization of the CGL hash function to dimensions $g \leq 3$. The main task in the hash function is the computation of a walk in the graph $\mathcal{G}_2(p, g)$. For security, they require the walks to be long enough for the target to be statistically close to the uniform distribution. Based on the bounds of [ATY24], they estimate in [KMM⁺25, Remark 26] the length of the walk to be at least $\approx 2^{18}$. As they do not state which statistical distance and which prime length this estimate uses, we present a detailed comparison in Example 4.6 targeting 128-bit security: statistical distance 2^{-128} , and $\log_2 p \approx \frac{128}{g-1}$ for $g \geq 2$. Then, using the bound $\lambda^{\text{nr}}(\mathcal{G}) \leq \left(\frac{2}{\ell^{g/2} + \ell^{-g/2}} \right)$:

1. For $g = 2$, walks have length 698 (while [ATY24] leads to 382610 steps);
2. For $g = 3$, walks have length 383 (while [ATY24] leads to 545459 steps).

New algorithms for isogeny and endomorphism problems. In Section 6, we present several algorithmic applications of rapid mixing.

1. *Splitting varieties (Theorem 6.9).* Following a heuristic strategy of Costello and Smith [CS20], we propose a rigorous algorithm which given any superspecial PPAV X , finds an isogeny from X to a product of elliptic curves, and runs in time $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$ (the exponent $\text{poly}(g)$ depends on the model in which the PPAV is represented — see Remark 1.3). We modify the original strategy to avoid failure cases. Our analysis combines the new spectral bounds with a computation of the mass of the graph, Lemma 6.5.
2. *The superspecial isogeny problem (Theorem 6.11).* The above splitting strategy from [CS20] was originally presented to resolve the superspecial ℓ -isogeny path problem: given two superspecial PPAVs, find an ℓ^n -isogeny between them. As an application of Theorem 6.9, we deduce the first rigorous algorithm to resolve this problem in time $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$ (note that the cofactor $(\ell + \log p)^{\text{poly}(g)}$ was hidden by a $\tilde{O}$ -notation in the heuristic algorithm of [CS20]).
3. *The superspecial endomorphism ring problem (Theorem 6.15).* We present the first algorithm for the superspecial endomorphism ring problem in any dimension, given a superspecial PPAV X , find a basis of its endomorphism ring, with complexity $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$. Our algorithm takes inspiration from the heuristic algorithm for $g = 2$ proposed in [ABI⁺25, Corollary 5.20].

Remark 1.3. All these algorithms are generic, in the sense that they do not require the input PPAVs to be represented in any particular way: one can use

any “efficient representation” of X , in the sense of [CPW26]. The final explicit complexity may depend on the choice of representation; for instance the aforementioned complexities $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$ become $(\ell + \log p)^{O(g^2)} \cdot p^{g-1}$ in the level-4 theta model.

We expect further applications as varied as in the dimension-1 case, constructive and cryptanalytic, and exploiting the full machinery of level structures.

Acknowledgements. The first and last authors were supported by the European Research Council under grant No. 101116169 (AGATHA CRYPTY). The second and last authors were supported by the Agence Nationale de la Recherche under grant ANR-22-PETQ-0008 (PQ-TLS). The third author was supported by the Agence Nationale de la Recherche under grant ANR-25-CE40-4664 (PPAL). Generative AI was not used for any aspect of this work.

2 Preliminaries

2.1 Basic notions and notation

Rings and Fields. We write $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ and $\mathbb{C}$ for the ring of integers, the fields of rational, real, and complex numbers. For any prime ℓ , $\mathbb{Z}_\ell$ and $\mathbb{Q}_\ell$ denote the ring of ℓ -adic integers and the field of ℓ -adic numbers. For any prime power q , $\mathbb{F}_q$ is the finite field with q elements. For any field k , $\bar{k}$ is its algebraic closure. For any set S , $\#S$ denotes its cardinality. We use R for a commutative ring with unity. *Asymptotic notation.* We write $f(x) = O(g(x))$ if there exist positive constants C and x_0 such that $|f(x)| \leq Cg(x)$, for all $x \geq x_0$, and equivalently $f(x) = \Omega(g(x))$. We also use the soft- O notation $\tilde{O}(g) = \log(g)^{O(1)}O(g)$. We have the convention that $\log(0) = -\infty$.

Probabilities. Given two probability distributions $P, Q : V \rightarrow [0, 1]$ on a discrete set V , their *statistical distance* (also known as the *total variation distance*) is $\frac{1}{2}\|P - Q\|_1 = \frac{1}{2} \sum_{v \in V} |P(v) - Q(v)| = \sup_{S \subseteq V} |P(S) - Q(S)|$.

2.2 Principally polarized abelian varieties

This subsection fixes notation for principally polarized abelian varieties and polarized isogenies: the objects and the morphisms out of which our graphs are built. General references are Milne [Mil86] and Mumford [Mum70].

Let p be a prime number and $k = \bar{\mathbb{F}}_p$ be the algebraic closure of the finite field $\mathbb{F}_p$. Let $\underline{X}/k$ be an abelian variety over k of dimension g , and let $\underline{X}^\vee = \text{Pic}^0(\underline{X})$ denote the dual abelian variety of $\underline{X}$ (as in [Mum70, p.125]).

An abelian variety $\underline{X}/k$ of dimension g is called *supersingular* if there exists an isogeny $\underline{X} \sim E^g = E \times \cdots \times E$, where E/k is a supersingular elliptic curve, and it is called *superspecial* if $\underline{X} \simeq E^g$, for some supersingular elliptic curve E/k ; see §5.1 for the alternative definitions.

We refer to [Mum70, p. 60] for the notion of a *polarization* $\lambda : \underline{X} \rightarrow \underline{X}^\vee$. A polarization is called *principal* if it is an isomorphism. An abelian variety $\underline{X}$

together with a principal polarization λ is called a *principally polarized abelian variety*, abbreviated *PPAV*, and denoted by $X := (\underline{X}, \lambda)$.

Definition 2.1. A (polarized) isogeny between PPAVs $X_1 = (\underline{X}_1, \lambda_1)$ and $X_2 = (\underline{X}_2, \lambda_2)$ of dimension g is an isogeny $f : \underline{X}_1 \rightarrow \underline{X}_2$ such that $f^* \lambda_2 := f^\vee \lambda_1 f = d \lambda_1$, where d denotes the multiplication-by- d map for some integer $d > 0$, and $f^\vee : \underline{X}_2^\vee \rightarrow \underline{X}_1^\vee$ denotes the dual isogeny. In this case, we have $\deg(f) = d^g$. We call $\text{rdeg}(f) := d$ the reduced degree of f , and f is called a *d-isogeny*.

Superspecial PPAVs. If $X = (\underline{X}, \lambda)$ is a PPAV, where $\underline{X}$ is a superspecial (supersingular) abelian variety, we call X a *superspecial (supersingular) PPAV*. Superspecial PPAVs are the objects of the following category, which underlies the superspecial isogeny graph defined in §3.

Definition 2.2. Let $N \in \mathbb{Z} > 0$. We let $\mathbf{SS}_N(p, g)$ denote the category with

- objects: superspecial PPAVs of dimension g over $\mathbb{F}_p$;
- morphisms in $\text{Hom}_N(X, Y)$: polarised isogenies of degree coprime to Np .

We also define the category $\mathbf{SS}(p, g) := \mathbf{SS}_1(p, g)$.

Adjoint isogeny. For any two PPAVs $X = (\underline{X}, \lambda)$ and $X' = (\underline{X}', \lambda')$, if $f : X \rightarrow X'$ is a polarized isogeny of (reduced) degree d , then we have the *adjoint* isogeny $f^\dagger := \lambda^{-1} f^\vee \lambda' : X' \rightarrow X$ such that $f f^\dagger = d = f^\dagger f$. If $X = (\underline{X}, \lambda)$ is a PPAV, then the adjoint operator $f^\dagger = \lambda^{-1} f^\vee \lambda$ is an involution of $\text{End}(\underline{X})$, called the *Rosati involution* with respect to λ .

Maximal isotropic subgroups. Let $X = (\underline{X}, \lambda)$ be a PPAV and $\underline{X}'$ be an abelian variety. Suppose that we have an isogeny $f : \underline{X} \rightarrow \underline{X}'$. In general, there does not exist a principal polarization of $\underline{X}'$ such that f is indeed a polarized isogeny. We can give a criterion for the existence of such an isogeny in terms of the maximal isotropic subgroups. To define these groups, let us recall the Weil pairing first. Throughout this subsection, d denotes a positive integer.

Remark 2.3. Let $X = (\underline{X}, \lambda)/k$ be a PPAV and let d be a prime to p . The polarization λ induces a nondegenerate alternating bilinear pairing $e_d^\lambda : \underline{X}[d] \times \underline{X}[d] \rightarrow \mu_d$, called the *d-Weil pairing* (or simply *pairing*), where $\underline{X}[d]$ is the kernel of the multiplication-by- d map and μ_d is the group of d -th roots of unity in $k^\times$; see [Mum70, p.228]. We write $X[d]$ for the pair $(\underline{X}[d], e_d^\lambda)$.

Definition 2.4. A subgroup scheme $G \subseteq \underline{X}[d]$ of a PPAV $(\underline{X}, \lambda)$ is called *isotropic with respect to the pairing e_d^λ* if $e_d^\lambda(P, Q) = 1$ for all $P, Q \in G(R)$ (where R is any k -algebra), and *maximally isotropic* if it is not properly contained in any isotropic subgroup of $\underline{X}[d]$.

Note that an isotropic subgroup of $\underline{X}[d]$ is maximal if and only if it is of order d^g . If in addition $d = \ell$ is a prime not equal to p , G has the structure $G \simeq (\mathbb{Z}/\ell\mathbb{Z})^g$ as an abelian group scheme. Here is the criterion for the existence of a polarized isogeny by [Kan97, Proposition 1.1] (see also [Mum70, §23]):

Proposition 2.5. *Let $X = (\underline{X}, \lambda)$ be a PPAV over k . Let $f : \underline{X} \rightarrow \underline{X}'$ be an isogeny between two abelian varieties. Then $\ker(f)$ is a maximally isotropic subgroup scheme of $\underline{X}[d]$ with respect to e_d^λ if and only if there is a (unique) principal polarization λ' on $\underline{X}'$ such that $f^*\lambda' = d\lambda$, i.e., f is a polarized isogeny.*

Definition 2.6. *Let $X = (\underline{X}, \lambda)$ be a PPAV. Two d -isogenies $f_i : X \rightarrow X_i = (\underline{X}_i, \lambda_i)$, for $i = 1, 2$, from X are called equivalent if there is an isomorphism $\phi : \underline{X}_1 \rightarrow \underline{X}_2$ with $\phi^*\lambda_2 = \lambda_1$ such that $f_2 = \phi f_1$.*

We write $\text{Isog}_d(X)$ for the set of equivalence classes $[f]$ of d -isogenies f from X , and $\text{Isot}_d(X)$ for the set of maximal isotropic subgroups of $X[d]$ (with respect to e_d^λ). By Proposition 2.5, we obtain a bijection $\text{Isog}_d(X) \xrightarrow{\sim} \text{Isot}_d(X) : [f] \mapsto \ker(f)$. If X/k is a superspecial PPAV of dimension g and $\ell \neq p$ is prime, the number of equivalence classes of ℓ -isogenies from X depends only on ℓ , g and is

$$N_g(\ell) := \# \text{Isot}_\ell(X) = \prod_{i=1}^g (\ell^i + 1), \quad (1)$$

This will be the (out)-degree of every vertex in our graphs.

2.3 Graph theory

A directed multigraph $\mathcal{G}$ is a tuple $\mathcal{G} = (V, E, s, t)$ consisting of a set V of vertices, a set of directed edges E and two maps, the *source* map, $s : E \rightarrow V$, which assigns a starting vertex to every edge and the *target* map, $t : E \rightarrow V$, which assigns an ending vertex to every edge.

Let $m(u, v)$ denote the number of edges from u to v , for vertices u, v in V . For any vertex $u \in V$, we write the (out-)degree as $\deg(u) = \sum_{v \in V} m(u, v)$, where the sum runs over all target vertices in V . If we have that $\deg(u) = d$, for any vertex $u \in V$, then the graph $\mathcal{G}$ is called *d-regular* graph.

The graph $\mathcal{G}$ is called *strongly connected* if there is a path in each direction between each pair of vertices of the graph. The gcd of the lengths of the cycles of $\mathcal{G}$ is called its *period*, and if the graph $\mathcal{G}$ has period 1, it is called *aperiodic*.

Throughout this section, $\mathcal{G}$ denotes a finite d -regular directed multigraph for some $d > 0$, unless explicitly stated otherwise.

Adjacency operator. A *Hilbert space* H is a complex inner product space that is complete with respect to the norm $\|\cdot\|$ induced by its Hermitian inner product $\langle \cdot, \cdot \rangle$. Recall that if H is finite-dimensional, it is automatically complete. Let H be a finite-dimensional Hilbert space. The *adjoint* of a linear operator $T : H \rightarrow H$ is the unique operator T^* satisfying $\langle Tv, w \rangle = \langle v, T^*w \rangle$ for all $v, w \in H$. A *normal*

operator is an operator that commutes with its adjoint. The *operator norm* of an operator T is $\max_{v \neq 0} \frac{\|Tv\|}{\|v\|}$.

For a weight function $\mu : V \rightarrow \mathbb{R}_{>0}$, we let $L^2(\mathcal{G}, \mu)$ be the finite-dimensional Hilbert space of complex-valued functions defined on the vertices of $\mathcal{G}$, equipped with the weighted inner product

$$\langle f, g \rangle_\mu = \sum_{v \in V} \mu(v) f(v) \overline{g(v)}. \quad (2)$$

The *adjacency operator* $A_{\mathcal{G}}$ on $L^2(\mathcal{G}, \mu)$ is

$$(A_{\mathcal{G}}f)(u) = \sum_{v \in V} m(u, v) f(v) = \sum_{u \rightarrow v} f(v), \quad (3)$$

where the latter sum is over all edges whose source is u .

Note that the adjacency operator $A_{\mathcal{G}}$ is not necessarily normal since the graph $\mathcal{G}$ is directed. However, the adjacency operators of the specific directed graphs studied in this work are normal with respect to the weighted inner product.

Consider a *random walk* $v_0, v_1, \dots$ on $\mathcal{G}$: we start at a vertex v_0 ; if at the n -th step we are at a vertex $v_n = u$, we move to a neighbor v with probability: $\Pr(v_{n+1} = v \mid v_n = u) = \frac{m(u, v)}{\deg(u)} = \frac{m(u, v)}{d}$.

While we primarily rely on $A_{\mathcal{G}}$ for theoretical purposes due to its close relationship with Hecke operators, we define an alternative adjacency operator $B_{\mathcal{G}}$ to track the probability distribution of a random walk:

$$(B_{\mathcal{G}}f)(v) = \sum_{u \in V} m(u, v) f(u) = \sum_{u \rightarrow v} f(u),$$

where the latter sum is over all edges whose target is v . The probability distribution of the random walk on the d -regular graph $\mathcal{G}$ at step n is given by:

$$P_n = \frac{1}{d} B_{\mathcal{G}} P_{n-1} = \left(\frac{1}{d} B_{\mathcal{G}} \right)^n P_0,$$

where $P_i(v) = \Pr(v_i = v)$ is the probability that the walk is at vertex v at step i , and P_0 is the initial distribution.

Remark 2.7. Two operators $A_{\mathcal{G}}$ and $B_{\mathcal{G}}$ are naturally related via the adjoint. More precisely, we have that $\langle A_{\mathcal{G}}f, g \rangle_\mu = \langle f, A_{\mathcal{G}}^*g \rangle_\mu$ where $A_{\mathcal{G}}^* = M^{-1}B_{\mathcal{G}}M$ and M is defined as $(Mf)(u) = \mu(u)f(u)$.

Spectral properties. We now define the spectral gap of a graph, which is a measure of how well-connected the graph is. The spectral gap is defined in terms of the eigenvalues of the adjacency operator. Since $\mathcal{G}$ is a d -regular graph, it is clear that d is an eigenvalue of the adjacency operator $A_{\mathcal{G}}$ (corresponding to the constant function). Any eigenvalue of $A_{\mathcal{G}}$ with absolute value d is called a *trivial* eigenvalue, and all other ones are called non-trivial. We might have more than one trivial eigenvalue, and this is discussed in greater detail in Section 2.4.

Definition 2.8 (Spectral gap). If $A_{\mathcal{G}}$ has no non-trivial eigenvalue, let $\lambda(\mathcal{G}) = 0$; otherwise let $\lambda(\mathcal{G}) = \max |\lambda|$, where λ runs over the non-trivial eigenvalues of $A_{\mathcal{G}}$. We set $\text{gap}(\mathcal{G}) = 1 - \frac{\lambda(\mathcal{G})}{d}$, and call $\text{gap}(\mathcal{G})$ the spectral gap.

A stationary distribution of the graph $\mathcal{G}$ is a probability distribution P on V such that $B_{\mathcal{G}}P = dP$. Let $\mu : V \rightarrow \mathbb{R}_{>0}$ be a weight function. The mass of any subset $S \subseteq V$ is $\mu(S) := \sum_{x \in S} \mu(x)$. By Remark 2.7, if the operator $A_{\mathcal{G}}$ is normal in $L^2(\mathcal{G}, \mu)$, then $B_{\mathcal{G}}$ is normal in $L^2(\mathcal{G}, \mu^{-1})$. Consider a random walk $v_0, \dots, v_n$ on $\mathcal{G}$, and denote P_0 as the probability distribution at v_0 and P_n as the probability distribution after n steps. Since $B_{\mathcal{G}}$ is normal, we can write the spectral decomposition of the P_0 as: $P_0 = \sum_{\lambda} P_{0,\lambda}$, where $P_{0,\lambda}$ is the projection of the P_0 onto the eigenspace associated with the eigenvalue λ of $B_{\mathcal{G}}$. Let us consider “periodic stationary distribution”

$$Q_n = \sum_{|\lambda|=d} \frac{\lambda^n}{d^n} P_{0,\lambda}. \quad (4)$$

Observe that if there exists a unique λ such that $|\lambda| = d$ (equivalently, if the graph $\mathcal{G}$ is strongly connected and aperiodic), then $Q_n = P_{0,d}$ is indeed the (unique) stationary distribution of the graph. The aim is now to find a bound on the statistical distance between P_n and Q_n .

Proposition 2.9. We have the following bound on the statistical distance:

$$\frac{1}{2} \|P_n - Q_n\|_1 \leq \frac{1}{2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n C \quad (5)$$

where $C = \mu(V)^{\frac{1}{2}} \cdot \mu(v_0)^{-\frac{1}{2}}$. Therefore, we obtain that $\frac{1}{2} \|P_n - Q_n\|_1 \leq \epsilon$, for

$$n \geq \frac{\log(C) - \log(2\epsilon)}{\log(d) - \log(\lambda(\mathcal{G}))}. \quad (6)$$

Proof. For technical reasons, we first suppose that $\lambda(\mathcal{G}) \neq 0$. The norm induced by the inner product is $\|f\| = \langle f, f \rangle_{\mu^{-1}}^{\frac{1}{2}} = (\sum_{x \in V} \mu^{-1}(x) |f(x)|^2)^{\frac{1}{2}}$. Since $B_{\mathcal{G}}$ is normal in $L^2(\mathcal{G}, \mu^{-1})$, we can obtain by the spectral theorem that

$$\|P_n - Q_n\| \leq \left(\frac{\lambda(\mathcal{G})}{d} \right)^n \|P_0 - Q_n\|. \quad (7)$$

Indeed, since we have $P_n = d^{-n} B_{\mathcal{G}}^n P_0 = d^{-n} B_{\mathcal{G}}^n (\sum_{\lambda} P_{0,\lambda}) = \sum_{\lambda} \left(\frac{\lambda}{d} \right)^n P_{0,\lambda}$, it follows that $P_n - Q_n = \sum_{|\lambda| < d} \left(\frac{\lambda}{d} \right)^n P_{0,\lambda}$. Note that since $\lambda(\mathcal{G}) \neq 0$, there exists at least one λ with $|\lambda| < d$. By the spectral theorem,

$$\|P_n - Q_n\|^2 = \left\| \sum_{|\lambda| < d} \left(\frac{\lambda}{d} \right)^n P_{0,\lambda} \right\|^2 = \sum_{|\lambda| < d} \left| \frac{\lambda}{d} \right|^{2n} \|P_{0,\lambda}\|^2 \leq \left(\frac{\lambda(\mathcal{G})}{d} \right)^{2n} \sum_{|\lambda| < d} \|P_{0,\lambda}\|^2.$$

Since $P_0 - Q_n = \sum_{|\lambda| < d} P_{0,\lambda}$, we see that

$$\sum_{|\lambda| < d} \|P_{0,\lambda}\|^2 = \left\| \sum_{|\lambda| < d} P_{0,\lambda} \right\|^2 = \|P_0 - Q_n\|^2,$$

and so,

$$\|P_n - Q_n\| \leq \left(\frac{\lambda(\mathcal{G})}{d} \right)^n \|P_0 - Q_n\|.$$

We now note that:

$$\frac{1}{2} \|P_n - Q_n\|_1 = \frac{1}{2} \sum_{x \in V} |P_n(x) - Q_n(x)| = \frac{1}{2} \sum_{x \in V} \left(|P_n(x) - Q_n(x)| \cdot \mu^{-1}(x)^{\frac{1}{2}} \right) \cdot \mu(x)^{\frac{1}{2}}.$$

By the Cauchy-Schwarz inequality, we have

$$\begin{aligned} \frac{1}{2} \sum_{x \in V} \left(|P_n(x) - Q_n(x)| \cdot \mu^{-1}(x)^{\frac{1}{2}} \right) \cdot \mu(x)^{\frac{1}{2}} &\leq \frac{1}{2} \left(\sum_{x \in V} |P_n(x) - Q_n(x)|^2 \mu^{-1}(x) \right)^{\frac{1}{2}} \mu(V)^{\frac{1}{2}} \\ &= \frac{1}{2} \|P_n - Q_n\| \mu(V)^{\frac{1}{2}}. \end{aligned}$$

Also, it is easy to see that $\|P_0 - Q_n\| \leq \|P_0\| = \mu^{-1}(v_0)^{\frac{1}{2}}$, and so, it follows by combining the inequalities that

$$\frac{1}{2} \|P_n - Q_n\|_1 \leq \frac{1}{2} \|P_n - Q_n\| \mu(V)^{\frac{1}{2}} \leq \frac{1}{2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n \|P_0 - Q_n\| \mu(V)^{\frac{1}{2}} \leq \frac{C}{2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n,$$

which proves (5). If $\frac{C}{2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n \leq \epsilon$, then we see that $n \log \left(\frac{\lambda(\mathcal{G})}{d} \right) \leq \log \left(\frac{2\epsilon}{C} \right)$.

Since $\lambda(\mathcal{G}) < d$, the value $\log \left(\frac{\lambda(\mathcal{G})}{d} \right) < 0$, and so

$$n \geq \frac{\log \left(\frac{2\epsilon}{C} \right)}{\log \left(\frac{\lambda(\mathcal{G})}{d} \right)} = \frac{\log(C) - \log(2\epsilon)}{\log(d) - \log(\lambda(\mathcal{G}))},$$

which proves (6). We now complete the proof by briefly discussing the easier case $\lambda(\mathcal{G}) = 0$. In this case, (5) is clear because we have that $P_n = Q_n$. The equation (6) also follows in this case by the convention that $\log(0) = -\infty$. $\square$

Probability of landing in a prescribed subset. We finish this section by presenting the following lemma, which bounds the probability of landing in a prescribed subset S of vertices after n steps, i.e., $\sum_{x \in S} P_n(x)$. This will be useful for our applications in Section 6.

Lemma 2.10. *Let $S \subseteq V$ be a subset of the vertex set. Assume that the only eigenvalue of magnitude d is d itself with multiplicity one, so that the unique stationary distribution is $Q(x) = \frac{\mu(x)}{\mu(V)}$. Then,*

$$\left| \sum_{x \in S} P_n(x) - \frac{\mu(S)}{\mu(V)} \right| \leq \mu(v_0)^{-1/2} \mu(S)^{1/2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n. \quad (8)$$

Proof. This comes from the standard arguments used in the proof of Proposition 2.9. Let 1_S denote the indicator function of the set S , i.e., $1_S(x) = 1$ if $x \in S$ and $1_S(x) = 0$ otherwise. Then we write the probability of landing in S (in terms of μ^{-1} inner product) as

$$\sum_{x \in S} P_n(x) = \sum_{x \in V} P_n(x) (\mu(x) 1_S(x)) \mu^{-1}(x) = \langle P_n, \mu 1_S \rangle.$$

Similarly, we compute that

$$\frac{\mu(S)}{\mu(V)} = \sum_{x \in S} Q_n(x) = \langle Q_n, \mu 1_S \rangle.$$

Using both equations we have that

$$\left| \sum_{x \in S} P_n(x) - \frac{\mu(S)}{\mu(V)} \right| = |\langle P_n - Q_n, \mu 1_S \rangle|.$$

It follows from the Cauchy-Schwarz inequality that

$$|\langle P_n - Q_n, \mu 1_S \rangle| \leq \|P_n - Q_n\| \|\mu 1_S\|.$$

We have already obtained in the proof of Proposition 2.9 that $\|P_n - Q_n\| \leq \left(\frac{\lambda(\mathcal{G})}{d} \right)^n \mu(v_0)^{-1/2}$. Furthermore, we easily compute the norm of $\mu 1_S$:

$$\|\mu 1_S\| = \left(\sum_{x \in V} |\mu(x) 1_S(x)|^2 \mu^{-1}(x) \right)^{1/2} = \left(\sum_{x \in S} \mu(x) \right)^{1/2} = \mu(S)^{1/2}.$$

Combining these inequalities yields the desired bound. $\square$

2.4 Graph-theoretic interpretation of trivial eigenvalues

Recall that the trivial eigenvalues of $\mathcal{G}$ are the eigenvalues of $A_{\mathcal{G}}$ with absolute value d (e.g., d itself, the eigenvalue if the constant function). We might have more than one trivial eigenvalue, and we now discuss how these arise.

We say that a vertex v is *reachable* from a vertex u if there exists a path $u \rightarrow \cdots \rightarrow v$ (possibly of length zero). We assume for the rest of this section that the relation “is reachable from” is symmetric, so that it is an equivalence relation on the set of vertices of $\mathcal{G}$. (This will be the case for all the isogeny graphs considered in this paper.) We call the equivalence classes connected components of the graph, and implicitly consider them as induced subgraphs. Our assumption implies that they are strongly connected.

Lemma 2.11. *Assume that $\mathcal{G}$ is non-empty and connected. Consider, for a vertex u of $\mathcal{G}$, the gcd $M(u)$ of the lengths of the cycles $u \rightarrow \cdots \rightarrow u$.*

1. *This integer $M(u)$ does not depend on the choice of u . We denote it by $M(\mathcal{G})$.*
2. *For any vertices u, v of $\mathcal{G}$ and any two paths from u to v of lengths k_1 and k_2 we have $k_1 \equiv k_2 \pmod{M(\mathcal{G})}$.*

Proof. Let u, v be vertices of $\mathcal{G}$. Let k_1, k_2 be lengths of paths from u to v . There exists a path from v to u , say of length ℓ . By concatenating paths we find that $M(u)$ divides both $k_1 + \ell$ and $k_2 + \ell$. In particular, $k_1 \equiv k_2 \pmod{M(u)}$. Any cycle from v to v , say of length ℓ , can be inserted in such a concatenation, and we deduce that $M(u)$ divides ℓ , so $M(u)$ divides $M(v)$. By symmetry we have $M(u) = M(v)$. $\square$

Definition 2.12. *Recall that we assume that the graph $\mathcal{G} = (V, E, s, t)$ is d -regular, where $d > 0$, and that the relation “is reachable from” is symmetric. Write $\mathcal{G}$ as the disjoint union of connected components $\mathcal{G}_i$. Define $M(\mathcal{G}_i)$ as in Lemma 2.11. Consider the equivalence relation $\sim$ on $\mathcal{G}$ defined by $u \sim v$ if u and v lie in the same connected component, say $\mathcal{G}_i$, and the length of a (equivalently, any) path from u to v is divisible by $M(\mathcal{G}_i)$.*

We define $\mathcal{G}_{\text{triv}}$ to be the directed graph whose set of vertices is $V/\sim$ and set of edges the set of pairs $([u], [v]) \in (V/\sim)^2$ where u and v lie in the same connected component, say $\mathcal{G}_i$, and there exists a path from u to v of length $\equiv 1 \pmod{M(\mathcal{G}_i)}$. We define $\pi : \mathcal{G} \rightarrow \mathcal{G}_{\text{triv}}$ to be the obvious morphism of graphs.

Note that the decomposition of $\mathcal{G}_{\text{triv}}$ into connected components is $\mathcal{G}_{\text{triv}} = \bigsqcup_i \mathcal{G}_{i, \text{triv}}$ and that each $\mathcal{G}_{i, \text{triv}}$ is cyclic with $M(\mathcal{G}_i)$ vertices. The operator $A_{\mathcal{G}_{i, \text{triv}}}$ on $L^2(\mathcal{G}_{i, \text{triv}})$ has eigenvalues the $M(\mathcal{G}_i)$ -th roots of unity, each with multiplicity one.

Proposition 2.13. *The morphism of graphs $\pi : \mathcal{G} \rightarrow \mathcal{G}_{\text{triv}}$ induces an embedding $\pi^* : L^2(\mathcal{G}_{\text{triv}}) \rightarrow L^2(\mathcal{G}, \mu)$ with image the direct sum of the eigenspaces for $A_{\mathcal{G}}$ corresponding to trivial eigenvalues. Moreover we have $A_{\mathcal{G}} \circ \pi^* = d \times (\pi^* \circ A_{\mathcal{G}_{\text{triv}}})$. In particular, any trivial eigenvalue of $A_{\mathcal{G}}$ is of the form ζd where ζ is a $M(\mathcal{G}_i)$ -th root of unity for some index i .*

Proof. The morphism π is surjective on vertices so π^* is an embedding. The relation $A_{\mathcal{G}} \circ \pi^* = (d \times \pi^*) \circ A_{\mathcal{G}_{\text{triv}}}$ is clear, and since $A_{\mathcal{G}_{i, \text{triv}}}$ is semi-simple with eigenvalues the $M(\mathcal{G}_i)$ -th roots of unity, the image of π^* is included in the direct sum of the eigenspaces for $A_{\mathcal{G}}$ corresponding to trivial eigenvalues.

Suppose ζd is an eigenvalue of $A_{\mathcal{G}}$ with $|\zeta| \geq 1$, and let $f \in L^2(\mathcal{G}, \mu)$ be an associated eigenvector. Let $\mathcal{G}_i$ be a connected component on which f does not vanish identically. Choose a vertex $u \in V$ that maximizes $|f|$ on $\mathcal{G}_i$, in particular $f(u) \neq 0$. By this maximality, taking the absolute value on both sides of $A_{\mathcal{G}}f(u) = \zeta df(u)$ implies

$$d|\zeta||f(u)| = |A_{\mathcal{G}}f(u)| = \left| \sum_{u \rightarrow v} f(v) \right| \leq \sum_{u \rightarrow v} |f(v)| \leq \sum_{u \rightarrow v} |f(u)| = d|f(u)|,$$

and so we have $|\zeta| = 1$ and equality must hold throughout. The condition $\sum |f(v)| = d|f(u)|$ forces $|f(v)| = |f(u)|$ for all target vertices v with source u , and thus, $f(v) = \zeta f(u)$ for every edge $u \rightarrow v$, in particular $|f|$ is also maximal at v . Applying this argument inductively along the edges of $\mathcal{G}$, we find that if a vertex w is reachable from u via a path of length k , it follows that $f(w) = \zeta^k f(u)$. Taking $w = u$ we find that ζ is a $M(\mathcal{G}_i)$ -th root of unity. The fact that f factors through π follows easily. $\square$

Proposition 2.14. *The properties in Proposition 2.13 characterize $\pi : \mathcal{G} \rightarrow \mathcal{G}_{\text{triv}}$: if $\pi' : \mathcal{G} \rightarrow \mathcal{G}'$ is a morphism of directed graphs which is surjective on vertices, with $\mathcal{G}'$ a disjoint union of (directed) cycles, satisfying $A_{\mathcal{G}} \circ (\pi')^* = d \times ((\pi')^* \circ A_{\mathcal{G}'})$ and such that the image of $(\pi')^* : L^2(\mathcal{G}') \rightarrow L^2(\mathcal{G}, \mu)$ is the direct sum of all eigenspaces for $A_{\mathcal{G}}$ corresponding to trivial eigenvalues, then π' is isomorphic to π .*

Proof. We can reduce to the case where $\mathcal{G}'$ is just one cycle. Then $\mathcal{G}$ is connected: otherwise the image of $(\pi')^*$ would not contain the whole eigenspace for d (e.g. locally constant functions which vanish on all but one connected component). Given two vertices u, v of $\mathcal{G}$ and two paths from u to v (in $\mathcal{G}$) of length k_1 and k_2 then applying π' we get $k_1 \equiv k_2 \pmod{M(\mathcal{G})}$. We deduce that $M(\mathcal{G}')$ divides $M(\mathcal{G})$ and that there is unique morphism of graphs $\rho : \mathcal{G}_{\text{triv}} \rightarrow \mathcal{G}'$ such that $\rho \circ \pi = \pi'$. This morphism ρ is clearly surjective, if it is not injective then $(\pi')^*$ cannot have the same image as π^* . Thus ρ is bijective on vertices, and since it is a morphism between directed cycles it is an isomorphism. $\square$

3 Superspecial isogeny graphs

The graphs we study have superspecial PPAVs as vertices and polarized isogenies as edges, with each vertex carrying an additional piece of information: a *level structure*. Rather than considering each kind of level structure separately, we encode a choice of extra (finite) data on a vertex X by a functor $\mathcal{F}$ from $\mathbf{SS}_N(p, g)$ to Sets, by means of the concept of the *category of elements* of $\mathcal{F}$ (see Definition 3.1), and attach a graph to each of such functor. This construction is too general to give interesting results for an arbitrary functor, and we therefore isolate a compatibility condition, namely the *(mod N)-congruence property* (see Definition 3.8). We illustrate three interesting graphs with the level structure with this property in Examples 3.2–3.4. Note that taking a functor $\mathcal{F}$ with constant value a one-point set recovers the (classical) superspecial ℓ -isogeny graph with no level structure.

3.1 Level structures via categories of elements

We consider the category $\mathbf{SS}_N(p, g)$ enriched with some additional information, such as torsion points. To this end, we use the notion of a *category of elements*. For a brief overview on general categories, see Appendix A; a general reference is [ML98].

Definition 3.1. Let $\mathbf{Sets}$ be the category of sets. Let $\mathcal{C}$ be a category and $\mathcal{F} : \mathcal{C} \rightarrow \mathbf{Sets}$ be a functor. The category of elements $\mathbf{El}(\mathcal{F})$ is the category with

- objects: pairs (c, x) where $c \in \mathcal{C}$ and $x \in \mathcal{F}(c)$;
- morphisms $(c, x) \rightarrow (c', x')$: those $f \in \mathbf{Hom}_{\mathcal{C}}(c, c')$ with $\mathcal{F}(f)(x) = x'$.

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

To illustrate this definition, we give three examples.

Example 3.2. Consider the functor $\mathbf{End}/N : \mathbf{SS}_N(p, g) \rightarrow \mathbf{Sets}$ defined by $(\mathbf{End}/N)(X) = \mathbf{End}(\underline{X})/N \mathbf{End}(\underline{X})$, and $\varphi : X \rightarrow Y$ by $(\mathbf{End}/N)(\varphi) : \alpha \mapsto \varphi \alpha \varphi^\dagger$. The category $\mathbf{El}(\mathbf{End}/N)$ is the category of pairs (X, α) , where $\alpha \in \mathbf{End}(\underline{X})/N \mathbf{End}(\underline{X})$.

Example 3.3. Consider the functor $\mathbf{TorPoints}_N : \mathbf{SS}_N(p, g) \rightarrow \mathbf{Sets}$ defined by $\mathbf{TorPoints}_N(X) = \underline{X}[N]$, and $\varphi : X \rightarrow Y$ by $\mathbf{TorPoints}_N(\varphi) : P \mapsto \varphi(P)$. The category $\mathbf{El}(\mathbf{TorPoints}_N)$ is the category of pairs (X, P) with $P \in \underline{X}[N]$.

Example 3.4. Consider the functor $\mathbf{MaxIso}_N : \mathbf{SS}_N(p, g) \rightarrow \mathbf{Sets}$ defined by $\mathbf{MaxIso}_N(X) = \{ G \subseteq \underline{X}[N] \text{ maximal isotropic} \}$, and $\varphi : X \rightarrow Y$ by $\mathbf{MaxIso}_N(\varphi) : G \mapsto \varphi(G)$. The category $\mathbf{El}(\mathbf{MaxIso}_N)$ is the category of pairs (X, G) with $G \subseteq \underline{X}[N]$ maximal isotropic.

Let $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \mathbf{Sets}$ be a functor. We say that two morphisms f_1 and f_2 in the category $\mathbf{El}(\mathcal{F})$ are *equivalent* if f_1 and f_2 are equivalent isogenies (in the sense of Definition 2.6).

Definition 3.5. Let $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \mathbf{Sets}$ be a functor with $\mathcal{F}(X)$ finite for all X . We define the graph $\mathcal{G}^{\mathcal{F}}$ with:

- (i) Vertices: isomorphism classes of objects in $\mathbf{El}(\mathcal{F})$;
- (ii) Edges: for an object (X, x) of $\mathbf{El}(\mathcal{F})$, the set of edges with source the isomorphism class of (X, x) is the set of equivalence classes of isogenies $f \in \mathbf{Hom}_N((X, x), (Y, y))$. The target of an edge $f \in \mathbf{Hom}_N((X, x), (Y, y))$ is the isomorphism class of (Y, y) .

Remark 3.6. (i) Since $\mathbf{SS}_N(p, g)$ has only finitely many isomorphism classes of objects (e.g., see [Eke87]), $\mathbf{El}(\mathcal{F})$ has only finitely many isomorphism classes of objects as well. Thus the graph $\mathcal{G}^{\mathcal{F}}$ is finite.

- (ii) In Definition 3.5, we are implicitly choosing a representative (X, x) of its isomorphism class: the set of edges depends on this choice, but different choices would give isomorphic graphs.
- (iii) The edges of $\mathcal{G}^{\mathcal{F}}$ are labeled by the reduced degree of the corresponding isogeny, which is constant on equivalence classes. For a prime ℓ , we write $\mathcal{G}_\ell^{\mathcal{F}}$ for the subgraph with the same vertices and containing only the edges labeled by ℓ . In this article, we will only consider prime $\ell \nmid pN$ (note that if $\ell \mid N$, the graph has no edge, by the definition of $\mathbf{SS}_N(p, g)$).

The number $m_\ell^{\mathcal{F}}((X, x), (Y, y))$ of edges whose (reduced) isogeny degree is ℓ from (X, x) to (Y, y) is

$$m_\ell^{\mathcal{F}}((X, x), (Y, y)) = \frac{\#\{\ell\text{-isogenies } f : X \rightarrow Y \text{ such that } \mathcal{F}(f)(x) = y\}}{\#\text{Aut}(Y, y)}. \quad (9)$$

To see this, note that $\text{Aut}(Y, y)$ acts by post-composition on the set of ℓ -isogenies $f : X \rightarrow Y$ with $\mathcal{F}(f)(x) = y$, and that the orbits of this action are exactly the edges from (X, x) to (Y, y) . Since this action is free, Equation (9) follows. Note that one may equivalently regard $\mathcal{G}_\ell^{\mathcal{F}}$ as a *weighted* graph on the same vertex set, in which the ordered pair $((X, x), (Y, y))$ has the weight $m_\ell^{\mathcal{F}}((X, x), (Y, y))$.

Definition 3.7 (The superspecial ℓ -isogeny graph). *Consider $N = 1$ and the zero functor $\text{Zero} : \mathbf{SS}(p, g) \rightarrow \mathbf{Sets}$ defined by $X \mapsto \{0\}$. The graph $\mathcal{G}_\ell^{\text{Zero}}$ is called the superspecial ℓ -isogeny graph (without level structure), denoted $\mathcal{G}_\ell(p, g)$.*

The construction of a category of elements is too general; to obtain interesting results, the additional information must be in a way “compatible” with the structure of the abelian varieties in the following sense.

Definition 3.8. *Let $\text{End}_N(X) := \text{End}_{\mathbf{SS}_N(p, g)}(X)$, for an object X in $\mathbf{SS}_N(p, g)$. Let $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \mathbf{Sets}$ be a functor. We say that $\mathcal{F}$ has the $(\text{mod } N)$ -congruence property if for every object X of $\mathbf{SS}_N(p, g)$ and any $\varphi, \psi \in \text{End}_N(X)$ such that $\varphi - \psi \in N \text{End}(\underline{X})$, we have $\mathcal{F}(\varphi) = \mathcal{F}(\psi)$.*

We now illustrate that the functors in Examples 3.2, 3.3 and 3.4 have the $(\text{mod } N)$ -congruence property.

Example 3.9 (Continuation of Example 3.2). The functor End/N has the $(\text{mod } N)$ -congruence property. Indeed, if $\varphi, \psi \in \text{End}_N(\underline{X})$ such that $\varphi - \psi = N\beta$, for some $\beta \in \text{End}(\underline{X})$, then for any $\alpha \in \text{End}(\underline{X})$ we have that $\text{End}/N(\varphi)(\alpha) = \varphi\alpha\varphi^\dagger = (N\beta + \psi)\alpha(N\beta + \psi)^\dagger = N\gamma + \psi\alpha\psi^\dagger$, for some $\gamma \in \text{End}(\underline{X})$, and so $\text{End}/N(\varphi)(\alpha) = N\gamma + \psi\alpha\psi^\dagger = \text{End}/N(\psi)(\alpha)$, so that $\text{End}/N(\varphi) = \text{End}/N(\psi)$.

We prove a lemma to give two other examples:

Lemma 3.10. *Let X be any object of $\mathbf{SS}_N(p, g)$ and let $\varphi, \psi \in \text{End}_N(X)$ be such that $\varphi - \psi \in N \text{End}(\underline{X})$. Then φ and ψ agree on the set $\underline{X}[N]$.*

Proof. Let us write $\varphi - \psi = N\beta$, for some $\beta \in \text{End}(\underline{X})$. For $P \in \underline{X}[N]$, we have that $\varphi(P) - \psi(P) = N\beta(P) = \beta(NP) = 0$, which proves the assertion. $\square$

Example 3.11 (Continuation of Examples 3.3 and 3.4). The functors TorPoints_N and MaxIso_N have the $(\text{mod } N)$ -congruence property by Lemma 3.10.

4 Statement of main theorem

Let $(X_1, x_1), (X_2, x_2)$ be two vertices in $\mathcal{G}^{\mathcal{F}}$. We say that (X_1, x_1) and (X_2, x_2) are *deg*-equivalent if there exists $f \in \text{Hom}_N(X_1, X_2)$ with $\mathcal{F}(f)(x_1) = x_2$ such that the reduced degree $\text{rdeg}(f) \equiv 1 \pmod{N}$. We define the space $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$ as the subspace of all functions in $L^2(\mathcal{G}^{\mathcal{F}}, \mu)$ that are constant on the deg-equivalence classes, where μ is the weight function $\mu(u) = \frac{1}{\#\text{Aut}(u)}$. Also, we define $L_0^2(\mathcal{G}^{\mathcal{F}}, \mu)$ as the orthogonal complement of $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$ with respect to the scalar product defined in (2).

We simply denote the adjacency operator of the graph $\mathcal{G}_{\ell}^{\mathcal{F}}$ by $A_{\ell} := A_{\mathcal{G}_{\ell}^{\mathcal{F}}}$. Note that for any $f \in L^2(\mathcal{G}^{\mathcal{F}}, \mu)$, the operator acts as

$$(A_{\ell}f)(u) = \sum_v m_{\ell}^{\mathcal{F}}(u, v)f(v), \quad \text{where } m_{\ell}^{\mathcal{F}}(u, v) \text{ as in (9).}$$

Our following main theorem generalizes [PW24b, Theorem 3.10] which concerns ℓ -isogeny graph of supersingular elliptic curves with the level structure. Informally, it says that random walks on $\mathcal{G}_{\ell}^{\mathcal{F}}$ equidistribute rapidly, up to unavoidable obstructions. There are two such obstructions, and both are well understood and controlled: the graph may be disconnected, and it may be periodic. They are accounted for by the space $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$, on which A_{ℓ} has the eigenvalues of maximal absolute value (“trivial” eigenvalues).

The significant part of the theorem is a bound for the operator norm of A_{ℓ} on the orthogonal complement $L_0^2(\mathcal{G}^{\mathcal{F}}, \mu)$, which allows to control the mixing rate. When the space $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$ is one-dimensional (e.g., when $\mathcal{F} = \text{Zero}$ and $\mathcal{G}_{\ell}^{\mathcal{F}} = \mathcal{G}_{\ell}(p, g)$), it is generated by the constant function, and the theorem says that random walks on $\mathcal{G}_{\ell}(p, g)$ converge rapidly to the unique stationary distribution.

Theorem 4.1. *Let p be a prime and $N \geq 1$ an integer coprime to p . Let $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \text{Sets}$ be a functor which has the $(\text{mod } N)$ -congruence property, and such that $\mathcal{F}(X)$ is finite for any object $X \in \mathbf{SS}_N(p, g)$. Then, for all primes $\ell \nmid pN$, the adjacency operator A_{ℓ} is a normal operator on $L^2(\mathcal{G}^{\mathcal{F}}, \mu)$ which stabilises the subspaces $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$ and $L_0^2(\mathcal{G}^{\mathcal{F}}, \mu)$. Moreover, we have that:*

1. *The A_{ℓ} for primes ℓ not dividing pN pairwise commute.*
2. *All eigenvalues of A_{ℓ} on $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$ have absolute value $\prod_{i=1}^g (1 + \ell^i)$.*
3. *The operator norm of A_{ℓ} on $L_0^2(\mathcal{G}^{\mathcal{F}}, \mu)$ is at most $\prod_{i=1}^g (1 + \ell^i) \cdot \left(\frac{2\sqrt{\ell}}{\ell+1}\right)^g$.*
4. *Conditional on the “weighted fundamental lemma” (see Remark 5.35) the operator norm of A_{ℓ} on $L_0^2(\mathcal{G}^{\mathcal{F}}, \mu)$ is at most $\prod_{i=1}^g (1 + \ell^i) \cdot \left(\frac{2}{\ell^{g/2} + \ell^{-g/2}}\right)$.*

The following companion proposition provides additional detail. Recall that the two obstructions mentioned above (disconnectedness and periodicity) are controlled by the trivial eigenvalues (see Section 2.3), so by the subspace $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$. When $L_{\text{deg}}^2(\mathcal{G}^{\mathcal{F}}, \mu)$ has dimension greater than one random walks do not converge to a unique limit: they converge to limits in a potentially periodic pattern.

The following result, which generalizes [PW24b, Proposition 3.11], makes Theorem 4.1 easier to apply in such cases by giving further details regarding $\mathcal{G}_\ell^\mathcal{F}$ alongside an explicit characterization of $L_{\deg}^2(\mathcal{G}^\mathcal{F}, \mu)$.

Proposition 4.2. *Let $X = (\underline{X}, \lambda)$ be any object of $\mathbf{SS}_N(p, g)$. With the same hypotheses as in Theorem 4.1:*

1. *For every morphism f in $\mathbf{SS}_N(p, g)$, the map $\mathcal{F}(f)$ is a bijection;*
2. *Let $\mathrm{GSp}(X[N])$ be the group of pairs $(g, s) \in \mathrm{GL}(\underline{X}[N]) \times (\mathbb{Z}/N\mathbb{Z})^\times$ such that $e_N^\lambda(g(P), g(Q)) = se_N^\lambda(P, Q)$ for all $P, Q \in \underline{X}[N]$. The morphism $\mathrm{End}_N(X) \rightarrow \mathrm{GSp}(X[N])$ is surjective, inducing an action of $\mathrm{GSp}(X[N])$ on $\mathcal{F}(X)$.*

We have the following description of the canonical projection $\pi : \mathcal{G}_\ell^\mathcal{F} \rightarrow (\mathcal{G}_\ell^\mathcal{F})_{\mathrm{triv}}$ of Proposition 2.13, which encodes connected components and periods of $\mathcal{G}$, uniformly in ℓ . Choose representatives $x_1, \dots, x_n$ of the orbits for the action of $\mathrm{GSp}(X[N])$ on $\mathcal{F}(X)$ and let H_i denote the stabilizer of x_i in $\mathrm{GSp}(X[N])$. Denote by $\nu : \mathrm{GSp}(X[N]) \rightarrow (\mathbb{Z}/N\mathbb{Z})^\times$ the second projection. Let $\mathcal{G}_{\deg}$ denote the graph with edges labelled by elements of $(\mathbb{Z}/N\mathbb{Z})^\times$ and with

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

Then:

3. *there exists a unique morphism of graphs*

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

such that for all i we have $\mathrm{Deg}(X, x_i) = 1 \in (\mathbb{Z}/N\mathbb{Z})^\times / \nu(H_i)$ and for every edge f of $\mathcal{G}^\mathcal{F}$, the edge $\mathrm{Deg}(f)$ is labelled by $\nu(f) \pmod{N}$;

4. *the map Deg is surjective and $L_{\deg}^2(\mathcal{G}^\mathcal{F}, \mu)$ is the space of functions that factor through Deg . In fact, a stronger statement holds: for any prime $\ell \nmid pN$, the graph $\mathcal{G}_{\deg, \ell}$, obtained by keeping only edges labelled by $\ell \pmod{N}$, together with Deg is isomorphic to the canonical quotient $(\mathcal{G}_\ell^\mathcal{F})_{\mathrm{triv}}$.*

Remark 4.3. By Remark 2.7, the adjacency operators A_ℓ and B_ℓ on the graph $\mathcal{G}^\mathcal{F}$ satisfy the relation $A_\ell^* = M^{-1}B_\ell M$, where $(Mf)(u) = (\#\mathrm{Aut}(u))^{-1}f(u)$. Therefore, B_ℓ inherits properties similar to those found for A_ℓ in Theorem 4.1. For instance, the operator B_ℓ is normal, the operators B_ℓ commute for all primes $\ell \nmid pN$, and they share the exact same eigenvalues and operator norms as A_ℓ .

4.1 Graph-theoretic rapid mixing consequence

In this section, we consider a graph $\mathcal{G}^\mathcal{F}$ as in Theorem 4.1. To quantify the rapid mixing of a random walk in $\mathcal{G}_\ell^\mathcal{F}$, we establish the explicit walk length that suffices to reach a stationary distribution within a small error tolerance ϵ . To this end, we apply Proposition 2.9 to the graph $\mathcal{G}_\ell^\mathcal{F}$. Since A_ℓ is normal with the weight function $\mu(v) = \frac{1}{\#\mathrm{Aut}(v)}$, we see that B_ℓ is normal with the weight

function $\mu^{-1}(v) = \#\text{Aut}(v)$ (see Remark 4.3). The mass of the graph is the quantity $\mu(V) := \sum_{v \in V} \#\text{Aut}(v)^{-1}$, where V denotes the set of vertices. Fix an arbitrary starting point $v_0 = (X, x) \in V$ for a random walk. By using the (conditional) bound for $\lambda(\mathcal{G}_\ell^{\mathcal{F}})$ in Theorem 4.1, we obtain from Proposition 2.9 with the notation therein that:

$$\frac{1}{2} \|P_n - Q_n\|_1 \leq \frac{1}{2} \left(\frac{\lambda(\mathcal{G}_\ell^{\mathcal{F}})}{d} \right)^n \left(\frac{\mu(V)}{\mu(v_0)} \right)^{\frac{1}{2}} \leq \frac{1}{2} \left(\frac{2}{\ell^{g/2} + \ell^{-g/2}} \right)^n (\mu(V) \#\text{Aut}(v_0))^{\frac{1}{2}}.$$

For the statistical distance to be bounded by a small error tolerance $\epsilon > 0$, we require $\frac{1}{2} \left(\frac{2}{\ell^{g/2} + \ell^{-g/2}} \right)^n \mu(V)^{\frac{1}{2}} (\#\text{Aut}(v_0))^{\frac{1}{2}} \leq \epsilon$. By (6), it suffices to take:

$$n \geq \frac{\frac{1}{2} \log(\mu(V) \cdot \#\text{Aut}(v_0)) - \log(2\epsilon)}{\log(\ell^{g/2} + \ell^{-g/2}) - \log(2)}. \quad (10)$$

To get an explicit bound in terms of p, ℓ, g , we need to bound $\mu(V)$ and $\#\text{Aut}(v_0)$.

Lemma 4.4. *Let $m = \#\mathcal{F}(X_0)$. The mass of the graph $\mathcal{G}^{\mathcal{F}}$ is*

$$\mu(V) = m \cdot \prod_{i=1}^g \frac{|B_{2i}|}{4i} (p^i + (-1)^i) \leq m \cdot p^{\frac{g(g+1)}{2}} \cdot \Psi(g)$$

where $\Psi(g) = \left(\frac{3}{2}\right)^g \prod_{i=1}^g \frac{|B_{2i}|}{4i} \leq e^{\text{poly}(g)}$.

Proof. Let X and X' be any two objects in $\mathbf{SS}_N(p, g)$. Since $\mathcal{G}_{\ell'}(p, g)$ is connected for any $\ell' \nmid Np$, there exists an isogeny $f \in \text{Hom}_N(X, X')$. By Proposition 4.2, $\mathcal{F}(f) : \mathcal{F}(X) \rightarrow \mathcal{F}(X')$ is a bijection between finite sets, and so $\#\mathcal{F}(X) = \#\mathcal{F}(X')$, in particular, $\#\mathcal{F}(X) = \#\mathcal{F}(X_0) = m$, for any X .

We now show that the mass $\mu(V)$ of the set of the vertices of $\mathcal{G}_\ell^{\mathcal{F}}$ satisfies $\mu(V) = m\mu(\text{SS}(p, g))$. Fix an object $X \in \mathbf{SS}_N(p, g)$, and let $\mathcal{F}(X) = \{x_1, \dots, x_m\}$. By definition, (X, x_i) and (X, x_j) represent the same vertex if and only if there exists $\varphi \in \text{Aut}(X)$ with $\mathcal{F}(\varphi)(x_i) = x_j$. Thus, the size of the orbit $\text{Orbit}(x_i)$ of x_i under the action of $\text{Aut}(X)$ on $\mathcal{F}(X)$ gives the number of vertices in the equivalence class of (X, x_i) . Moreover, it is easy to see that the stabilizer group $\text{Stab}_x(\text{Aut}(X))$ is equal to $\text{Aut}(X, x)$. By the Orbit-Stabilizer Theorem, we have that $\#\text{Orbit}(x) \cdot \#\text{Stab}_x(\text{Aut}(X)) = \#\text{Aut}(X)$, so:

$$\mu(V) = \sum_{(X, x) \in V} \frac{1}{\#\text{Aut}(X, x)} = m \sum_{X \in \text{SS}(p, g)} \frac{1}{\#\text{Aut}(X)} = m\mu(\text{SS}(p, g)).$$

By the classical mass formula of $\text{SS}(p, g)$ (see Theorem 6.4), we have that $\mu(\text{SS}(p, g)) = \beta_g \prod_{i=1}^g (p^i + (-1)^i)$, where $\beta_g := \prod_{i=1}^g \frac{|B_{2i}|}{4i}$, and so,

$$\mu(\text{SS}(p, g)) \leq \beta_g \prod_{i=1}^g (p^i + 1) = \beta_g p^{g(g+1)/2} \prod_{i=1}^g (p^{-i} + 1) \leq \beta_g p^{g(g+1)/2} \left(\frac{3}{2}\right)^g,$$

and the assertion then follows. $\square$

We are now ready to prove Theorem 1.1, a specialization of the above bounds to the connected, aperiodic case.

Proof of Theorem 1.1. Since the graph is connected and aperiodic, it has a unique trivial eigenvalue, hence there is a unique stationary distribution. The theorem follows from Equation (10), Lemma 4.4, and the fact that $\#\text{Aut}(X) \leq ((6.31)2g)^{2g}$ by [Sil92, Theorem 3.2]. $\square$

To conclude this section, we give the following characterization of strongly connected and aperiodic graphs.

Lemma 4.5. *If $\mathcal{G}_{\text{deg}}$ has a single vertex, $\mathcal{G}_\ell^\mathcal{F}$ is strongly connected and aperiodic.*

Proof. The graph $\mathcal{G}_{\text{deg}}$ has a single vertex so, using the notation of Proposition 4.2, there is a single orbit and $\nu(H_1) = (\mathbb{Z}/N\mathbb{Z})^\times$. Thus $L_{\text{deg}}^2(\mathcal{G}^\mathcal{F}, \mu)$ is one-dimensional, spanned by the constant function, and so $d = \prod_{i=1}^g (1 + \ell^i)$ is the only eigenvalue of A_ℓ of the absolute value d . Hence, $\mathcal{G}_\ell^\mathcal{F}$ is (strongly) connected and aperiodic. $\square$

We now give explicit examples, with bounds for walks in $\mathcal{G}_2(p, g)$ for $g \leq 3$, and in $\mathcal{G}_2^{\text{MaxIso}_N}$ for $g = 2$.

Example 4.6. Let us consider the graph $\mathcal{G}_2(p, 2)$, i.e., $\ell = g = 2$. Let $\epsilon = p^{-r}$, for some $r > 0$. Since ℓ and g are fixed, we can get an explicit bound for n by Theorem 1.1. Indeed, since $\beta_2 = \frac{1}{5760}$, we obtain that $n \geq (4.66 + 3.11r) \log(p)$.

This implies that for a standard 256-bit cryptographic prime ($p \approx 2^{256}$), it takes roughly 1997 steps for a random walk to reach the stationary distribution with a high security tolerance of $r = 1$. For comparison, relying on the alternative upper bound of 0.9996 from [ATY24] would require over 1108800 steps, demonstrating that our explicit bound is over 550 times more efficient.

If we instead consider the graph $\mathcal{G}_2(p, 3)$, i.e., $\ell = 2, g = 3$, we obtain a bound $n \geq (4.48 + 1.49r) \log(p)$ (since $\beta_3 = \frac{1}{2903040}$), and so the number of steps providing the same security is 1536. Since our bound decreases from $g = 2$ to $g = 3$, even though the graph is considerably larger, the number of steps decreases.

Example 4.7. (Continuation of example Example 3.4) Let $\mathcal{F} = \text{MaxIso}_N$ be the functor as in Example 3.4. We have from Example 3.11 that it has the mod N -congruence property, and so Proposition 4.2 applies to the graph $\mathcal{G}^\mathcal{F}$. In fact, this result gives that $\mathcal{G}^\mathcal{F}$ is connected and aperiodic. To show this, it suffices to show that $\mathcal{G}_{\text{deg}}$ is trivial by the discussion of the trivial eigenvalues in Section 2.3.

Indeed, the group $\text{GSp}(X[N])$ acts transitively on the set of maximal isotropic subspaces $\mathcal{F}(X)$ because any basis of a maximal isotropic subspace can be extended to a standard symplectic basis for $X[N]$. So, there is only one orbit, say x_1 . Furthermore, since the stabilizer H_1 of x_1 contains block matrices of the form $\text{diag}(I_g, cI_g)$ for every scalar $c \in (\mathbb{Z}/N\mathbb{Z})^\times$, we find that $\nu(H_1) = (\mathbb{Z}/N\mathbb{Z})^\times$. Therefore, the graph $\mathcal{G}_{\text{deg}}$ consists of a single vertex.

By Theorem 1.1, we now can get an explicit bound for the number n of steps for a random walk on $\mathcal{G}^\mathcal{F}$ to reach the stationary distribution. Note that $\#\mathcal{F}(X) = \prod_{i=1}^g (N^i + 1)$, for any object X by Equation (1). Concretely, if we take $\ell = g = 2$, then $n \geq (4.66 + 3.11r) \log(p) + 4.66 \log N$. With the same error tolerance in Example 4.6, we get that $n \geq 2589$ for $N = 2^{127} - 1$.

5 Proof of Results

In this section, we prove our main results: Theorem 4.1 and Proposition 4.2. We begin by fixing some notation and definitions needed only for the proof, then we give an overview of the proof strategy, and finally we give the proof.

5.1 Notation and conventions

We denote $\widehat{\mathbb{Z}} = \varprojlim_N \mathbb{Z}/N\mathbb{Z}$ the profinite completion of $\mathbb{Z}$, which decomposes as $\prod'_\ell \mathbb{Z}_\ell$ (restricted product over all prime numbers ℓ). We denote $\mathbb{A}_f = \mathbb{Q} \otimes_{\mathbb{Z}} \widehat{\mathbb{Z}}$ (finite adèles) and $\mathbb{A} = \mathbb{R} \times \mathbb{A}_f$, which decomposes as a restricted product $\prod'_v \mathbb{Q}_v$ over all places v of $\mathbb{Q}$. For an integer $N > 1$ we denote $\mathbb{Z}_N = \prod_{\ell|N} \mathbb{Z}_\ell$, similarly for a set of primes Σ we denote $\mathbb{Z}_\Sigma = \prod'_{\ell \in \Sigma} \mathbb{Z}_\ell$. Similarly we denote $\mathbb{A}_N = \mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z}_N$ and $\mathbb{A}_\Sigma = \mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z}_\Sigma$, which again decomposes as a restricted product.

For a (commutative) ring R and an integer $n \geq 1$ we denote by $\mathbf{GL}_{n,R}$ the algebraic group scheme over R defined by $\mathbf{GL}_{n,R}(R') := \mathrm{GL}_n(R')$ where R' is any (commutative) R -algebra.

A *symplectic space* over a ring R is a pair (V, ψ) where V is a finite free R -module and ψ is a perfect¹ alternating form $V \times V \rightarrow R$. In this situation we denote by $\mathrm{GSp}(V, \psi)$ the group of *symplectic similitudes* (also known as *general symplectic group*) of (V, ψ) :

$$\mathrm{GSp}(V, \psi) := \{(g, s) \in \mathrm{GL}(V) \times R^\times \mid \forall x, y \in V, \psi(g(x), g(y)) = s\psi(x, y)\}.$$

More generally we denote by $\mathbf{GSp}(V, \psi)$ the corresponding linear algebraic group over R : for an R -algebra R' we define $\mathbf{GSp}(V, \psi)(R') := \mathrm{GSp}(R' \otimes_R V, \psi)$. This algebraic group is reductive (in the sense of [DG70, Exposé XIX]). We have a morphism $\nu : \mathbf{GSp}(V, \psi) \rightarrow \mathbf{GL}_{1,R}$ mapping $(g, s) \in \mathbf{GSp}(V, \psi)(R')$ to s . Its kernel is the *symplectic group* $\mathbf{Sp}(V, \psi)$, also a reductive (even semi-simple) algebraic group over R . We denote $\mathrm{Sp}(V, \psi) = \mathbf{Sp}(V, \psi)(R)$. For an integer $g \geq 0$ we denote by $\mathbf{GSp}_{2g,R}$ the general symplectic group associated to $V = R^{2g}$

and the alternating form with Gram matrix $\begin{pmatrix} 0 & I_g \\ -I_g & 0 \end{pmatrix}$. We denote by $\mathbf{Sp}_{2g,R}$ the associated symplectic group. Denote $\mathrm{GSp}_{2g}(R) = \mathbf{GSp}_{2g,R}(R)$ (naturally isomorphic to $\mathbf{GSp}_{2g,\mathbb{Z}}(R)$) and $\mathrm{Sp}_{2g}(R) = \mathbf{Sp}_{2g,R}(R)$.

For an algebraic group G over a (commutative) ring R , and for an R -algebra R' we denote $G_{R'} := G \times_{\mathrm{Spec} R} \mathrm{Spec} R'$ (base change).

5.2 Overview of the proof strategy

We first give an overview of our proof strategy to help a reader less familiar with the theory of automorphic forms to navigate the various statements.

¹i.e., it induces an isomorphism between V and its dual $\mathrm{Hom}_R(V, R)$.

Step 1: Translation to adèlic world. We begin by translating the geometric world of PPAVs to the algebraic world of finite adèles, allowing us to do ℓ -adic linear algebra for all ℓ at the same time. The translation proceeds in two steps.

First step of translation. First, we develop a lattice-theoretic description of isogenies. Fix superspecial $X_0 = (\underline{X}_0, \lambda_0)$ defined over $\overline{\mathbb{F}}_p$ and a prime $\ell \neq p$. A polarised (quasi-)isogeny $f : X_0 \rightarrow X$ induces a $\mathbb{Q}_\ell$ -linear isomorphism of Tate modules $V_\ell f : V_\ell \underline{X}_0 \xrightarrow{\sim} V_\ell \underline{X}$, where $V_\ell \underline{X}_0 := \mathbb{Q}_\ell \otimes_{\mathbb{Z}_\ell} T_\ell \underline{X}_0$ and similarly for $V_\ell \underline{X}$. The isomorphism $V_\ell f$ is a symplectic similitude, scaling the polarization pairing by the reduced degree of f . Pulling the lattice $T_\ell \underline{X} \subset V_\ell \underline{X}$ back along $V_\ell(f)^{-1}$ produces a lattice $L_\ell(f) := V_\ell(f)^{-1}(T_\ell \underline{X}) \subset V_\ell \underline{X}_0$ inside the fixed space $V_\ell \underline{X}_0$. We handle the case $\ell = p$ via the Dieudonné module in place of the Tate module; it is here that superspeciality enters, guaranteeing that the Dieudonné module of X_0 is as rigid as its Tate modules, in the sense that it admits a standard basis adapted to the polarization (Lemma 5.7). This allows us to assemble lattice data over all primes ℓ to obtain $L(f) := (L_\ell(f))_{\ell \text{ prime}}$ and obtain a bijection (for fixed X_0) between (quasi-)isogenies $f : X_0 \rightarrow X$ and certain families of lattices $L(f)$. This is proved in Proposition 5.2 for polarised isogenies and Corollary 5.3 for polarised quasi-isogenies. The latter bijection is the one that we will use. Indeed, unlike isogenies, quasi-isogenies can be composed *and* inverted. Thus for a fixed target X , the set of polarised quasi-isogenies $f : X_0 \rightarrow X$ carries a natural group action. In the second step, we will use this to remove the redundancy of different quasi-isogenies with the same target X . In the end, we can recover honest polarised isogenies of prescribed degree by restricting the allowed denominators.

Second step of translation. Next, we relate these lattices to double cosets of linear algebraic groups over $\mathbb{Z}$. A linear algebraic group G over $\mathbb{Z}$ is defined as

$$G(R) := \left\{ (g, s) \in (R \otimes_{\mathbb{Z}} \text{End}(\underline{X}_0)) \times R^\times \mid g^\dagger g = s \otimes \text{id}_{\underline{X}_0} \right\}$$

for any $\mathbb{Z}$ -algebra R . It can be viewed as defining a group of similitudes over R .

Let ℓ be a prime as before. For $\ell \neq p$, we observe that $G(\mathbb{Q}_\ell)$ acts transitively on $L(f)$ with stabilizer $G(\mathbb{Z}_\ell)$ (exactly the subgroup fixing $T_\ell \underline{X}_0$). There is a unique coset $g_\ell G(\mathbb{Z}_\ell)$ with $g_\ell(T_\ell \underline{X}_0) = L_\ell(f)$, namely $g_\ell = V_\ell(f)^{-1}$.

For $\ell = p$, we again work with the Dieudonné module instead. Putting together this data over all ℓ recovers a bijection between quasi-isogenies $X_0 \xrightarrow{f} X$ with X superspecial and the global quotient $G(\mathbb{A}_f)/G(\widehat{\mathbb{Z}})$ given by

$$(X, f) \mapsto gG(\widehat{\mathbb{Z}}), \text{ where } g = (g_\ell) \text{ as defined above.}$$

To get a bijection with the set of isomorphism classes of superspecial PPAVs over $\overline{\mathbb{F}}_p$ of dimension g , we have to take care of one redundancy: two targets X, X' can be isomorphic as abelian varieties even if arise from different quasi-isogenies f, f' out of X_0 . We remove this redundancy by quotienting on the left by $G(\mathbb{Q})$: any such f, f' differ by the unique element $f^{-1} \circ f' \in G(\mathbb{Q})$. Finally, we note that for this parametrization to reach *every* isomorphism class, we must show

that each superspecial PPAV over $\bar{\mathbb{F}}_p$ admits a polarised quasi-isogeny from the fixed base point X_0 . This is Theorem 5.10, which holds for supersingular PPAVs in general.

In conclusion, in Theorem 5.12 we obtain a bijection between the set of isomorphism classes of superspecial PPAVs over $\bar{\mathbb{F}}_p$ of dimension g and the double coset space $G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / G(\widehat{\mathbb{Z}})$.

Restricting the degrees of quasi-isogenies. We can go further and restrict the degrees of the quasi-isogenies we consider. To handle this case, we have to slightly modify the double coset space. Namely, we work with the subsets $\mathbb{Z}[\Sigma^{-1}]$ (instead of $\mathbb{Q}$), $\mathbb{Z}_\Sigma$ (instead of $\widehat{\mathbb{Z}}$) and $A_\Sigma := \mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z}_\Sigma$ (instead of $\mathbb{A}_f$). Then, in Theorem 5.12 we show that the image of the embedding

$$G(\mathbb{Z}[\Sigma^{-1}]) \backslash G(\mathbb{A}_\Sigma) / G(\mathbb{Z}_\Sigma) \hookrightarrow G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / G(\widehat{\mathbb{Z}})$$

is in bijection with the set of isomorphism classes $[X]$ such that there exists a polarised isogeny $f : X_0 \rightarrow X$ of reduced degree $\text{rdeg}(g)$ whose prime divisors are contained in Σ .

Step 2: Extending the correspondence to the case with level structure.

We would now like to extend the correspondence above to the case of PPAVs with level structure. More formally, for $N \geq 1$, an integer coprime to p , we are given a functor $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \text{Sets}$ satisfying the $(\text{mod } N)$ -congruence property (see Definition 3.8).

Alternative characterisation of the $(\text{mod } N)$ -congruence property. To work with the $(\text{mod } N)$ -congruence property in the proof, we use an alternative characterisation. In Proposition 5.19, we show that any functor $\mathcal{F}$ satisfying the $(\text{mod } N)$ -congruence property factors through the torsion functor $\mathcal{T}_N$, which maps a PPAV $(\underline{X}, \lambda) \in \mathbf{SS}_N(p, g)$ to $(X[N], e_N^\lambda)$. Namely, we can write $\mathcal{F} = \mathcal{F}' \circ \mathcal{T}_N$ for some functor $\mathcal{F}'$. This factorization is possible because the congruence property guarantees that if two varieties share identical N -torsion data, the functor yields the exact same output. Consequently, one can rigorously define a secondary function $\mathcal{F}'$ that operates directly on the torsion structure. The advantage of this alternate characterisation is that to construct our bijection, we can work with the finite group $X[N]$ and Weil pairing e_N^λ , instead of the full abelian variety X .

Group action on level structures. Recall that we have a factorisation $\mathcal{F} = \mathcal{F}' \circ \mathcal{T}_N$. The group $\text{GSp}(X[N]) = \text{Aut}(\mathcal{T}_N(X))$ acts on $X[N]$, so via the functor $\mathcal{F}'$ it acts on $\mathcal{F}'(X[N])$. We can ‘lift’ this finite action: For $X = X_0$ and a given $x_0 \in \mathcal{F}'(X_0[N])$ we have a natural map $G(\mathbb{Z}_N) \rightarrow \text{GSp}(X_0[N])$ which is surjective with kernel J_N . So, $G(\mathbb{Z}_N)/J_N$ acts on $\mathcal{F}'(X_0[N])$ and the stabilizer of x_0 in $G(\mathbb{Z}_N)$ is a subgroup containing J_N . In general, $\mathcal{F}(X_0)$ decomposes into several orbits under this action. We handle each orbit separately, and this is the source of the direct sum appearing in **Step 3**.

The bijection. We fix a base variety X_0 and a reference level structure x_0 . As before, we measure other varieties relative to this base point via quasi-isogenies. To add the data of the level structure to the bijection from **Step 1**, we use the action of $G(\mathbb{Z}_N)$ on x_0 , quotienting by the stabilizer K_N of x_0 in $G(\mathbb{Z}_N)$ to remove redundancies. Note that, from our discussion above, K_N contains J_N .

Concretely, on top of the coordinate in $G(\mathbb{A}_\Sigma)/G(\mathbb{Z}_\Sigma)$ recording the quasi-isogeny $f: X_0 \rightarrow X$ as in **Step 1**, we record a class $g_N K_N \in G(\mathbb{Z}_N)/K_N$ expressing how the level structure x on X compares, through f , to the reference level structure x_0 . In Theorem 5.21, we obtain a bijection between isomorphism classes of superspecial PPAVs equipped with level structure and the double coset:

$$G(\mathbb{Z}[\Sigma^{-1}]) \backslash (G(\mathbb{A}_\Sigma) \times G(\mathbb{Z}_N)) / (G(\mathbb{Z}_\Sigma) \times K_N).$$

Step 3: Relating the operators on each space. The bijections of **Steps 1–2** identify the vertices of the graph $\mathcal{G}_\ell^\mathcal{F}$ with an adèlic double coset space. In this step we upgrade this identification of sets to an identification of *operators*: under this correspondence, the adjacency operator A_ℓ of the graph becomes a classical object of the theory of automorphic forms, namely a Hecke operator T_ℓ .

Functions on the graph as automorphic forms. Denote by $\mathcal{A}^2(G)$ the space of automorphic forms for G , that is the space of $G(\mathbb{R}) \times K$ -finite functions $G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0} \rightarrow \mathbb{C}$ (where K is any compact open subgroup of $G(\mathbb{A}_f)$). The space $\mathcal{A}^2(G)$ is a unitary representation of $G(\mathbb{A})$. Since **Step 2** identifies the vertices of $\mathcal{G}_\ell^\mathcal{F}$ with the points of a double coset space, a function on the vertices is the same thing as an automorphic form invariant under the action of the relevant compact subgroups. In Proposition 5.23 we make this precise and construct an isomorphism

$$\beta: \bigoplus_{i=1}^r \mathcal{A}^2(G)^{G(\mathbb{R}) \times K_i} \longrightarrow L^2(\mathcal{G}_\ell^\mathcal{F}, \mu),$$

where the direct sum accounts for the fact that the set of level structures $\mathcal{F}(X_0)$ may decompose into several orbits $x_1, \dots, x_r$ under the action of $\mathrm{GSp}_{2g}(\mathbb{Z}/N\mathbb{Z})$, each orbit contributing a compact subgroup K_i built from the stabilizer of x_i as in **Step 2**. Moreover, β is an *isometry* (Proposition 5.24), provided each vertex (X, x) of the graph is weighted by the mass $\mu(X, x) = 1/\#\mathrm{Aut}(X, x)$ as in Equation (2). Indeed, the automorphisms of (X, x) are precisely the rational elements of G fixing the corresponding coset, so the portion of $G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0}$ lying above the vertex (X, x) has volume $1/\#\mathrm{Aut}(X, x)$.

ℓ -isogenies as a local double coset. Next we encode the edges of the graph. Let $D_\ell \subset G(\mathbb{Q}_\ell)$ be the subset of elements whose similitude factor is ℓ up to units and whose inverse is integral. We first show that D_ℓ is a *single* double coset $G(\mathbb{Z}_\ell) x G(\mathbb{Z}_\ell)$ (Lemma 5.25). In lattice terms this expresses the fact that all ℓ -isogenies look alike locally.

Lemma 5.26 then provides the dictionary for edges: given quasi-isogenies $f: X_0 \rightarrow X$ and $f': X_0 \rightarrow Y$, the ‘difference’ $f'f^{-1}: X \rightarrow Y$ is an ℓ -isogeny if

and only if the adèlic coordinates of X and Y agree at every place $q \neq \ell$ and differ at ℓ by an element of D_ℓ . In other words, taking one step in the graph $\mathcal{G}_\ell^\mathcal{F}$ is a purely local move at the place ℓ , prescribed by D_ℓ .

The adjacency operator is a Hecke operator. In Proposition 5.27 we show that β intertwines the two operators: $A_\ell \circ \beta = \beta \circ T_\ell$. The proof consists in unwinding the dictionary above, and noting that the edges of reduced degree ℓ leaving a vertex (X, x) are in bijection with the cosets of D_ℓ . Indeed, noting that the set D_ℓ decomposes into finitely many cosets $tG(\mathbb{Z}_\ell)$, and the associated Hecke operator T_ℓ sends an automorphic form f to the function $[g] \mapsto \sum_t f([gt])$, we can match the sum defining A_ℓ with the sum defining T_ℓ term by term.

This identification has a first important consequence: the operator A_ℓ is *normal* (Corollary 5.28). This is not automatic, since the graph is directed and A_ℓ is in general not self-adjoint. On the automorphic side, however, the adjoint of T_ℓ is the Hecke operator attached to D_ℓ^{-1} , i.e., the operator attached to *dual* ℓ -isogenies. An explicit computation gives $D_\ell^{-1} = \ell D_\ell$, so that T_ℓ^* differs from T_ℓ by translation by the *central* element $\ell \in G(\mathbb{Q}_\ell)$ and therefore commutes with T_ℓ . The same argument shows that the operators A_ℓ for varying ℓ pairwise commute, as the corresponding Hecke operators act at distinct places. Normality brings us to **Step 4**: A_ℓ is diagonalisable in an orthogonal basis of $L^2(\mathcal{G}_\ell^\mathcal{F}, \mu)$, and its eigenvalues govern the convergence of random walks in the graph.

Step 4: Bounding the eigenvalues of the operators. Due to **Step 3**, bounding the eigenvalues of A_ℓ is now a problem about automorphic forms, more precisely the problem of bounding the eigenvalue of T_ℓ corresponding to the local component π_ℓ at ℓ of an automorphic representation π occurring in $\mathcal{A}_0^2(G)$, the space of automorphic forms $G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0} \rightarrow \mathbb{C}$ orthogonal to all automorphic forms factoring through the map $G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0} \rightarrow \mathbb{Q}^\times \backslash \mathbb{A}^\times / \mathbb{R}_{>0}$ induced by the similitude character.

For such a representation π_ℓ of $G(\mathbb{Q}_\ell) \simeq \mathrm{GSp}_{2g}(\mathbb{Q}_\ell)$ we have that π_ℓ is irreducible and unitary and does not have non-zero invariants under $\mathrm{Sp}_{2g}(\mathbb{Q}_\ell)$ (Lemma 5.31). In Section 5.5 we simply apply a theorem of Oh, giving a bound on Hecke eigenvalues in irreducible unitary unramified representations of $\mathrm{GSp}_{2g}(\mathbb{Q}_\ell)$ which are not of dimension 1.

It is well-known that local components of automorphic representations should be much more constrained than arbitrary unitary representations, e.g., for $g = 1$ the Ramanujan conjecture (proved by Deligne) and the Jacquet–Langlands correspondence imply π_ℓ is *tempered*. Unfortunately it is known that this does not hold for arbitrary groups. Arthur made very precise conjectures [Art89] explaining and quantifying this failure. In our case G^1 is an inner form of $\mathrm{Sp}_{2g, \mathbb{Q}}$, for which Arthur proved in [Art13] (an unconditional formulation of) his conjectures. Using the stabilization of the trace formula (to compare automorphic spectra of inner forms), known cases of the Ramanujan–Petersson conjecture *for general linear groups*, and simple combinatorics we derive a better bound in Section 5.6.

Following this blueprint, we provide the proof of the main results.

5.3 Translation to automorphic forms

We first recall the translation between functions on isogeny graphs and spaces of automorphic forms for certain inner forms of general symplectic group. The technique is well-known as it can be understood as a special case of a step in the Ihara–Langlands–Kottwitz method (this is at least as old as [Mil79], the results that we will need may be seen as special cases of results from [Kot92]). For clarity we find it useful to spell out the details in the superspecial case, which also relies on Lemma 5.7.

Equivalent definitions of superspeciality. We denote by α_p the group scheme over $\mathbb{F}_p$ defined by $\alpha_p(R) = \{x \in R \mid x^p = 0\}$ for an $\mathbb{F}_p$ -algebra R , with the addition law inherited from R .

There are several equivalent definitions of “superspecial” for abelian varieties over $\mathbb{F}_p$, by results of Oort, Deligne and Ogus [LO98, §1.6]. Let $\underline{X}$ be an abelian variety of dimension g over $\mathbb{F}_p$. The following are equivalent (possible definitions for “superspecial”):

1. $\underline{X}$ is isomorphic to a product of g supersingular elliptic curves,
2. the p -divisible group $\underline{X}[p^\infty]$ is isomorphic to a product of g copies of the p -divisible group of a supersingular elliptic curve,
3. the a -number $a(\underline{X}) := \dim_{\mathbb{F}_p} \text{Hom}(\alpha_{p, \mathbb{F}_p}, \underline{X}[p])$ is equal to g .

Moreover for $g > 1$ all superspecial $\underline{X}$ abelian varieties of dimension g over $\mathbb{F}_p$ are isomorphic.

For the rest of Section 5 we work with the (a priori weaker) definition $a(\underline{X}) = g$. As we point out in Remark 5.16, simpler versions of the arguments in Section 5.3 allow one to recover Oort’s theorem (that all supersingular abelian varieties of dimension $g > 1$ are isomorphic).

Lattice-theoretic description of isogenies. For the translation we need to fix a basepoint, that is a principally polarized supersingular abelian variety $X_0 = (\underline{X}_0, \lambda_0)$ over $\mathbb{F}_p$ of dimension g . Ultimately (after Theorem 5.10) we will need to assume that $\underline{X}_0$ is superspecial so for simplicity we make this assumption for the rest of Section 5.3 (for concreteness one can take $\underline{X}_0 = E^g$ for some supersingular elliptic curve E , with the product polarization).

The pairings that we consider as in Remark 2.3 allow us to take the limit of the pairings $e_{\ell^n}^\lambda$ on the Tate module, where ℓ denotes a prime number not equal to p . More precisely, let $\mathbb{Z}_\ell(1) := T_\ell \mathbb{G}_{m, k} = \varprojlim \mu_{\ell^n}$, and then the principal polarization λ on $\underline{X}$ defines a perfect alternating bilinear form (see [Mil86, p.132]):

$$e_{\mathbb{Z}_\ell}^\lambda : T_\ell \underline{X} \times T_\ell \underline{X} \rightarrow \mathbb{Z}_\ell(1), \quad \text{defined by} \quad (v, w) \mapsto (\bar{e}_{\ell^n}(v_n, \lambda(w_n)))_{n \geq 1}, \quad (11)$$

where $v = (v_n)_{n \geq 1}$ with $v_n \in \underline{X}[\ell^n]$. We have by [Mil86, p.137] a useful property that

$$e_{\mathbb{Z}_\ell}^\lambda(\alpha u, v) = e_{\mathbb{Z}_\ell}^\lambda(u, \alpha^\dagger v), \quad \text{for any } u, v \in T_\ell \underline{X}, \quad \text{and } \alpha \in \text{End}(\underline{X}). \quad (12)$$

We recall the parametrization of isomorphism classes of pairs (X, f) , where X is a principally polarized abelian variety and $f : X_0 \rightarrow X$ is a polarized isogeny, in terms of lattices.

Consider first a prime $\ell \neq p$. Any polarized isogeny $f : X_0 \rightarrow X$ induces $T_\ell(f) : T_\ell X_0 \rightarrow T_\ell X$ satisfying

$$e_{\mathbb{Z}_\ell}^\lambda((T_\ell f)(v), (T_\ell f)(w)) = \text{rdeg}(f) \cdot e_{\mathbb{Z}_\ell}^{\lambda_0}(v, w).$$

In particular $T_\ell(f)$ is injective and induces an isomorphism $V_\ell(f)$ between the $\mathbb{Q}_\ell$ -vector spaces $V_\ell X_0 := \mathbb{Q}_\ell \otimes_{\mathbb{Z}_\ell} T_\ell X_0$ and $V_\ell X$.

For a $\mathbb{Z}_\ell$ -lattice L in $V_\ell X_0$ define its dual

$$L^\sharp := \{v \in V_\ell X_0 \mid e_{\mathbb{Z}_\ell}^{\lambda_0}(v, L) \subset \mathbb{Z}_\ell\}.$$

The $\mathbb{Z}_\ell$ -lattice $L_\ell(f) = V_\ell(f)^{-1}(T_\ell(X))$ in $V_\ell X_0$ contains $T_\ell X_0$ and satisfies $L_\ell(f)^\sharp = \text{rdeg}(f)L_\ell(f)$.

Now we consider the $\ell = p$ case. Let $W = W(\bar{\mathbb{F}}_p)$ be the ring of Witt vectors of $\bar{\mathbb{F}}_p$. Denote by τ the automorphism of W lifting the Frobenius automorphism $x \mapsto x^p$ of $\bar{\mathbb{F}}_p$. Let $K = W[\frac{1}{p}]$.

Definition 5.1. *A Dieudonné module, or $W\{F, V\}$ -module, is a finitely generated W -module M endowed with a τ -semilinear endomorphism F and a τ^{-1} -semilinear endomorphism V satisfying $FV = VF = p\text{Id}_M$.*

Associated to an abelian variety X of dimension g over $\bar{\mathbb{F}}_p$ (via its p -divisible group $X[p^\infty]$) is the (contravariant) Dieudonné module $M(X) = M(X[p^\infty])$, a Dieudonné module which as W -module is free of rank $2g$ (see [Dem72, Theorem III.8] or [Fon77, Théorème III.1]), which can be identified with the projective limit of the Dieudonné modules of finite length $M(X_0[p^n])$. This Dieudonné module replaces the Tate module for $\ell = p$.

As in the $\ell \neq p$ case we have (using [Mum70, Theorem 15.1], [Dem72, Theorem III.8 (4)], [Fon77, §III.5]) for a principally polarized $X = (\underline{X}, \lambda)$ a perfect alternating bilinear form

$$e_W^\lambda : M(\underline{X}) \times M(\underline{X}) \longrightarrow W$$

satisfying $e_W^\lambda(F(x), y) = \tau(e_W^\lambda(x, V(y)))$ for all $x, y \in M(\underline{X})$. We denote by $M(X)$ the Dieudonné module $M(\underline{X})$ endowed with this pairing. A polarized isogeny $f : X_0 \rightarrow X$ of reduced degree $\text{rdeg}(f)$ induces a morphism of Dieudonné modules $M(f) : M(\underline{X}) \rightarrow M(\underline{X}_0)$ satisfying

$$e_W^{\lambda_0}(M(f)(x), M(f)(y)) = \text{rdeg}(f)e_W^\lambda(x, y)$$

for all $x, y \in M(\underline{X})$. The morphism $M(f)$ is clearly injective. Let $L_p(f) := M(f)(M(\underline{X}))$, a sub-Dieudonné module of $M(\underline{X}_0)$ satisfying $L_p(f)^\sharp = \text{rdeg}(f)^{-1}L_p(f)$.

Proposition 5.2. *We have a bijection $[X_0 \xrightarrow{f} X] \mapsto L(f) := (L_\ell(f))_{\ell \text{ prime}}$ between the set of isomorphism classes of polarized isogenies between X_0 and a principally polarized abelian variety X and the set of lattices $(L_\ell)_{\ell \text{ prime}}$ satisfying:*

- for $\ell \neq p$ the $\mathbb{Z}_\ell$ -lattice L_ℓ of $V_\ell \underline{X}_0$ contains $T_\ell \underline{X}_0$ and satisfies $L_\ell^\sharp = \ell^{r_\ell} L_\ell$ for some integer r_ℓ (necessarily ≥ 0),
- the W -submodule L_p of $M(\underline{X}_0[p^\infty])$ has rank $2g$, is stable under F and V and satisfies $L^\sharp = p^{-r_p} L$ for some integer r_p (necessarily ≥ 0),
- for almost all primes $\ell \neq p$ we have $L_\ell = T_\ell \underline{X}_0$ (equivalently $r_\ell = 0$).

In this bijection the reduced degree $\text{rdeg}(f)$ of f is equal to $\prod_{\ell \text{ prime}} \ell^{r_\ell}$.

Proof. Fix an integer $c > 0$, written as $\prod_{\ell} \ell^{r_\ell}$. We know by Definition 2.6 that isomorphism classes of polarized isogenies $f : X_0 \rightarrow X$ are parametrized by subgroup schemes of $\underline{X}_0[c]$ of order c^g which are maximal isotropic for the Weil pairing $\underline{X}_0[c] \times \underline{X}_0[c] \rightarrow \mu_c$, by mapping f to $\ker f$. We have a canonical decomposition $\underline{X}_0[c] \simeq \prod_{\ell|c} \underline{X}_0[\ell^{r_\ell}]$, so we are left to check that for any prime $\ell|c$ isotropic subgroup schemes K_ℓ of order $\ell^{r_\ell g}$ of $\underline{X}_0[\ell^{r_\ell}]$ are parametrized by lattices L_ℓ as in the proposition.

For $\ell \neq p$ we have a bijection between overlattices L_ℓ of $T_\ell \underline{X}_0$ and finite subgroups K_ℓ of $\underline{X}_0[\ell^\infty]$ characterized by the property: for n large enough the image of $\ell^n L_\ell$ in $\underline{X}_0[\ell^n]$ is K_ℓ . In this bijection we have $K_\ell \simeq L_\ell / T_\ell \underline{X}_0$. For any integer $r \geq 0$ we have a commutative diagram

$$\begin{array}{ccc} \underline{X}_0[\ell^r] \times \underline{X}_0[\ell^r] & \xrightarrow{e_{\ell^r}^{\lambda_0}} & \mu_{\ell^r} \\ \downarrow \sim & & \downarrow \sim \\ (\ell^{-r} T_\ell \underline{X}_0 / T_\ell \underline{X}_0) \times (\ell^{-r} T_\ell \underline{X}_0 / T_\ell \underline{X}_0) & \xrightarrow{\ell^r e_{\mathbb{Z}_\ell}^{\lambda_0}} & (\ell^{-r} \mathbb{Z}_\ell / \mathbb{Z}_\ell) \otimes_{\mathbb{Z}_\ell} \mathbb{Z}_\ell(1) \end{array}$$

and we easily deduce that we have $L_\ell^\sharp = \ell^r L_\ell$ if and only if K_ℓ is a maximally isotropic subgroup of $\underline{X}_0[\ell^r]$.

For $\ell = p$ finite subgroup schemes of $\underline{X}_0[p^\infty]$ correspond via [Fon77, Théorème III.1] to quotients of finite length of $M := M(\underline{X}_0[p^\infty])$ by sub- W -modules stable under F and V ; the subgroup $\underline{X}_0[p^{r_p}]$ corresponds to $p^{r_p} M$ and the pairing $\underline{X}_0[p^{r_p}] \times \underline{X}_0[p^{r_p}] \rightarrow \mu_{p^{r_p}}$ corresponds to the pairing $M/p^{r_p} M \times M/p^{r_p} M \rightarrow p^{-r_p} W/W$ induced by $p^{-r_p} \langle \cdot, \cdot \rangle$, so we conclude essentially as in the $\ell = p$ case. $\square$

Recall that a (polarized) quasi-isogeny between X_0 and X is an element f of $\mathbb{Q} \otimes_{\mathbb{Z}} \text{Hom}(X_0, X)$ such that there exists an integer $n > 0$ such that nf belongs to $\text{Hom}(X_0, X)$ and is a polarized isogeny. The reduced (virtual) degree $\text{rdeg}(f) \in \mathbb{Q}_{>0}$ of f is defined as $\text{rdeg}(nf)/n^2$ (which does not depend on n).

Corollary 5.3. *We have a bijection $[X_0 \xrightarrow{f} X] \mapsto L(f) := (L_\ell(f))_{\ell \text{ prime}}$ between the set of isomorphism classes of polarized quasi-isogenies between X_0 and a principally polarized abelian variety X and the set of lattices $(L_\ell)_\ell$ satisfying:*

- for $\ell \neq p$ the $\mathbb{Z}_\ell$ -lattice L_ℓ of $V_\ell X_0$ satisfies $L_\ell^\sharp = \ell^{r_\ell} L_\ell$ for some integer r_ℓ ,
- the W -lattice L_p in $K \otimes_W M(\underline{X}_0[p^\infty])$ has rank $2g$, is stable under F and V and satisfies $L^\sharp = p^{-r_p} L$ for some integer r_p ,
- for almost all $\ell \neq p$ we have $L_\ell = T_\ell \underline{X}_0$ (equivalently $r_\ell = 0$).

In this bijection the reduced degree $\text{rdeg}(f)$ of f is equal to $\prod_{\ell} \ell^{r_\ell}$.

Supersingular Dieudonné modules and polarizations.

Lemma 5.4. *Let L be a Dieudonné module which is free of rank $2g$ as W -module. Denote $L^{(2)} = \{x \in L \mid F^2(x) = -px\}$. The following are equivalent:*

1. $\dim_{\mathbb{F}_p} L/(F(L) + V(L)) \geq g$,
2. *the natural map $W \otimes_{W(\mathbb{F}_{p^2})} L^{(2)} \rightarrow L$ is an isomorphism.*

If these conditions are satisfied then we have $\dim_{\mathbb{F}_p} L/(F(L) + V(L)) = g$.

Proof. Denote by $\overline{F}$ and $\overline{V}$ the induced (semi-linear) endomorphisms of L/pL . We have $\text{im } \overline{V} = \ker \overline{F}$ and $\text{im } \overline{F} = \ker \overline{V}$: one inclusion follows from $\overline{V}\overline{F} = \overline{F}\overline{V} = 0$, the other from $VF = FV = p\text{Id}_L$ and the fact that L is p -torsion free. In particular we have $\dim_{\mathbb{F}_p} \text{im } \overline{V} + \dim_{\mathbb{F}_p} \text{im } \overline{F} = 2g$.

Assume $\dim_{\mathbb{F}_p} L/(F(L) + V(L)) \geq g$, i.e. $\dim_{\mathbb{F}_p} (\text{im } \overline{V} + \text{im } \overline{F}) \leq g$. We deduce $\text{im } \overline{V} = \text{im } \overline{F}$ and that this subspace of L/pL has dimension g . We also deduce $F^2(L) \subset pL$ and $V^2(L) \subset pL$, and so $p^{-1}F^2$ and $p^{-1}V^2$ are (semi-linear) endomorphisms of L which are inverse of each other. By a well-known argument (see [Fon90, p.259]) the natural morphism $W \otimes_{W(\mathbb{F}_{p^2})} L^{(2)} \rightarrow L$ is an isomorphism.

Conversely assume that $W \otimes_{W(\mathbb{F}_{p^2})} L^{(2)} \rightarrow L$ is an isomorphism. We have $V|_{L^{(2)}} = -F|_{L^{(2)}}$ and we deduce $\text{im } \overline{V} = \text{im } \overline{F}$. Combining with the equality $\dim_{\mathbb{F}_p} \text{im } \overline{V} + \dim_{\mathbb{F}_p} \text{im } \overline{F} = 2g$ proved above we conclude $\dim_{\mathbb{F}_p} L/(F(L) + V(L)) = g$. $\square$

Definition 5.5. *A Dieudonné module which is free of rank $2g$ (for some $g \geq 0$) as W -module and which satisfies the equivalent conditions of Lemma 5.4 is called *superspecial*.*

Remark 5.6. The definition is justified by the fact that an abelian variety $\underline{X}$ of dimension g over $\mathbb{F}_p$ is superspecial (i.e. $\dim_{\mathbb{F}_p}(\alpha_p, \underline{X}[p]) = g$) if and only if its Dieudonné module $L := M(\underline{X})$ is superspecial. Indeed the (contravariant) equivalence between finite p -group schemes and Dieudonné modules of finite length gives

$$\text{Hom}(\alpha_p, \underline{X}[p]) = \text{Hom}(\alpha_p, \underline{X}) \simeq \text{Hom}_{W,F,V}(L, M(\alpha_p))$$

and $M(\alpha_p)$ is isomorphic to $\mathbb{F}_p$ with $F = V = 0$, so this last space is naturally isomorphic to the dual of the $\mathbb{F}_p$ -vector space $L/(F(L) + V(L))$.

Although we will not need it one can deduce from Lemma 5.4 that all superspecial Dieudonné modules of rank $2g$ are isomorphic, more precisely any such module admits a basis $(e_1, f_2, \dots, e_g, f_g)$ satisfying $F(e_i) = -f_i$ and $F(f_i) = pe_i$ (take a family $(e_1, \dots, e_g)$ in $L^{(2)}$ lifting a basis of $L/(F(L) + V(L))$). Instead, we directly prove the following lemma which also includes alternating bilinear forms (related to polarizations of abelian varieties).

Lemma 5.7. *Let L be superspecial Dieudonné module of rank $2g$ endowed with a perfect alternating bilinear form $\langle \cdot, \cdot \rangle : L \times L \rightarrow W$ satisfying $\langle F(x), y \rangle = \tau(\langle x, V(y) \rangle)$ for all $x, y \in L$. Then there exists a basis $(e_1, f_1, \dots, e_g, f_g)$ of L such that the planes $We_i \oplus Wf_i$ are pairwise orthogonal for $\langle \cdot, \cdot \rangle$ and for all $1 \leq i \leq g$ we have $\langle e_i, f_i \rangle = 1$, $F(e_i) = -f_i$ and $F(f_i) = pe_i$.*

Proof. We prove the lemma by induction on g . The initial case $g = 0$ is trivial. So assume $g > 0$. Assume (arguing by contradiction) that for any $x \in L^{(2)}$ we have $\langle x, F(x) \rangle \in pW$. Polarizing this identity we obtain, for all $x, y \in L^{(2)}$ and $\mu \in W(\mathbb{F}_{p^2})^\times$

$$\mu^{p-1} \langle x, F(y) \rangle + \langle y, F(x) \rangle \in pW.$$

Since the map $\mathbb{F}_{p^2} \rightarrow \mathbb{F}_{p^2}$, $\mu \mapsto \mu^{p-1}$ is not constant, this implies $\langle x, F(y) \rangle \in pW$ for all $x, y \in L^{(2)}$, and thus also for all $x, y \in L$ by bilinearity. By perfection of $\langle \cdot, \cdot \rangle$ this implies $F(L) \subset pL$, in contradiction with $\dim_{\mathbb{F}_p} \text{im } F = g$. Thus there exists $x \in L^{(2)}$ satisfying $\langle x, F(x) \rangle \in W^\times$. We have $V|_{L^{(2)}} = -F|_{L^{(2)}}$ and so

$$\langle F(x), x \rangle = \tau(\langle x, V(x) \rangle) = -\tau(\langle x, F(x) \rangle),$$

i.e. $\langle x, F(x) \rangle \in \mathbb{Z}_p^\times$. For $\mu \in W(\mathbb{F}_{p^2})^\times$ we have

$$\langle \mu x, F(\mu x) \rangle = \mu \tau(\mu) \langle x, F(x) \rangle$$

and the map

$$W(\mathbb{F}_{p^2})^\times \longrightarrow \mathbb{Z}_p^\times, \quad \mu \longmapsto \mu \tau(\mu).$$

is surjective so up to replacing x by μx we may assume $\langle x, F(x) \rangle = 1$. Now define $e_1 = x$ and $f_1 = -F(x)$. The plane $P = We_1 \oplus Wf_1 \subset L$ is clearly stable under F and V , and the restriction of $\langle \cdot, \cdot \rangle$ to P is perfect. Let L' be the orthogonal of P . The compatibility between $\langle \cdot, \cdot \rangle$ and (F, V) easily implies that L' is also stable under F and V , and we conclude using the induction hypothesis. $\square$

Group-theoretic parametrization of superspecial PPAVs. Recall from §2.2 that the Rosati involution with respect to λ_0 on $\text{End}(\underline{X}_0)$ defined by $f \mapsto f^\dagger := \lambda^{-1} f^\vee \lambda$. Let G be the linear algebraic group over $\mathbb{Z}$ defined by

$$G(R) = \left\{ (g, s) \in (R \otimes_{\mathbb{Z}} \text{End}(\underline{X}_0)) \times R^\times \mid g^\dagger g = s \otimes \text{id}_{\underline{X}_0} \right\}$$

for any $\mathbb{Z}$ -algebra R . We have a morphism

$$\begin{aligned} \nu : G &\longrightarrow \mathbf{GL}_{1, \mathbb{Z}} \\ (g, s) &\longmapsto s \end{aligned}$$

and we denote by G^1 its kernel.

Recall from (11) that for any prime $\ell \neq p$ the polarization λ_0 defines a perfect alternating bilinear form $e_{\mathbb{Z}_\ell}^{\lambda_0}$ on $T_\ell \underline{X}_0$.

Lemma 5.8. *For any prime $\ell \neq p$ the natural morphism (of group schemes over $\mathbb{Z}_\ell$) $T_\ell : G_{\mathbb{Z}_\ell} \xrightarrow{\sim} \mathbf{GSp}(T_\ell \underline{X}_0, e_{\mathbb{Z}_\ell}^{\lambda_0})$ is an isomorphism.*

Proof. There exists a power q of p and a supersingular abelian variety A over $\mathbb{F}_q$ such that $\underline{X}_0$ is isomorphic to $\overline{\mathbb{F}}_p \times_{\mathbb{F}_q} A$. Up to replacing q by an even power of q we may assume that the Frobenius endomorphism acts by multiplication by (the integer) $q^{1/2}$ on $T_\ell A$. Tate's isogeny theorem [Tat66] tells us that for any integer $r \geq 1$ the natural map

$$\mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{End}(\mathbb{F}_{q^r} \times_{\mathbb{F}_q} A) \longrightarrow \mathrm{End}(T_\ell A)^{\mathrm{Frob}_q^r}$$

is an isomorphism. Noting that the target does not depend on r we deduce that the (finite free) $\mathbb{Z}$ -modules $\mathrm{End}(\mathbb{F}_{q^r} \times_{\mathbb{F}_q} A)$ are all equal and that the natural map

$$\mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{End}(\underline{X}_0) \longrightarrow \mathrm{End}(T_\ell \underline{X}_0)$$

is an isomorphism. This map is compatible with duality in the obvious sense and the lemma follows. $\square$

The following proposition easily follows.

Proposition 5.9. *The group scheme $G_{\mathbb{Z}[p^{-1}]} \rightarrow \mathrm{Spec} \mathbb{Z}[p^{-1}]$ is reductive.*

Theorem 5.10. *Let $X = (\underline{X}, \lambda)$ be a principally polarized supersingular abelian variety of dimension g over $\overline{\mathbb{F}}_p$. Then there exists a polarized isogeny $X_0 \rightarrow X$.*

Proof. Consider the functor T mapping a commutative ring (with unit) R to the set of pairs (f, s) with $f \in R \otimes_{\mathbb{Z}} \mathrm{Hom}(\underline{X}_0, \underline{X})$ and $s \in R^\times$ satisfying $f^* \lambda = s \lambda_0$. The theorem is equivalent to $T(\mathbb{Q}) \neq \emptyset$. The functor T is representable by a scheme of finite type over $\mathbb{Z}$ because $\mathrm{Hom}(\underline{X}_0, \underline{X})$ is a finite free $\mathbb{Z}$ -module and $f \mapsto f^\vee$ is linear. Moreover $T_{\mathbb{Q}}$ is a right $\overline{G}_{\mathbb{Q}}$ -torsor: the action is clearly free and transitive and we have $T(\mathbb{Z}_\ell) \neq \emptyset$ by Tate's isogeny theorem (similarly to the proof of Lemma 5.8) so T is not empty and $T(\mathbb{Q}) \neq \emptyset$. Thus T defines a class $\alpha \in H^1(\mathbb{Q}, G)$ and we will show that it is trivial, implying that $T(\mathbb{Q})$ is not empty. For every prime ℓ we have $H^1(\mathbb{Q}_\ell, G) = 1$ by Kneser's theorem and Hilbert 90. We have $\ker^1(\mathbb{Q}, G) = 1$ by [Kot84, Lemma 4.3.1 (b)] and Hilbert 90. Thus it is enough to show that the image of α in $H^1(\mathbb{R}, G)$ is trivial. This is [Kot92, Lemma 17.1] in a slightly obfuscated language. For the convenience of the reader we briefly explain the proof.

For a field E of characteristic zero consider a class $[c]$ in the Galois cohomology set $H^1(E, G)$. By Hilbert 90 there exists $e \in \overline{E} \otimes_{\mathbb{Z}} \mathrm{End}(\underline{X}_0)$ is such that $c(\sigma) = e^{-1} \sigma(e)$ for any $\sigma \in \mathrm{Gal}(\overline{E}/E)$. We can associate to $[c]$ the anti-involution

$$x \in \overline{E} \otimes_{\mathbb{Z}} \mathrm{End}(\underline{X}_0) \mapsto ee^\dagger x^\dagger (ee^\dagger)^{-1}$$

which turns out to be defined over E . In this way $H^1(E, G)$ parametrizes the anti-involutions of $E \otimes_{\mathbb{Z}} \mathrm{End}(\underline{X}_0)$ of the same type (symmetric or alternating, see [Kot92, §1]; this type is an invariant of the anti-involution after extending

scalars to $\overline{E}$) as the Rosati involution up to conjugation by an inner automorphism. Now by [Mum70, §21 Theorem 1] the Rosati involutions associated to the polarizations λ_0 and λ are positive. Translating the fact that $(\mathbb{R} \otimes_{\mathbb{Z}} \text{End}(\underline{X}_0))^\times$ acts transitively on the set of positive anti-involutions of the same type gives us that the image of α in $H^1(\mathbb{R}, G)$ is trivial. $\square$

Our next goal is to reformulate Corollary 5.3 in terms of the algebraic group G , at least when we restrict to superspecial PPAVs (this will be achieved in Theorem 5.12 below). At the primes $\ell \neq p$ this is straightforward: it is well-known that symplectic spaces of dimension $2g$ over a ring are all isomorphic, and we easily deduce a bijection $xG(\mathbb{Z}_\ell) \mapsto x(T_\ell \underline{X}_0)$ between $G(\mathbb{Q}_\ell)/G(\mathbb{Z}_\ell)$ and the set of $\mathbb{Z}_\ell$ -lattices $L \subset V_\ell \underline{X}_0$ satisfying $L^\sharp = \ell^r L$ for some integer r .

Now we consider the case $\ell = p$. Recall from [WM71, §II] that Tate's theorem also holds in this setting, and so denoting $M_0 := M(\underline{X}_0[p^\infty])$ the natural map

$$\mathbb{Z}_p \otimes_{\mathbb{Z}} \text{End}(\underline{X}_0) \longrightarrow \text{End}_{W,F,V}(M_0)$$

is an isomorphism of $\mathbb{Z}_p$ -algebras. Denoting $M_0[\frac{1}{p}] := K \otimes_W M_0$ we obtain anti-isomorphisms

$$M: G(\mathbb{Z}_p) \simeq \left\{ (g, s) \in \text{Aut}_{W,F,V}(M_0) \times \mathbb{Z}_p^\times \mid \forall x, y \in M_0, e_W^{\lambda_0}(g(x), g(y)) = s e_W^{\lambda_0}(x, y) \right\}, \quad (13)$$

$$M_K: G(\mathbb{Q}_p) \simeq \left\{ (g, s) \in \text{Aut}_{K,F,V}(M_0[\frac{1}{p}]) \times \mathbb{Q}_p^\times \mid \forall x, y \in M_0[\frac{1}{p}], e_W^{\lambda_0}(g(x), g(y)) = s e_W^{\lambda_0}(x, y) \right\}. \quad (14)$$

We have a bijection $gG(\mathbb{Z}_p) \mapsto M_K(g^{-1})(M_0)$ between $G(\mathbb{Q}_p)/G(\mathbb{Z}_p)$ and the set of W -modules $L \subset M_0[\frac{1}{p}]$ satisfying the following conditions:

- L is free of rank $2g$ and stable under F and V ,
- there exists $r \in \mathbb{Z}$ for which $L^\sharp = p^r L$ (equivalently the alternating bilinear form $p^r e_W^{\lambda_0}$ on $L \times L$ takes values in W and is perfect),
- the symplectic Dieudonné module $(L, F, V, p^r e_W^{\lambda_0})$ is isomorphic to $(M_0, F, V, e_W^{\lambda_0})$.

By Lemma 5.7 this image is precisely the set of Dieudonné submodules of rank $2g$ of $M_0[\frac{1}{p}]$ which are superspecial and essentially self-dual. Knowing this is crucial for the following definition.

Definition 5.11. Let ι_{X_0} be the map from the set of isomorphism classes of (polarized) quasi-isogenies $X_0 \xrightarrow{f} X$ with $\underline{X}$ superspecial to $G(\mathbb{A}_f)/G(\widehat{\mathbb{Z}})$ mapping the class of (X, f) to $gG(\widehat{\mathbb{Z}})$ where

- for $\ell \neq p$ the class $g_\ell G(\mathbb{Z}_\ell)$ is defined by $V_\ell(g_\ell)(T_\ell \underline{X}_0) = V_\ell(f)^{-1}(T_\ell \underline{X})$,
- the class $g_p G(\mathbb{Z}_p)$ is defined by $M_K(g_p^{-1})(M(\underline{X}_0[p^\infty])) = M_K(f)(M(\underline{X}[p^\infty]))$.

Theorem 5.12. The map ι_{X_0} is a bijection and induces a bijection $\bar{\iota}_{X_0}$ between

- the set of isomorphism classes of principally polarized superspecial abelian varieties over $\mathbb{F}_p$ of dimension g , and

– the finite² set $G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / G(\widehat{\mathbb{Z}})$.

For a set of primes Σ the image of the embedding

$$G(\mathbb{Z}[\Sigma^{-1}]) \backslash G(\mathbb{A}_\Sigma) / G(\mathbb{Z}_\Sigma) \hookrightarrow G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / G(\widehat{\mathbb{Z}})$$

corresponds under $\bar{\iota}_{X_0}$ to the set of isomorphism classes $[X]$ such that there exists a polarized isogeny $f : X_0 \rightarrow X$ such that the set of prime divisors of the reduced degree $\text{rdeg}(f)$ is contained in Σ .

Proof. The fact that ι_{X_0} is a bijection is a reformulation of Corollary 5.3. The fact that $\bar{\iota}_{X_0}$ is a bijection follows using Theorem 5.10: for any superspecial PPAV X of dimension g the set of quasi-isogenies $X_0 \rightarrow X$ is a right torsor under $G(\mathbb{Q})$. The last statement of the theorem is simply a variant proved in the same way. Note that the map is an embedding because, denoting by Σ^c the complement of Σ in the set of prime numbers, we have $G(\mathbb{Q}) \cap G(\mathbb{A}_\Sigma) G(\mathbb{Z}_{\Sigma^c}) = G(\mathbb{Z}[\Sigma^{-1}])$. $\square$

Theorem 5.13. *Let X be a principally polarized supersingular abelian variety over $\mathbb{F}_p$ of dimension g . Let $\ell \neq p$ be a prime number. There exists a polarized isogeny $X_0 \rightarrow X$ of reduced degree ℓ^{2r} for some integer $r \geq 0$.*

Proof. Taking $\Sigma = \{\ell\}$ in the last part of Theorem 5.12 we are left to show that the natural map

$$G^1(\mathbb{Q}_\ell) / G^1(\mathbb{Z}_\ell) \longrightarrow G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / G(\widehat{\mathbb{Z}})$$

is surjective. This follows from Lemma 5.14 below, applied with $N = 1$, $\mathcal{Y} = \emptyset$ and $q = \ell$. $\square$

Lemma 5.14. *Let N be a positive integer. Denote by J_N the kernel of the surjective morphism $G(\mathbb{Z}_N) \rightarrow G(\mathbb{Z}/N\mathbb{Z})$. Let $\mathcal{Y}$ be a set of prime numbers that do not divide pN . Assume that $\mathcal{Y}$ generates $(\mathbb{Z}/N\mathbb{Z})^\times$. Let q be a prime that does not divide pN (we could have $q \in \mathcal{Y}$). Denote $K^N = G(\prod_{\ell \nmid N} \mathbb{Z}_\ell)$. Then we have*

$$G(\mathbb{Q}) G^1(\mathbb{Q}_q) G(\mathbb{A}_\mathcal{Y}) K^N J_N = G(\mathbb{A}_f).$$

Proof. Our goal is to show that $G(\mathbb{Q}) G^1(\mathbb{Q}_q) G(\mathbb{A}_\mathcal{Y}) \backslash G(\mathbb{A}_f) / K^N J_N$ is a singleton. Consider the map induced by ν

$$G(\mathbb{Q}) G^1(\mathbb{Q}_q) G(\mathbb{A}_\mathcal{Y}) \backslash G(\mathbb{A}_f) / K^N J_N \xrightarrow{\bar{\nu}} \nu(G(\mathbb{Q})) \nu(G(\mathbb{A}_\mathcal{Y})) \backslash \mathbb{A}_f^\times / \nu(K^N J_N).$$

We have a short exact sequence

$$1 \rightarrow G_{\mathbb{Q}}^1 \rightarrow G_{\mathbb{Q}} \xrightarrow{\nu} \mathbf{GL}_{1,\mathbb{Q}} \rightarrow 1$$

which for any field F of characteristic zero yields an exact sequence (part of the long exact sequence in Galois cohomology)

$$G(F) \xrightarrow{\nu} F^\times \longrightarrow H^1(F, G^1). \quad (15)$$

²see [God64, Théorème 5]

Denote $C^N := \prod_{\ell \nmid N} \mathbb{Z}_\ell^\times$ and $C_N := \ker(\mathbb{Z}_N^\times \rightarrow (\mathbb{Z}/N\mathbb{Z})^\times)$. We have $\nu(G(\mathbb{A}_T)) = \mathbb{A}_T^\times$ and $\nu(K^N J_N) = C^N C_N$: at primes $\ell \neq p$ this is clear using the isomorphism $G_{\mathbb{Z}_\ell} \simeq \mathbf{GSp}_{2g, \mathbb{Z}_\ell}$, and at p this reduces to the $g = 1$ case (which is easy) using Lemma 5.7 and (13), (14). The Hasse principle says that

$$H^1(\mathbb{Q}, G^1) \longrightarrow H^1(\mathbb{R}, G^1)$$

is injective, so with the exact sequence (15) for $F = \mathbb{Q}$ we conclude $\nu(G(\mathbb{Q})) = \mathbb{Q}_{>0}$. We conclude that $\tilde{\nu}$ has target

$$\mathbb{Q}_{>0} \mathbb{A}_T^\times \backslash \mathbb{A}_f^\times / C^N C_N.$$

Denote by $\langle \mathcal{Y} \rangle \subset \mathbb{Q}_{>0}$ the subgroup generated by $\mathcal{Y}$. The natural map

$$\langle \mathcal{Y} \rangle \backslash \mathbb{Z}_N^\times / C_N \longrightarrow \mathbb{Q}_{>0} \mathbb{A}_T^\times \backslash \mathbb{A}_f^\times / C^N C_N$$

is an isomorphism, essentially because the multiplication map $\mathbb{Q}_{>0} \times \widehat{\mathbb{Z}}^\times \rightarrow \mathbb{A}_f^\times$ is an isomorphism. We have a natural isomorphism (reduction modulo N)

$$\langle \mathcal{Y} \rangle \backslash \mathbb{Z}_N^\times / C_N \simeq \langle \mathcal{Y} \rangle \backslash (\mathbb{Z}/N\mathbb{Z})^\times.$$

Since $\mathcal{Y}$ generates $(\mathbb{Z}/N\mathbb{Z})^\times$ this quotient has just one element, so the target of $\tilde{\nu}$ has just one element. Thus the natural map

$$G^1(\mathbb{Q}) G^1(\mathbb{Q}_q) \backslash G^1(\mathbb{A}_f) / (G^1(\widehat{\mathbb{Z}}) \cap (K^N J_N)) \longrightarrow G(\mathbb{Q}) G(\mathbb{A}_T) \backslash G(\mathbb{A}_f) / K^N J_N$$

is surjective. The source of this map is a singleton by the strong approximation theorem [PR94, Theorem 7.12]: the algebraic group $G_{\mathbb{Q}}^1$ is simply connected and $G_{\mathbb{Q}_q}^1 \simeq \mathbf{Sp}_{2g, \mathbb{Q}_q}$ is isotropic. $\square$

Remark 5.15. One can show that Theorem 5.13 also holds for $\ell = p$ if $g > 1$ using essentially the same argument: the point is that the algebraic group $G_{\mathbb{Q}_p}^1$ is isomorphic to the non-split inner form of $\mathbf{Sp}_{2g, \mathbb{Q}_p}$, which is isotropic if $g > 1$ (this easily reduces to the $g = 2$ case, which via the exceptional isomorphism $\mathbf{PGSp}_4 \simeq \mathbf{SO}_5$ reduces to the well-known fact that quadratic forms over $\mathbb{Q}_p$ in at least 5 variables are isotropic).

Remark 5.16. In [JZ23, Theorem 44] the authors obtain similar results using Oort's theorem which implies that the category of superspecial PPAV's is equivalent to its subcategory of polarizations on a fixed superspecial abelian variety (e.g. E^g for a supersingular elliptic curve E). In the end our arguments differ from theirs essentially in the order in which they are used. Indeed Oort's theorem can be proved by arguments similar to the ones above, removing polarizations. More precisely we can consider the algebraic group H of automorphisms of $\underline{X}_0$: for any commutative ring R define $H(R) = (R \otimes_{\mathbb{Z}} \text{End}(\underline{X}_0))^\times$. It contains G as a subgroup. Now for $g > 1$ the derived subgroup $H_{\mathbb{Q}}^1$ of $H_{\mathbb{Q}}$ is simply connected and $H_{\mathbb{R}}^1$ is isotropic (unlike $G_{\mathbb{R}}^1$), so strong approximation tells us that $H^1(\mathbb{Q})$ is dense in $H^1(\mathbb{A}_f)$ and it is easy to deduce that $H(\mathbb{Q}) \backslash H(\mathbb{A}_f) / H(\widehat{\mathbb{Z}})$, which parametrizes isomorphism classes of supersingular abelian varieties over $\overline{\mathbb{F}}_p$ of dimension g , has only one element.

Adding level structure. Now we add “elements of $\mathcal{F}$ ”. Let $N \geq 1$ be an integer coprime to p . In this section we denote by Σ be the set of primes which do not divide pN and write $\text{Hom}_\Sigma(\cdot, \cdot)$ for sets of homomorphisms in $\mathbf{SS}_N(p, g)$.

Proposition 5.17. *Point 2 in Proposition 4.2 holds true: for any object X of $\mathbf{SS}_N(p, g)$ the morphism $\text{End}_N(X) \rightarrow \text{End}(X[N])$ is surjective.*

Proof. Since X_0 is any object of $\mathbf{SS}_N(p, g)$ it is enough to prove this in the case $X = X_0$, in other words that $G(\mathbb{Z}[\Sigma^{-1}]) \rightarrow G(\mathbb{Z}/N\mathbb{Z})$ is surjective. Let $h_N \in G(\mathbb{Z}/N\mathbb{Z})$ and choose a lift $\tilde{h}_N \in G(\mathbb{Z}_N)$. Applying Lemma 5.14 with $\Upsilon = \Sigma$ and any $q \in \Sigma$ we may write (in $G(\mathbb{A}_{pN})$) $\tilde{h}_N = \gamma k$ where $\gamma \in G(\mathbb{Q})$, and $k \in G(\mathbb{Z}_p)J_N$. We clearly have $\gamma \in G(\mathbb{Z}[\Sigma^{-1}])$ and $\tilde{h}_N J_N = \gamma J_N$ in $G(\mathbb{A}_N)$, i.e. γ maps to h_N in $G(\mathbb{Z}/N\mathbb{Z})$. $\square$

Definition 5.18. *Let Alt_N be the groupoid of free $\mathbb{Z}/N\mathbb{Z}$ -modules of rank $2g$ endowed with a perfect alternating bilinear pairing taking values in $\mu_N(\mathbb{F}_p)$, with morphisms the linear isomorphisms which preserve pairings up to $(\mathbb{Z}/N\mathbb{Z})^\times$.*

Proposition 5.19. *Let $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \text{Sets}$ be a functor. Assume that for any object X of $\mathbf{SS}_N(p, g)$ we have that $\text{End}_N(X) \rightarrow \text{End}(X[N], e_N^\lambda)$ is surjective. Then, $\mathcal{F}$ has the $(\text{mod } N)$ -congruence property if and only if it is isomorphic to a functor of the form $\mathcal{F}' \circ \mathcal{T}_N$, where $\mathcal{T}_N$ is the functor*

$$\begin{aligned} \mathbf{SS}_N(p, g) &\longrightarrow \text{Alt}_N \\ (X, \lambda_X) &\longmapsto (X[N], e_N^\lambda). \end{aligned}$$

Proof. First suppose that $\mathcal{F}' \circ \mathcal{T}_N \cong \mathcal{F}$. Let $X \in \mathbf{SS}(p, g)$ and $\varphi, \psi \in \text{End}_N(X)$ such that $\varphi - \psi \in N \text{End}(X)$. We have $\mathcal{T}_N(\varphi) = \mathcal{T}_N(\psi)$, hence $(\mathcal{F}' \circ \mathcal{T}_N)(\varphi) = (\mathcal{F}' \circ \mathcal{T}_N)(\psi)$, and from the isomorphism we deduce $\mathcal{F}(\varphi) = \mathcal{F}(\psi)$.

Now, suppose that $\mathcal{F}$ has the $(\text{mod } N)$ -congruence property. Fix an element $X \in \mathbf{SS}(p, g)$, and for each $Y \in \text{Alt}_N$, fix an isomorphism $\varepsilon_Y : Y \rightarrow (X[N], e_N^\lambda)$. Define the functor $\mathcal{F}' : \text{Alt}_N \rightarrow \text{Sets}$ as follows:

- For any $Y \in \text{Alt}_N$, let $\mathcal{F}'(Y) = \mathcal{F}(X)$, and
- For any morphisms $f : Y \rightarrow Y'$ in Alt_N , let $\hat{f} \in \text{End}_N(X)$ such that $\hat{f} \equiv \varepsilon_{Y'} \circ f \circ \varepsilon_Y^{-1} \pmod{N}$ (such an isogeny exists because we assume that $\text{End}_N(X) \rightarrow \text{End}(X[N], e_N^\lambda)$ is surjective), and define $\mathcal{F}'(f) = \mathcal{F}(\hat{f})$. It is well-defined because of the $(\text{mod } N)$ -congruence property.

We get an isomorphism $\eta : \mathcal{F} \rightarrow \mathcal{F}' \circ \mathcal{T}_N$ defined as $\eta_X = \mathcal{F}'(\varepsilon_{X[N]}) : \mathcal{F}(X) \rightarrow \mathcal{F}(X)$ for any $X \in \mathbf{SS}(p, g)$. $\square$

Let $\mathcal{F} : \mathbf{SS}_N(p, g) \rightarrow \text{Sets}$ be a functor isomorphic to $\mathcal{F}' \circ \mathcal{T}_N$, where $\mathcal{T}_N$ is the functor

$$\begin{aligned} \mathbf{SS}_N(p, g) &\longrightarrow \text{Alt}_N \\ X = (\underline{X}, \lambda) &\longmapsto X[N] := (\underline{X}[N], e_N^\lambda) \end{aligned}$$

and $\mathcal{F}'$ is a functor from Alt_N to Sets . We fix an isomorphism $\mathcal{F} \simeq \mathcal{F}' \circ \mathcal{T}_N$, allowing us to pretend that this is an equality. For objects X, Y of we denote by $\Sigma^{-1} \text{Hom}_\Sigma(X, Y)$ the set of Σ -quasi-isogenies from X to Y , i.e. the set of $f \in \mathbb{Z}[\Sigma^{-1}] \otimes_{\mathbb{Z}} \text{Hom}(\underline{X}, \underline{Y})$ such that there exists an integer $n > 0$ whose set of prime divisors is contained in Σ and such that nf belongs to $\text{Hom}_\Sigma(X, Y)$. Since Alt_N is a groupoid the functor $\mathcal{F}'$ extends uniquely to Σ -quasi-isogenies. We may decompose $\mathcal{F}'(X_0[N])$ into orbits under $\text{GSp}(X_0[N]) = \text{Aut}(\mathcal{T}_N(X_0))$, and for an object X of $\mathbf{SS}_N(p, g)$ any choice of $f \in \text{Hom}_\Sigma(X_0, X)$ (see Theorem 5.13) induces a bijection

$$\text{GSp}(X_0[N]) \backslash \mathcal{F}'(X_0[N]) \simeq \text{GSp}(X[N]) \backslash \mathcal{F}'(X[N])$$

which does not depend on f . Thus it makes sense to say that an element of $\mathcal{F}'(X[N])$ belongs to the orbit of an element of $\mathcal{F}'(X_0[N])$.

We fix $x_0 \in \mathcal{F}'(X_0[N])$, denote its orbit by Ω_0 and its stabilizer in $G(\mathbb{Z}_N)$ by K_N . In particular K_N contains the kernel J_N of the surjective morphism

$$G(\mathbb{Z}_N) \longrightarrow G(\mathbb{Z}/N\mathbb{Z}) \simeq \text{GSp}(X_0[N]).$$

We generalize Definition 5.11.

Definition 5.20. Let ι_{X_0, x_0} be the map from the set of isomorphism classes of triples (X, f, x) where

- X is an object of $\mathbf{SS}_N(p, g)$,
- $f \in \Sigma^{-1} \text{Hom}_\Sigma(X_0, X)$ and
- $x \in \mathcal{F}(X)$ belongs to the orbit $\mathcal{T}_N(f)(\Omega_0)$,

to the quotient

$$(G(\mathbb{A}_\Sigma) \times G(\mathbb{Z}_N)) / (G(\mathbb{Z}_\Sigma) \times K_N),$$

mapping the class of (X, f, x) to $(\iota_{X_0}(X, f), g_N K_N)$ where ι_{X_0} is understood as an element of $G(\mathbb{A}_\Sigma)/G(\mathbb{Z}_\Sigma)$ and $g_N \in G(\mathbb{Z}_N)$ is defined by $x = \mathcal{F}(f)(g_N \cdot x_0)$.

Theorem 5.21. The map ι_{X_0, x_0} of Definition 5.20 is a bijection, and it induces a bijection $\bar{\iota}_{X_0, x_0}$ between the set of isomorphism classes of objects (X, x) in $\text{El}(\mathcal{F})$ such that x belongs to the orbit Ω_0 of x_0 and the double quotient

$$G(\mathbb{Z}[\Sigma^{-1}]) \backslash (G(\mathbb{A}_\Sigma) \times G(\mathbb{Z}_N)) / (G(\mathbb{Z}_\Sigma) \times K_N). \quad (16)$$

Proof. The proof is essentially the same as for Theorem 5.12, using Theorem 5.13 to know that $\text{Hom}_\Sigma(X_0, X)$ is inhabited. $\square$

Denoting $K^N = \prod_{\ell \nmid N} G(\mathbb{Z}_\ell)$ and $K = K^N \times K_N$ the double quotient (16) naturally embeds in the set $G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / K$ and we still denote

$$\{(X, x)\} / \simeq \xrightarrow{\bar{\iota}_{X_0, x_0}} G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / K. \quad (17)$$

Proposition 5.22. The embedding

$$G(\mathbb{Z}[\Sigma^{-1}]) \backslash (G(\mathbb{A}_\Sigma) \times G(\mathbb{Z}_N)) / (G(\mathbb{Z}_\Sigma) \times K_N) \hookrightarrow G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / K$$

is surjective, in particular the embedding (17) is a bijection.

Proof. This follows from Lemma 5.14 applied with $\Upsilon = \Sigma$ and any $q \in \Sigma$. $\square$

5.4 Automorphic forms

We continue with notation from the previous section. Let $\mathcal{A}^2(G)$ be the space of automorphic forms for G , i.e. the space of $(G(\mathbb{R}) \times K)$ -finite continuous functions $G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0} \rightarrow \mathbb{C}$ (for any compact subgroup K of $G(\mathbb{A}_f)$). The superscript 2 in $\mathcal{A}^2(G)$ is redundant and only serves to remind us that we consider $\mathcal{A}^2(G)$ as a unitary representation of $G(\mathbb{A})$. We have a decomposition into isotypic components

$$\mathcal{A}^2(G) \simeq \bigoplus_{\pi \in \Pi(G)} \pi^{\oplus m(\pi)} \quad (18)$$

ranging over a (countable) set $\Pi(G)$ of unitary irreducible admissible representations of $G(\mathbb{A})$, where each $m(\pi)$ is a positive integer. Each $\pi \in \Pi(G)$ admits a factorization as a restricted tensor product $\pi \simeq \otimes'_v \pi_v$ ranging over all places of $\mathbb{Q}$, where each π_v is a unitary irreducible admissible representation of $G(\mathbb{Q}_v)$ (trivial on $\mathbb{R}_{>0}$ for $v = \infty$).

The following proposition follows formally from Theorem 5.21 and Proposition 5.22.

Proposition 5.23. *Assume for simplicity that $\mathcal{F}(X_0)$ is finite, and choose representatives $x_1, \dots, x_r$ for the action of $\text{Aut}(\mathcal{T}_N(X_0)) \simeq \text{GSp}_{2g}(\mathbb{Z}/N\mathbb{Z})$ on $\mathcal{F}(X_0)$, with corresponding stabilizers $K_{N,1}, \dots, K_{N,r}$ in $G(\mathbb{Z}_N)$. Denote $K_i = K^N K_{N,i} \subset G(\mathbb{A}_f)$. We obtain an isomorphism (which depends on the choice of X_0 and $(x_i)_i$)*

$$\beta : \bigoplus_{i=1}^r \mathcal{A}^2(G)^{G(\mathbb{R}) \times K_i} \longrightarrow L^2(\mathcal{G}_\ell^\mathcal{F}, \mu)$$

$$(f_i)_{1 \leq i \leq r} \longmapsto ((X, x) \mapsto f_i(\bar{\iota}_{X_0, x_{i(x)}}(X, x)))$$

where $i(x)$ is defined by the property that x belongs to the orbit of $x_{i(x)}$.

Proposition 5.24. *For $1 \leq i \leq r$ endow $G(\mathbb{A}) / \mathbb{R}_{>0} = G(\mathbb{R}) / \mathbb{R}_{>0} \times G(\mathbb{A}_f)$ with the Haar measure μ_i which is the product of the Haar measure on $G(\mathbb{R}) / \mathbb{R}_{>0}$ of total volume 1 with the Haar measure on $G(\mathbb{A}_f)$ giving K_i volume 1. Let $\bar{\mu}_i$ be the quotient measure on $G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0}$ (of μ_i by the counting measure). Endow each $\mathcal{A}^2(G)^{G(\mathbb{R}) \times K_i} \subset L^2(G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0})$ with the hermitian inner product associated to this measure: $\langle f, g \rangle_i := \int_{G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0}} f(h) \overline{g(h)} d\bar{\mu}_i(h)$. Then the isomorphism β of Proposition 5.23 is an isometry.*

Proof. The proposition reduces to checking that for an object (X, x) of $\text{El}(\mathcal{F})$ such that x belongs to the orbit of x_i , denoting $\bar{\iota}_{X_0, x_i}(X, x) = G(\mathbb{Q})gK_i$ we have

$$\int_{G(\mathbb{Q}) \backslash G(\mathbb{A}) / \mathbb{R}_{>0}} 1_{G(\mathbb{Q})gK_i}(h) d\bar{\mu}_i(h) = \frac{1}{\# \text{Aut}(X, x)}. \quad (19)$$

Choosing $f \in \Sigma^{-1} \text{Hom}_\Sigma(X_0, X)$, so that we may assume $gK_i = \iota_{X_0, x_i}(X, f, x)$, we have an identification

$$\begin{aligned} \text{Aut}(X, x) &\xrightarrow{\sim} G(\mathbb{Q}) \cap gK_i g^{-1} \\ h &\longmapsto f^{-1} h f \end{aligned}$$

and (19) follows. $\square$

Now we interpret the operators $(A_\ell)_{\ell \in \Sigma}$ in this context. Fix a prime $\ell \in \Sigma$ (i.e. not dividing pN). Let $D_\ell \subset G(\mathbb{Q}_\ell)$ be the set of elements g satisfying $\nu(g) \in \ell^{-1}\mathbb{Z}_\ell^\times$ and $g^{-1} \in \mathbb{Z}_\ell \otimes_{\mathbb{Z}} \text{End}(\underline{X}_0)$. It is obviously invariant under left and right multiplication by $G(\mathbb{Z}_\ell)$.

Lemma 5.25. *The set D_ℓ is not empty and for any $x \in D_\ell$ we have $D_\ell = G(\mathbb{Z}_\ell)xG(\mathbb{Z}_\ell)$.*

Proof. Denote $L_0 = T_\ell X_0$. For $x \in D_\ell$ the lattice $L = x(L_0)$ contains L_0 and satisfies $L^\sharp = \ell L$. This implies $\ell L \subset L_0^\sharp = L_0$, and the $\mathbb{F}_\ell$ -subspace $\ell L / \ell L_0$ of $L_0 / \ell L_0$ is maximal isotropic (for $e_{\mathbb{Z}_\ell}^{\lambda_0} \bmod \ell$) and has dimension g . It is well-known that $G(\mathbb{F}_\ell)$ (on which $G(\mathbb{Z}_\ell)$ surjects by formal smoothness of $G_{\mathbb{Z}_\ell}$) acts transitively on the set of such subspaces. $\square$

Lemma 5.26. *Consider Σ -quasi-isogenies $f : X_0 \rightarrow X$ and $f' : X_0 \rightarrow Y$. Denoting $\iota_{X_0}(X, f) = (g_q G(\mathbb{Z}_q))_{q \in \Sigma}$ and $\iota_{X_0}(Y, f') = (g'_q G(\mathbb{Z}_q))_{q \in \Sigma}$, the Σ -quasi-isogeny $h := f' f^{-1} : X \rightarrow Y$ is an ℓ -isogeny if and only if*

$$g_q G(\mathbb{Z}_q) = g'_q G(\mathbb{Z}_q) \text{ for } q \neq \ell \text{ and } g_\ell^{-1} g'_\ell \in D_\ell.$$

Proof. This translation is straightforward. $\square$

Consider the Hecke operator:

$$T_\ell : \mathcal{A}^2(G)^{G(\mathbb{R}) \times K_i} \longrightarrow \mathcal{A}^2(G)^{G(\mathbb{R}) \times K_i}$$

$$f \longmapsto ([g] \mapsto \sum_{tG(\mathbb{Z}_\ell) \in D_\ell / G(\mathbb{Z}_\ell)} f([gt])).$$

It acts diagonally on $\bigoplus_{i=1}^r \mathcal{A}^2(G)^{G(\mathbb{R}) \times K_i}$.

Proposition 5.27. *The isomorphism β of Proposition 5.23 intertwines A_ℓ and T_ℓ , i.e. we have*

$$A_\ell \circ \beta = \beta \circ T_\ell.$$

Proof. Recall

$$A_\ell(\beta((f_i)_i))([X, x]) = \sum_{h : (X, x) \rightarrow (Y, y)} \beta((f_i)_i)([Y, y])$$

where the sum is over equivalence classes of morphisms in $\text{El}(\mathcal{F})$ of reduced degree $c(h)$ equal to ℓ . Choose a Σ -quasi-isogeny $f : X_0 \rightarrow X$ and let i be the index such that x belongs to the orbit of $\mathcal{F}(f)(x_i)$ and denote $\iota_{X_0, x_i}(X, f, x) = gG(\mathbb{Z}_\Sigma)K_{N, i} \in G(\mathbb{A}_\Sigma)G(\mathbb{Z}_N)/G(\mathbb{Z}_\Sigma)K_{N, i}$. It follows from Lemma 5.26 that we have a bijection $h \mapsto \iota_{X_0, x_i}(Y, h \circ f, y)$ between equivalence classes of morphisms $h : (X, x) \rightarrow (Y, y)$ in $\text{El}(\mathcal{F})$ of reduced degree ℓ and classes $g'G(\mathbb{Z}_\Sigma)K_{N, i} \in G(\mathbb{A}_\Sigma)G(\mathbb{Z}_N)/G(\mathbb{Z}_\Sigma)K_{N, i}$ satisfying

1. for all $q \in \Sigma \setminus \{\ell\}$ we have $g'_q G(\mathbb{Z}_q) = g_q G(\mathbb{Z}_q)$,
2. $g_\ell^{-1} g'_\ell \in D_\ell$,
3. denoting $g_N = (g_q)_{q|N}$ and similarly for g'_N , we have $g_N K_{N,i} = g'_N K_{N,i}$.

Writing $g' = gt$ where $t \in D_\ell \times \{1\} \subset G(\mathbb{A}_\Sigma) \times G(\mathbb{Z}_N)$ we see that the set of such classes $g'G(\mathbb{Z}_\Sigma)K_{N,i}$ is in bijection with $D_\ell/G(\mathbb{Z}_\ell)$, and the intertwining relation follows. $\square$

Corollary 5.28. *The operator A_ℓ on $L^2(\mathcal{G}^\mathcal{F}, \mu)$ is normal.*

Proof. By Propositions 5.27 and 5.24 it is enough to show that T_ℓ is normal. The action of $G(\mathbb{A}_f)$ on $\mathcal{A}^2(G)^{G(\mathbb{R})}$ (the space of functions $G(\mathbb{Q}) \backslash G(\mathbb{A}_f) \rightarrow \mathbb{C}$ which are right K -invariant for some compact open subgroup K of $G(\mathbb{A}_f)$) is unitary, so the adjoint T_ℓ^* is the Hecke operator associated to $D_\ell^{-1} := \{x^{-1} \mid x \in D_\ell\}$. Via an isomorphism $G_{\mathbb{Z}_\ell} \simeq \mathbf{GSp}_{2g, \mathbb{Z}_\ell}$ (obtained by trivializing $T_\ell X_0$) the set D_ℓ is identified with

$$G(\mathbb{Z}_\ell) \begin{pmatrix} \ell^{-1} I_g & 0 \\ 0 & I_g \end{pmatrix} G(\mathbb{Z}_\ell) = G(\mathbb{Z}_\ell) \begin{pmatrix} I_g & 0 \\ 0 & \ell^{-1} I_g \end{pmatrix} G(\mathbb{Z}_\ell).$$

We deduce $D_\ell^{-1} = \ell D_\ell$, and so $T_\ell^* = T_\ell \langle \ell \rangle$ where the “diamond” operator $\langle \ell \rangle$ on $\mathcal{A}^2(G)^{G(\mathbb{R})}$ is right translation by $\ell \in G(\mathbb{Q}_\ell)$. Since this element is central in $G(\mathbb{Q}_\ell)$ we have $T_\ell \langle \ell \rangle = \langle \ell \rangle T_\ell$ and so T_ℓ^* commutes with T_ℓ . $\square$

In the decomposition (18) we are only interested in representations $\pi \in \Pi(G)$ satisfying $\pi_\infty \simeq \mathbb{C}$ (the trivial representation of $G(\mathbb{R})/\mathbb{R}_{>0}$) and $\pi^{K_i} \neq 0$ for some $1 \leq i \leq r$. In this case $\pi_\ell^{G(\mathbb{Z}_\ell)}$ is a line on which T_ℓ acts by a scalar (in other words π_ℓ is unramified).

Let $\mathcal{A}_{\text{deg}}^2(G) \subset \mathcal{A}^2(G)$ be the subspace of functions factoring through the map

$$G(\mathbb{Q}) \backslash G(\mathbb{A})/\mathbb{R}_{>0} \xrightarrow{\nu} \mathbb{Q}^\times \backslash \mathbb{A}^\times/\mathbb{R}_{>0} \simeq \widehat{\mathbb{Z}}^\times, \quad (20)$$

clearly a sub- $G(\mathbb{A})$ -representation.

Lemma 5.29. *The map (20) is surjective, and two elements of $G(\mathbb{Q}) \backslash G(\mathbb{A})/\mathbb{R}_{>0}$ have the same image if and only if they differ by right multiplication by an element of $G^1(\mathbb{A})$.*

Proof. We have already seen in the proof of Lemma 5.14: $\nu(G(\mathbb{A}_f)) = \mathbb{A}_f^\times$, $\nu(G(\mathbb{R})) = \mathbb{R}_{>0}$, $\nu(G(\mathbb{Q})) = \mathbb{Q}_{>0}$, so surjectivity follows from the fact that the natural map $\mathbb{Q}_{>0} \backslash (\mathbb{A}_f^\times \times \mathbb{R}_{>0})/\mathbb{R}_{>0} \rightarrow \mathbb{Q}^\times \backslash \mathbb{A}^\times/\mathbb{R}_{>0}$ is bijective. The second statement easily follows from the equalities $\nu(G(\mathbb{R})) = \mathbb{R}_{>0}$ and $\nu(G(\mathbb{Q})) = \mathbb{Q}_{>0}$. $\square$

The second part of the lemma implies that $\mathcal{A}_{\text{deg}}^2(G)$ is also the subspace of $\mathcal{A}^2(G)$ consisting of (right) $G^1(\mathbb{A})$ -invariant functions. The isotypic decomposition of $\mathcal{A}_{\text{deg}}^2(G)$ is straightforward:

$$\mathcal{A}_{\text{deg}}^2(G) \simeq \bigoplus_{\chi} \mathbb{C}_\chi$$

where the sum ranges over all smooth characters $\chi = \prod_v \chi_v$ of the profinite group $\mathbb{Q}^\times \backslash \mathbb{A}^\times / \mathbb{R}_{>0}$ and $\mathbb{C}_\chi \subset \mathcal{A}_{\deg}^2(G)$ is the line of $\chi \circ \nu$ -equivariant functions. In particular all eigenvalues of T_ℓ acting on $\mathcal{A}_{\deg}^2(G)$ are of the form $\zeta \prod_{i=1}^g 1 + \ell^i$ where $\zeta = \chi_\ell(\ell)^{-1}$ is a root of unity.

Proposition 5.30. *The isomorphism β identifies $\bigoplus_i \mathcal{A}_{\deg}^2(G)^{G(\mathbb{R}) \times K_i}$ with $L_{\deg}^2(\mathcal{G}^\mathcal{F}, \mu)$.*

Proof. For $1 \leq i \leq r$ and Σ -quasi-isogenies $f : X_0 \rightarrow X$ (resp. $f' : X_0 \rightarrow Y$) such that $x \in \mathcal{F}(X)$ (resp. $y \in \mathcal{F}(Y)$) belongs to the orbit of $\mathcal{F}(f)(a_i)$ (resp. $\mathcal{F}(f')(a_i)$), corresponding to $gK_i^{(p)} \in G(\mathbb{A}_\Sigma) \times K_{N,i}/K_i^{(p)}$ (resp. $g'K_i^{(p)}$) we are interested in understanding when the following conditions are satisfied:

1. The quasi-isogeny $f'f^{-1}$ is an isogeny,
2. its reduced degree $\text{rdeg}(f'f^{-1}) \equiv 1 \pmod{N}$,
3. $\mathcal{F}(f'f^{-1})(x) = y$.

These translate to

1. for all $\ell \in \Sigma$ we have $(g'_\ell)^{-1}g_\ell \in \mathbb{Z}_\ell \otimes \text{End}(\underline{X}_0)$,
2. $\prod_{\ell \in \Sigma} \ell^{v_\ell(\nu(g_\ell^{-1}g'_\ell))} \equiv 1 \pmod{N}$,
3. $g_N K_N = g'_N K_N$.

First consider $(f_i)_i \in \bigoplus_i \mathcal{A}_{\deg}^2(G)^{G(\mathbb{R}) \times K_i}$, i.e. each f_i factors through ν , and consider two objects (X, x) and (Y, y) of $\text{El}(\mathcal{F})$ and $h \in \text{Hom}_\Sigma(X, Y)$ such that $\mathcal{F}(h)(x) = y$ and $c(h) \equiv 1 \pmod{N}$. Choose a Σ -quasi-isogeny $f : X_0 \rightarrow X$, let i be the index such that x belongs to the orbit of $\mathcal{F}(f)(x_i)$, and let $f' = h \circ f$. The three conditions above are then satisfied by $gK_i^{(p)} := \iota_{X_0, x_i}(X, f, x)$ and $g'K_i^{(p)} := \iota_{X_0, x_i}(Y, f', y)$. The classes of $\nu(g)$ and $\nu(g')$ in $\mathbb{A}_f^\times / \nu(K_i)$ differ by $\prod_{\ell \in \Sigma} \ell^{v_\ell(\nu(g_\ell^{-1}g'_\ell))} \in \mathbb{Q}^\times$ thanks to the second and third conditions, so we conclude

$$\beta((f_i)_i)([X, x]) = f_i([g]) = f_i([g']) = \beta((f_i)_i)([Y, y]).$$

Conversely, assume that $(f_i)_i \in \bigoplus_i \mathcal{A}_{\deg}^2(G)^{G(\mathbb{R}) \times K_i}$ is such that $\beta((f_i)_i)$ is constant on equivalence classes, and consider $1 \leq i \leq r$ and $[g], [g'] \in G(\mathbb{Q}) \backslash G(\mathbb{A}_f) / K_i$ having the same image under ν . We can write $[g] = \bar{\iota}_{X_0, x_i}(X, x)$ and $[g'] = \bar{\iota}_{X_0, x_i}(Y, y)$. Choose a Σ -quasi-isogeny $f : X_0 \rightarrow X$, which amounts to choose $gK_i^{(p)} \in G(\mathbb{A}_\Sigma) \times K_{N,i}/K_i^{(p)}$ representing $[g]$ (namely $gK_i = \iota_{X_0, x_i}(X, f, x)$). The assumption $\nu([g]) = \nu([g'])$ implies the existence of a Σ -quasi-isogeny $f' : X_0 \rightarrow Y$ such that $g'K_i^{(p)} := \iota_{X_0, x_i}(Y, f', y)$ satisfies $\nu(gK_i^{(p)}) = \nu(g'K_i^{(p)})$ in $(\mathbb{A}_f^{(p)})^\times / \mathbb{Z}_\Sigma^\times \nu(K_{N,i})$, in particular the second condition above is satisfied. By strong approximation for any $\ell \in \Sigma$ the group $G^1(\mathbb{Z}[\ell^{-1}])$ is dense in $G(\mathbb{Z}_N)$, so up to replacing g' by an element of $G^1(\mathbb{Z}[\Sigma^{-1}])g'$ we may also arrange for $g_N K_N = g'_N K_N$. Let S be the (finite) set of $\ell \in \Sigma$ such that the first condition is not satisfied. Replacing g' by $\prod_{\ell \in S} \ell^{-k\phi(N)} g'$ for a large enough integer k , where $\phi(N) = \#(\mathbb{Z}/N\mathbb{Z})^\times$, allows us to assume that the first condition is satisfied (without changing the other two conditions). We obtain $f_i([g]) = f_i([g'])$, as desired. $\square$

Let $\mathcal{A}_0^2(G)$ be the orthogonal of $\mathcal{A}_{\deg}^2(G)$ in $\mathcal{A}^2(G)$, also a sub- $G(\mathbb{A})$ -representation. The following argument is well known.

Lemma 5.31. *For $\pi \in \Pi(G)$ if π_ℓ has non-zero invariants under $G^1(\mathbb{Q}_\ell)$ then the π -isotypic component of $\mathcal{A}^2(G)$ is contained in $\mathcal{A}_{\deg}^2(G)$.*

Proof. Since $G^1(\mathbb{Q}_\ell)$ is a normal subgroup of $G(\mathbb{Q}_\ell)$ this isotypic component consists of $G^1(\mathbb{Q}_\ell)$ -invariant functions, and by strong approximation for G^1 any continuous $G^1(\mathbb{Q}_\ell)$ -invariant function on $G(\mathbb{Q}) \backslash G(\mathbb{A})$ is even $G^1(\mathbb{A})$ -invariant. $\square$

The lemma implies in particular that each $\pi \in \Pi(G)$ appears either in $\mathcal{A}_{\deg}^2(G)$ or in $\mathcal{A}_0^2(G)$, not in both.

Summary. We have proved all statements in Theorem 4.1 and Proposition 4.2 except for the bounds on non-trivial eigenvalues in Theorem 4.1:

- The first and second points in Theorem 4.1 follow from Propositions 5.27 and 5.30.
- The first point of Proposition 4.2 follows from Proposition 5.19 and the fact that Alt_N is a groupoid.
- The second point of Proposition 4.2 follows from Lemma 5.14 applied with $\Upsilon = \Sigma$ and any $q \in \Sigma$.
- The third point of Proposition 4.2 follows from Proposition 5.30.
- The fourth point of Proposition 4.2 follows from Propositions 5.30 and 2.14.

5.5 Bounds from Oh's theorem

We specialize Oh's theorem [Oh02, Theorem 1.1] to unramified unitary irreducible representations of $\text{GSp}_{2g}(\mathbb{Q}_\ell)$.

Theorem 5.32. *Let $g > 1$ be an integer, and let (V, ρ) be an unramified unitary irreducible smooth representation of $\text{GSp}_{2g}(\mathbb{Q}_\ell)$. Assume $\dim_{\mathbb{C}} V > 1$. Then the eigenvalue x of the Hecke operator T_ℓ acting on the line $V^{\text{GSp}_{2g}(\mathbb{Z}_\ell)}$ satisfies*

$$|x| \leq \prod_{i=1}^g (1 + \ell^i) \times \left(\frac{2\sqrt{\ell}}{\ell + 1} \right)^g = N_g(\ell) \left(\frac{2\sqrt{\ell}}{\ell + 1} \right)^g. \quad (21)$$

Proof. Let $\mathcal{S}$ be a maximal strongly orthogonal system, e.g., $\mathcal{S} = \{2e_1, \dots, 2e_g\}$ in standard notation (see the appendix to [Oh02], case C_n where $n = g$), and let Ξ_{PGL_2} be the Harish-Chandra function of PGL_2 (see §2.2 of [Oh02]). Then it follows from [Oh02, Theorem 1.1] that

$$|(\rho(a)v, w)| \leq \prod_{\alpha \in \mathcal{S}} \Xi_{\text{PGL}_2}(\alpha(a)) \quad (22)$$

where $a \in A^+$, $(v, v) = (w, w) = 1$. We also obtain from 3.8 of [Oh02] that

$$\Xi_{\mathrm{PGL}_2}(\ell^{-k}) = \Xi_{\mathrm{PGL}_2}(\ell^k) = \ell^{-k/2} \frac{k(\ell-1) + \ell + 1}{\ell + 1} \text{ for } k \geq 0. \quad (23)$$

Since $\alpha(a) = \ell^\pm$, for each $\mathcal{S}$, we have $\prod_{\alpha \in \mathcal{S}} \Xi_{\mathrm{PGL}_2}(\alpha(a)) = \left(\frac{2\sqrt{\ell}}{\ell+1}\right)^g$ from (23). We take $a = \mathrm{diag}(\ell^{-1}I_g, I_g)$. The right hand side of (22) simplifies to $(2\ell^{1/2}/(\ell+1))^g$. Denote $K = \mathrm{GSp}_{2g}(\mathbb{Z}_\ell)$. We are interested in the eigenvalue x of the Hecke operator $[KaK]$ acting on the line V^K . Write the decomposition $KaK = \bigsqcup_{i \in I} h_i K$, where I is the index set so that $\#I = \prod_{i=1}^g (1 + \ell^i) = N_g(\ell)$.

Since V^K is one-dimensional, we have

$$xv = \rho([KaK])v = \mathrm{vol}(K)^{-1} \int_{KaK} \rho(h)vdh.$$

Also, since $(\rho(k_1 a k_2)v, w) = (\rho(a)v, w)$ for any $k_1, k_2 \in K$ by the definition of V^K , the integrand is constant on KaK , and so, for a unit vector v and $h \in KaK$

$$\begin{aligned} x &= x(v, v) = (v, xv) = (v, \mathrm{vol}(K)^{-1} \int_{KaK} \rho(h)vdh) \\ &= \mathrm{vol}(K)^{-1} (v, \int_{KaK} \rho(h)vdh) = \mathrm{vol}(K)^{-1} \mathrm{vol}(KaK) (v, \rho(a)v). \end{aligned}$$

As $\mathrm{vol}(K)^{-1} \mathrm{vol}(KaK)$ is the number of the cosets, it follows that $x = N_g(\ell)(\rho(a)v, v)$, and the assertion follows from (22). $\square$

Theorem 5.32 may now be applied to the component at ℓ of the irreducible constituents of $\mathcal{A}_0^2(G)$, thanks to Lemma 5.31. Translating via Propositions 5.23, 5.24, 5.27 and 5.30 we obtain the bound in point 3 of Theorem 4.1.

One can see from the proof of Theorem 5.32 that one cannot obtain a better upper bound for the eigenvalue $|x|$ than the one in (21) by applying Oh's Theorem. Therefore, we will use a different technique to get a better bound in the following section.

5.6 Refined bounds using Arthur's endoscopic classification

We now use Arthur's endoscopic classification for classical groups and known cases of the Ramanujan–Petersson conjecture to deduce point 4 of Theorem 4.1.

Let (V, π) be an irreducible sub- $G(\mathbb{A})$ -representation of $\mathcal{A}^2(G)$ which is not contained in $\mathcal{A}_{\mathrm{deg}}^2(G)$ (see Lemma 5.31). It decomposes as a restricted tensor product of $(V_v, \pi_v)_v$, and we consider a prime $\ell \neq p$ such that π_ℓ is unramified, i.e. $V_\ell^{G(\mathbb{Z}_\ell)} \neq 0$, in which case this space is a line, on which T_ℓ acts by a scalar $\pi_\ell(T_\ell) \in \mathbb{C}$. We want to estimate this eigenvalue. In terms of the Satake parameter $c(\pi_\ell)$ of π_ℓ (a semi-simple conjugacy class in the Langlands dual group $\widehat{G}(\mathbb{C}) \simeq \mathrm{GSpin}_{2g+1}(\mathbb{C})$) we have (see [Gro98, (3.15)])

$$\pi_\ell(T_\ell) = \omega_{\pi_\ell}(\ell)^{-1} \ell^{g(g+1)/4} \mathrm{Tr}(\mathrm{spin}(c(\pi_\ell)))$$

where ω_{π_ℓ} is the central character of π_ℓ and $\text{spin} : \widehat{G}(\mathbb{C}) \rightarrow \text{GL}_{2g}(\mathbb{C})$ is the spin representation. As ω_{π_ℓ} is the component at ℓ of a continuous character $\mathbb{Q}^\times \backslash \mathbb{A}^\times / \mathbb{R}_{>0}$, it has finite order and $|\omega_{\pi_\ell}(\ell)| = 1$.

Recall $G^1 = \ker(G \rightarrow \mathbf{GL}_{1,\mathbb{Z}})$. Let $\mathcal{A}^2(G^1)$ be the space of square-integrable automorphic forms $G^1(\mathbb{Q}) \backslash G^1(\mathbb{A}) \rightarrow \mathbb{C}$. We have a restriction map

$$\text{res}_{G^1}^G : \mathcal{A}^2(G) \longrightarrow \mathcal{A}^2(G^1).$$

The image of V is non-zero, and as a representation of $G^1(\mathbb{A})$ it is a quotient of $\otimes'_v \text{res}_{G^1(\mathbb{Q}_v)}^{G(\mathbb{Q}_v)}(V_v, \pi_v)$. Each restriction $\text{res}_{G^1(\mathbb{Q}_v)}^{G(\mathbb{Q}_v)}(V_v, \pi_v)$ is semi-simple and has finite length, so $\text{res}_{G^1}^G V$ is semi-simple as well.

Lemma 5.33. *Any irreducible sub- $G^1(\mathbb{A})$ -representation (V', π') of $\text{res}_{G^1}^G V$ is unramified at ℓ : there exists a hyperspecial compact subgroup K of $G^1(\mathbb{Q}_\ell)$ (i.e. $gG^1(\mathbb{Z}_\ell)g^{-1}$ for some $g \in G(\mathbb{Q}_\ell)$) satisfying $(V'_\ell)^{\pi'_\ell(K)} \neq 0$. Moreover the Satake parameter of (V'_ℓ, π'_ℓ) is equal to the image of that of (V_ℓ, π_ℓ) under $\widehat{G} \rightarrow \widehat{G}^1$.*

Proof. Let V'_ℓ be the irreducible sub- $G^1(\mathbb{Q}_\ell)$ -representation of $\text{res}_{G^1(\mathbb{Q}_\ell)}^{G(\mathbb{Q}_\ell)}(V_\ell, \pi_\ell)$ containing the line $V_\ell^{G(\mathbb{Z}_\ell)}$. (There is a unique such representation because V_ℓ embeds in an unramified principal series representation, whose space of $G^1(\mathbb{Z}_\ell)$ -invariants is one-dimensional as well.) We have $V_\ell = \sum_{g \in G(\mathbb{Q}_\ell)} \pi_\ell(g) V'_\ell$ and each term is an irreducible representation of $G^1(\mathbb{Q}_\ell)$, so there exists a finite subset $F_\ell \subset G(\mathbb{Q}_\ell)$ such that $\text{res}_{G^1(\mathbb{Q}_\ell)}^{G(\mathbb{Q}_\ell)}(V_\ell, \pi_\ell) \simeq \bigoplus_{g \in F_\ell} \pi_\ell(g) V'_\ell$. Note that each irreducible factor $\pi_\ell(g) V'_\ell$ is unramified (for some hyperspecial maximal compact subgroup of $G^1(\mathbb{Q}_\ell)$): we have $(\pi_\ell(g) V'_\ell)^{gG^1(\mathbb{Z}_\ell)g^{-1}} \neq 0$.

The second point follows from [Sat63, Proposition 7.1]. $\square$

We recall the main invariants associated to a cuspidal automorphic representation $\sigma \simeq \otimes'_v \sigma_v$ for $\mathbf{GL}_{m,\mathbb{Q}}$. At the real place the infinitesimal character $c(\sigma_\infty)$ is a semi-simple conjugacy class in $\mathfrak{gl}_m(\mathbb{C})$ (characterized by its characteristic polynomial, or its multi-set of eigenvalues). At each finite place (i.e., prime) ℓ where σ_ℓ is unramified (i.e., the space of invariants under $\mathbf{GL}_m(\mathbb{Z}_\ell)$ is non-zero) the Satake parameter $c(\sigma_\ell)$ is a semi-simple conjugacy class in $\text{GL}_m(\mathbb{C})$.

Theorem 5.34. *Let (V', π') be an irreducible representation of $G^1(\mathbb{A})$ occurring in $\mathcal{A}^2(G^1)$. There exists a unique formal Arthur parameter $\psi = \bigoplus_i \pi_i[d_i]$ (each π_i is a self-dual cuspidal automorphic representation for $\mathbf{GL}_{n_i,\mathbb{Q}}$, $d_i \geq 1$ is an integer and $\sum_i n_i d_i = 2g + 1$) such that for any $\ell \neq p$ at which (V', π') is unramified (for some hyperspecial subgroup of $G^1(\mathbb{Q}_\ell)$) each $\pi_{i,\ell}$ is unramified and the image of the Satake parameter $c(\pi'_\ell)$ under the standard representation $\text{Std}_{\widehat{G}^1} : \widehat{G}^1 \rightarrow \text{GL}_{2g+1}(\mathbb{C})$ is equal to*

$$\bigoplus_i c(\pi_{i,\ell}) \otimes \text{diag}(\ell^{(d_i-1)/2}, \dots, \ell^{(1-d_i)/2}).$$

Moreover the image of the infinitesimal character $c(\pi'_\infty)$ (seen as a semi-simple conjugacy class in the Lie algebra $\widehat{\mathfrak{g}}^1$ of $\widehat{G}^1$) under the differential of $\text{Std}_{\widehat{G}^1}$ is equal to

$$\bigoplus_i c(\pi_{i,\infty}) \otimes \text{diag}\left(\frac{d_i-1}{2}, \dots, \frac{1-d_i}{2}\right).$$

Proof. It follows similarly as in [Che22, App. A]; see also [Shi24]. $\square$

Remark 5.35. Theorem 5.34 relies on Arthur’s endoscopic classification [Art13], which depends on the stabilization of the twisted trace formula ([MW16a], [MW16b]) and so remains conditional on the “twisted weighted fundamental lemma” [MW16a, §II.4.4]. Recently Connor Halleck-Dubé announced a proof (generalizing work of Chaudouard-Laumon [CL12]) of this “lemma”. Thus, one can be hopeful that the results of [Art13] will soon become unconditional.

Theorem 5.36. *In the situation of Theorem 5.34, for any index i and any prime ℓ at which $\pi_{i,\ell}$ is unramified, all eigenvalues of the Satake parameter $c(\pi_{i,\ell})$ have absolute value 1.*

Proof. Since $G^1(\mathbb{R})$ is a compact Lie group the infinitesimal character $c(\pi'_\infty)$ is algebraic and regular: more concretely its eigenvalues in the standard representation of $\widehat{\mathfrak{g}}^1 \simeq \mathfrak{so}_{2g+1}(\mathbb{C})$ are $2g+1$ distinct integers (symmetric around 0). The compatibility of infinitesimal characters in Theorem 5.34 then implies that either π_i or $\pi_i|\det|^{1/2}$ is algebraic regular (in the sense of [Clo88, Definitions 1.8 and 3.12]). So Theorem 5.36 follows from [Clo13] or [Car12]. $\square$

Lemma 5.37. *In the situation of Theorem 5.34 if (V', π') occurs in $\text{res}_{G^1}^G \mathcal{A}_0^2(G)$ then the parameter ψ is non-trivial, i.e. it is not equal to $\pi_0[2g+1]$ where π_0 is the trivial automorphic representation for $\mathbf{GL}_{1,\mathbb{Q}}$.*

Proof. We may fix an irreducible subrepresentation (V, π) of $\mathcal{A}_0^2(G)$ and consider an irreducible constituent (V', π') of $\text{res}_{G^1}^G V$. If its Arthur–Langlands parameter is trivial then for almost all primes q we have that π'_q is the trivial representation of $G^1(\mathbb{Q}_q)$, in particular there exists $f \in V \setminus \{0\}$ such that $\text{res}_{G^1}^G(f)$ is non-zero and $G^1(\mathbb{Q}_q)$ -invariant. By irreducibility of (V, π) the $G(\mathbb{A})$ -translates of f generate V (as a complex vector space), and we see that all elements of V are $G^1(\mathbb{Q}_q)$ -invariant, contradicting Lemma 5.31. $\square$

We now require explicit parametrizations of maximal tori of $\widehat{G}$ and $\widehat{G}^1$. There is a well-known parametrization $\text{GL}_1^g \simeq \mathcal{T}_{\widehat{G}^1}$ of a maximal torus $\mathcal{T}_{\widehat{G}^1}$ of $\widehat{G}^1$ where the eigenvalues of $(t_1, \dots, t_g)$ in the standard representation of $\widehat{G}^1$ are $t_1^{\pm 1}, \dots, t_g^{\pm 1}$ and 1. There is a parametrization of the corresponding maximal torus of $\widehat{G}$

$$\mathcal{T}_{\widehat{G}} \simeq \{(t_1, \dots, t_g, s, \lambda) \in (\mathbb{C}^\times)^{g+2} \mid s^2 = t_1 \dots t_g\} / \langle (1, \dots, 1, -1, -1) \rangle$$

such that the natural projection $\mathcal{T}_{\widehat{G}} \rightarrow \mathcal{T}_{\widehat{G}^1}$ is $[t_1, \dots, t_g, s, \lambda] \mapsto (t_1, \dots, t_g)$ and the eigenvalues of $[t_1, \dots, t_g, s, \lambda]$ in the spin representation of $\widehat{G}$ are the 2^g complex numbers $(\lambda s / \prod_{j \in J} t_j)_{J \subset \{1, \dots, g\}}$.

Going back to our unramified irreducible representation π_ℓ of $G(\mathbb{Q}_\ell)$ (the factor at ℓ of an irreducible constituent of $\mathcal{A}_0^2(G)$), its Satake parameter $c(\pi_\ell)$ may be represented by $[t_1, \dots, t_g, s, \lambda] \in \mathcal{T}_{\widehat{G}}$. Considering the central character of π , that we can see as a character of the compact group $\mathbb{A}^\times/\mathbb{R}_{>0}$, we have $|\lambda| = 1$. Note that $|\pi_\ell(T_\ell)|$ only depends on $c(\pi'_\ell)$ (represented by $(t_1, \dots, t_g) \in \mathcal{T}_{\widehat{G}^1}$) as:

$$|\pi_\ell(T_\ell)| = \ell^{g(g+1)/4} |\mathrm{Tr}(\mathrm{spin}(c(\pi'_\ell)))| = \ell^{g(g+1)/4} \prod_{i=1}^g |t_i|^{1/2} |1 + t_i^{-1}|.$$

Write $\psi = \pi_0[d_0] \oplus \dots \oplus \pi_r[d_r]$. There is a unique index i such that $n_i d_i$ is odd, say $i = 0$. If there exists $i > 0$ such that d_i is odd (i.e. π_i is of orthogonal type) and $\det(c(\pi_{i,\ell})) = -1$ then $\mathrm{Std}_{\widehat{G}^1}(c(\pi'_\ell))$ admits -1 as an eigenvalue and we have $\mathrm{Tr}(\mathrm{spin}(c(\pi'_\ell))) = 0$. We may thus assume there is no such $i > 0$, in which case we have $\det(c(\pi_{0,\ell})) = +1$ as $\mathrm{Std}_{\widehat{G}^1}(c(\pi'_\ell))$ has determinant 1. Thus:

- for $i > 0$ the eigenvalues of $c(\pi_{i,\ell})$ are of the form $(\alpha_{i,1}^{\pm 1}, \dots, \alpha_{i,n_i/2}^{\pm 1})$,
- the eigenvalues of $c(\pi_{0,\ell})$ are of the form $(\alpha_{0,1}^{\pm 1}, \dots, \alpha_{0,(n_0-1)/2}^{\pm 1}, 1)$.

For the rest of the argument we only consider the local situation at ℓ . Up to reindexing, we may represent $c(\pi'_\ell)$ by $(t_1, \dots, t_g) \in \mathcal{T}_{\widehat{G}^1}$ with an equality of multi-sets

$$\{t_1, \dots, t_g\} = \{\ell^j \mid 1 \leq j \leq (d_0 - 1)/2\} \cup \bigcup_{1 \leq i \leq r} \{\alpha_i \ell^{(d_i-1)/2-j} \mid 0 \leq j \leq d_i - 1\}$$

where $|\alpha_i| = 1$ by Theorem 5.36. Now the integers $(d_i)_{0 \leq i \leq r}$ satisfy $2g + 1 = d_0 + 2 \sum_{i=1}^r d_i$ and Lemma 5.37 translates to $r > 0$. We compute

$$\begin{aligned} |\mathrm{Tr}(\mathrm{spin}(c(\pi'_\ell)))| &= \left(\prod_{j=1}^{(d_0-1)/2} \ell^{j/2} + \ell^{-j/2} \right) \times \left(\prod_{i=1}^r |\alpha_i|^{d_i/2} \prod_{j=(1-d_i)/2}^{(d_i-1)/2} |1 + \alpha_i^{-1} \ell^j| \right) \\ &\leq \left(\prod_{j=1}^{(d_0-1)/2} \ell^{j/2} + \ell^{-j/2} \right) \times \left(\prod_{i=1}^r \prod_{j=(1-d_i)/2}^{(d_i-1)/2} 1 + \ell^j \right) \end{aligned}$$

where the products $\prod_{j=(1-d_i)/2}^{(d_i-1)/2}$ range over $j \in \frac{d_i-1}{2} + \mathbb{Z}$ satisfying $|j| \leq (d_i-1)/2$. The last step in the proof of point 4 of Theorem 4.1 is the following lemma.

Lemma 5.38. *Recall that the family of positive integers $(d_i)_{0 \leq i \leq r}$ satisfies $r > 0$ and $2g + 1 = d_0 + 2 \sum_{i=1}^r d_i$. We have*

$$\frac{\left(\prod_{j=1}^{(d_0-1)/2} \ell^{j/2} + \ell^{-j/2} \right) \times \left(\prod_{i=1}^r \prod_{j=(1-d_i)/2}^{(d_i-1)/2} 1 + \ell^j \right)}{\prod_{j=1}^g \ell^{j/2} + \ell^{-j/2}} \leq \frac{2}{\ell^{g/2} + \ell^{-g/2}}.$$

Proof. Consider the multi-sets $A := \{j \in \mathbb{Z} \mid \frac{d_0+1}{2} \leq j \leq g\}$ and

$$B := \{|j| \mid 1 \leq i \leq r, j \in \frac{d_i-1}{2} + \mathbb{Z}, |j| \leq \frac{d_i-1}{2}\}.$$

They have the same cardinality $k := g - \frac{d_0-1}{2} = \sum_{i=1}^r d_i$, and we have

$$\frac{\left(\prod_{j=1}^{(d_0-1)/2} \ell^{j/2} + \ell^{-j/2}\right) \times \left(\prod_{i=1}^r \prod_{j=(1-d_i)/2}^{(d_i-1)/2} 1 + \ell^j\right)}{\prod_{j=1}^g \ell^{j/2} + \ell^{-j/2}} = \frac{\prod_{j \in B} \ell^{j/2} + \ell^{-j/2}}{\prod_{j \in A} \ell^{j/2} + \ell^{-j/2}}.$$

Write the second multi-set as $\{j_1 \leq \dots \leq j_k\}$. We have $j_{a+1} \leq j_a + 1$ for all $1 \leq a < k$. Note that the function $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ defined by $x \mapsto \ell^{x/2} + \ell^{-x/2}$ is increasing. Assume first that there exists $i > 0$ for which d_i is odd. In this case we have $j_1 = 0$ and $j_a \leq a - 1$ so we conclude

$$\frac{\prod_{j \in B} f(j)}{\prod_{j \in A} f(j)} = \frac{f(j_1)}{f(g)} \frac{\prod_{a=2}^k f(j_a)}{\prod_{j=(d_0+1)/2}^{g-1} f(j)} = \frac{2}{\ell^{g/2} + \ell^{-g/2}} \frac{\prod_{a=2}^k f(j_a)}{\prod_{j=(d_0+1)/2}^{g-1} f(j)} \leq \frac{2}{\ell^{g/2} + \ell^{-g/2}}.$$

We are left to consider the case where all $(d_i)_{1 \leq i \leq r}$ are even. Here, we have $k \geq 2$, $j_1 = j_2 = \frac{1}{2}$ and $j_a \leq \frac{3}{2} + a - 3 \leq a - 1$ for $3 \leq a \leq k$. We decompose

$$\frac{\prod_{j \in B} f(j)}{\prod_{j \in A} f(j)} = \frac{f(1/2)^2}{f((d_0+1)/2)f(g)} \frac{\prod_{a=3}^k f(j_a)}{\prod_{j=(d_0+3)/2}^{g-1} f(j)}$$

and conclude by

$$\frac{\prod_{a=3}^k f(j_a)}{\prod_{j=(d_0+3)/2}^{g-1} f(j)} \leq \frac{\prod_{j=2}^k f(j)}{\prod_{j=(d_0+3)/2}^{g-1} f(j)} \leq 1$$

$$\text{and } \frac{f(1/2)^2}{f((d_0+1)/2)f(g)} \leq \frac{f(1/2)^2}{f(1)f(g)} = \frac{f(0)+f(1)}{f(1)f(g)} < \frac{2f(1)}{f(1)f(g)} = \frac{2}{\ell^{g/2} + \ell^{-g/2}}. \quad \square$$

5.7 Comparison with previous bounds

We can compare the bounds we have obtained with two different techniques and the bound given by [ATY24] in the case that there is no level structure as we have already discussed in the introduction. To be more precise, the bound for the spectral gap $\lambda(\mathcal{G}_\ell(p, g))$ was given in [ATY24] is:

$$\frac{\lambda(\mathcal{G}_\ell(p, g))}{\deg(\mathcal{G}_\ell(p, g))} \leq 1 - \frac{1}{4(g+2)} \cdot \left(\frac{\ell-1}{2(\ell-1) + 3\sqrt{2\ell(\ell+1)}} \right)^2.$$

In Table 1, we compare that with our two bounds in Theorem 4.1: $(2\sqrt{\ell}/(\ell+1))^g$ and $2/(\ell^{g/2} + \ell^{-g/2})$.

Remark 5.39. For a certain g and ℓ , we can get a sharper bound by following the proof of Lemma 5.38. To illustrate this, let us consider the graph $\mathcal{G}_2(p, 2)$ (i.e., $g = \ell = 2$). In this case, we have from the last inequality of the proof that:

$$\frac{f(1/2)^2}{f((d_0+1)/2)f(g)} = \frac{f(1/2)^2}{f((1+1)/2)f(2)} = \frac{(2^{1/4} + 2^{-1/4})^2}{(2^{1/2} + 2^{-1/2})(2 + 2^{-1})} = \frac{6 + 4\sqrt{2}}{15}.$$

Thus, we obtain the following upper bound: $\lambda(\mathcal{G}_2(p, 2)) \leq 6 + 4\sqrt{2} \approx 11.657$. Note that there are several explicit examples reaching to the close value of this bound. For example, we have from Appendix A of [FS22] that $\mathcal{G}_2(137, 2) = 11.648$.

g	ℓ	$N_g(\ell)$	[ATY24]	Ours (Oh)	Ours (Arthur)
2	2	15	0.999593	0.888889	0.800000
	3	40	0.999285	0.750000	0.600000
	5	156	0.998975	0.555556	0.384615
	$\ell \rightarrow \infty$	∞	0.998396	0	0
3	2	135	0.999674	0.838052	0.628539
	3	1120	0.999428	0.649519	0.371154
	5	19656	0.999180	0.414087	0.177466
	$\ell \rightarrow \infty$	∞	0.998717	0	0
4	2	2295	0.999729	0.790123	0.470588
	3	91840	0.999523	0.562500	0.219512
	5	12304656	0.999317	0.308642	0.0798722
	$\ell \rightarrow \infty$	∞	0.998931	0	0
8	2	163204759575	0.999837	0.624295	0.124514
	3	234870301468364800	0.999714	0.316406	0.024687
	5	18342689145211210250854656	0.999590	0.095259	0.003199
	$\ell \rightarrow \infty$	∞	0.999358	0	0

Table 1. The last three columns are bounds on the absolute value of the non-trivial normalised eigenvalues, truncated to six digits.

6 Applications

We now present several applications of our rapid mixing result, focusing in particular on those concerning the underlying hardness assumptions of isogeny-based cryptography in dimension $g \geq 2$: the isogeny problem (Section 6.3), and the endomorphism ring problem (Section 6.4).

Computational framework to work with PPAVs. To present our algorithms, we first need to introduce a framework that allows us to work with PPAVs computationally following [CPW26]. To this end, we define the notion of an efficient model for PPAVs, which is a collection of algorithms that allow us to perform certain operations on PPAVs efficiently. We do not give the full definition here, but rather an informal one that captures the main idea. We highlight in particular the subroutines that are relevant for the algorithm we present later.

Notation. We use the same notational convention as [CPW26] for algorithms. When we define algorithms that have several subroutines, we say that $\mathcal{A}$ is an algorithm, and write $\mathcal{A}.\mathbf{xx}$ for its subroutine $\mathbf{xx}$. Here, the string $\mathbf{xx}$ can be interpreted as a first input to $\mathcal{A}$ which dictates what the algorithm is going to do, i.e., $\mathcal{A}.\mathbf{xx}(-) := \mathcal{A}(\mathbf{xx}, -)$. When we say that $\mathcal{A}$ is a probabilistic polynomial time algorithm, we mean that all of its subroutines $\mathcal{A}.\mathbf{xx}$ are probabilistic polynomial time algorithms (as a function of the length of the input).

Definition 6.1. (Informal). Consider a (probabilistic) algorithm $\mathcal{A}$ defined by a collection of algorithms that enable operations on the principally polarised abelian varieties, their points, and morphisms, which are encoded by binary strings.

For instance, an algorithm $\mathcal{A}.\text{isomorphic}$ would take as input the encodings of two PPAVs X, X' , and decide whether or not they are isomorphic. Another algorithm $\mathcal{A}.\text{diff}$ would take as input the encodings of two points $P, Q \in X$, and return the encoding of $P - Q$, computed using the group law on X . With appropriate bounds on the length of encodings and on the complexity of the family of algorithms $\mathcal{A}$, we will say that $\mathcal{A}$ is an efficient model for PPAVs. In particular, the algorithm $\mathcal{A}$ (i.e., each of its subroutines $\mathcal{A}.\mathbf{x}$) has expected complexity at most $\text{length}(\text{input})^{\text{poly}(\dim(\text{input}))}$.

Most relevant to us are the following subroutines that an efficient model $\mathcal{A}$ should have:

- $\mathcal{A}.\text{decompose}(X)$: given a PPAV X , it returns the encoding of the product $X_1 \times \cdots \times X_n$, where each X_i is indecomposable and $X \cong X_1 \times \cdots \times X_n$.
- $\mathcal{A}.\text{product}(X, Y)$: given PPAVs X, Y , it returns the encoding of the product $X \times Y$ with the product polarisation.
- $\mathcal{A}.\text{domcodom}(\varphi)$: given a morphism $\varphi : X \rightarrow Y$, it returns the encoding of the domain X and the codomain Y .
- $\mathcal{A}.\text{composition}(\varphi, \psi)$: given two morphisms $\varphi : X \rightarrow Y$ and $\psi : Y \rightarrow Z$, it returns the encoding of the composition $\psi \circ \varphi : X \rightarrow Z$.
- $\mathcal{A}.\text{torsion}(X, 1^N)$: given a PPAV X and a positive integer N , it returns the encoding of a basis of the N -torsion $X[N]$.
- $\mathcal{A}.\text{weil}(X, N, P, Q)$: given a PPAV X and two points $P, Q \in X[N]$, it returns the value of the Weil pairing $e_N(P, Q)$.
- $\mathcal{A}.\text{quotient}(X, P_1, \dots, P_{N^g})$: given a PPAV X and a maximal isotropic subgroup of $X[N]$ generated by $P_1, \dots, P_{N^g}$, it returns the encoding of the quotient PPAV $X/\langle P_1, \dots, P_{N^g} \rangle$ and of the quotient isogeny $X \rightarrow X/\langle P_1, \dots, P_{N^g} \rangle$.

The authors provide a concrete instantiation of this model: the (level-4) *theta model*, introduced by Mumford [Mum66]. Concretely, given a PPAV A defined over $\overline{\mathbb{F}}_p$, a (level-4) *theta structure* (or system of theta coordinates) is an injective map $\theta^A : A(\overline{\mathbb{F}}_p) \hookrightarrow \mathbb{P}^{4^g-1}(\overline{\mathbb{F}}_p)$, $P \mapsto \theta^A(P) := (\theta_0^A(P) : \cdots : \theta_{4^g-1}^A(P))$, where θ_i^A are theta functions (i.e., given by a canonical basis of the space of global sections of the line bundle $4\mathcal{L}$ where $\mathcal{L}$ is the ample line bundle associated to the polarisation on A). This allows us to view A as a subvariety of projective space $\mathbb{P}^{4^g-1}(\overline{\mathbb{F}}_p)$, and work with points P on A as projective points.

Theorem 6.2 ([CPW26, Corollary 7.7]). *The theta model is an efficient model for PPAVs.*

In the theta model, all the subroutines noted above can be implemented in time $\text{length}(\text{input})^{O(g)}$, except for $\mathcal{A}.\text{decompose}$ that terminates in $\text{length}(\text{input})^{O(g^2)}$.

6.1 Random walk

Every algorithm in this section is built around a central functionality: taking a random walk in $\mathcal{G}_\ell(p, g)$ for some fixed prime $\ell \neq p$. The walking procedure is itself folklore, and pseudocode is provided in Algorithm 1. Using Theorem 1.1, we get precise bounds for the length of a walk needed to achieve a distance ϵ to the stationary distribution.

Algorithm 1 Walk ^{$\mathcal{A}$}

Require: A superspecial PPAV X defined over $\mathbb{F}_p$ of dimension g , a prime $\ell \neq p$, an integer n .

Ensure: A random ℓ -isogeny path $X \rightarrow Y$ of length n .

```

1:  $\varphi \leftarrow$  an encoding of the identity morphism on  $X$ 
2:  $Y \leftarrow X$ 
3: for  $i$  from 1 to  $n$  do
4:    $Y[\ell] \leftarrow \mathcal{A}.\text{torsion}(Y, 1^\ell)$ 
5:    $\{P_1, \dots, P_{\ell^g}\} \leftarrow$  a uniformly random maximal isotropic subgroup of  $Y[\ell]$   $\triangleright$ 
   Using  $\mathcal{A}.\text{weil}$ 
6:    $\varphi' \leftarrow \mathcal{A}.\text{quotient}(Y, P_1, \dots, P_{\ell^g})$ 
7:    $\varphi \leftarrow \mathcal{A}.\text{composition}(\varphi, \varphi')$ 
8:    $\_, Y \leftarrow \mathcal{A}.\text{domcodom}(\varphi)$ 
9: end for
10: return  $\varphi$ 
```

Lemma 6.3. *Let $\epsilon > 0$. There is an algorithm that, given a superspecial PPAV X of dimension g over $\mathbb{F}_p$, computes a random walk of length $n = g^{O(1)} \log_\ell(p) - O(\log_\ell(\epsilon)/g)$ (the bound of Theorem 1.1) and returns an ℓ^n -isogeny $X \rightarrow Y$ where Y is a random PPAV at statistical distance at most ϵ from the stationary distribution. For any efficient model for PPAVs $\mathcal{A}$, it runs in time $n \cdot (\ell + \log p)^{\text{poly}(g)}$. If X is given in the theta model, it runs in time $n \cdot (\ell + \log p)^{O(g)}$.*

Proof. We claim that running Algorithm 1 with n gives us the required algorithm. Here we note that $\log(\#\mathcal{F}(X)) = 0$ since there is no level structure in our graph, and so the claim immediately follows from Theorem 1.1. By definition, each subroutine of $\mathcal{A}$ has expected complexity at most $(\ell + \log p)^{\text{poly}(g)}$. Each of these subroutines is called at most n times, immediately giving us the first statement. The second statement follows by observing that in the theta model, each of the subroutines run in Algorithm 1 has expected complexity at most $(\ell + \log p)^{O(g)}$. $\square$

6.2 An algorithm for finding a splitting

The goal of this section is to obtain an algorithm for finding a splitting of a superspecial PPAV into a product of supersingular elliptic curves. This will underlie both of our cryptanalytic applications.

The first step to construct such an algorithm is developing a routine that finds a splitting of a particular form. Namely, on input a superspecial PPAV X of dimension g , it should return a polarised isogeny $X \rightarrow E \times Y$ with E a supersingular elliptic curve and Y a superspecial PPAV of dimension $g - 1$. The strategy, going back to Costello and Smith [CS20], is simply to walk until we hit such a product: we run the random walk of Lemma 6.3 for n steps, and test whether the endpoint has a one-dimensional factor. If it does not, we restart. We present this algorithm (**FindSplitting**) in Algorithm 2.

Algorithm 2 FindSplitting ^{$\mathcal{A}$}

Require: A superspecial PPAV X defined over $\overline{\mathbb{F}}_p$ of dimension g , and a prime $\ell \neq p$.

Ensure: A PPAV Y_1 of dimension 1, a superspecial PPAV Y_2 of dimension $g - 1$, a positive integer m , and an ℓ^{2m} -isogeny $\varphi : X \rightarrow Y_1 \times Y_2$.

```

1: while true do
2:    $e \leftarrow$  the smallest even integer  $\geq n$ , with  $n$  as in Equation (25)
3:    $(\varphi : X \rightarrow Y) \leftarrow \text{Walk}^{\mathcal{A}}(X, \ell, e)$   $\triangleright$  A random walk computed by Algorithm 1
4:    $Y_1, \dots, Y_k \leftarrow \mathcal{A}.\text{decompose}(Y)$ 
5:   Up to reordering, ensure that  $\dim(Y_1) \leq \dim(Y_i)$  for all  $i = 2, \dots, k$ .
6:   if  $\mathcal{A}.\text{dimension}(Y_1) = 1$  then
7:     for  $i$  from 3 to  $k$  do
8:        $Y_2 \leftarrow \mathcal{A}.\text{product}(Y_2, Y_i)$   $\triangleright$  Construct  $Y_2 \times \dots \times Y_k$ 
9:     end for
10:     $m \leftarrow e/2$ 
11:    return  $Y_1, Y_2, m, \varphi$ 
12:   end if
13: end while

```

With the results from Lemma 6.3, we can analyse the runtime of the algorithm by computing the relative mass of the set $S \subseteq \text{SS}(p, g)$ of superspecial PPAVs with a one-dimensional factor: the probability of success of one iteration is $\mu(S)/\mu(V)$, with μ the mass function. We know $\mu(V)$ from Lemma 4.4, and using similar techniques, we compute $\mu(S)$.

Theorem 6.4. *For any prime p and positive integer g ,*

$$\mu(\text{SS}(p, g)) = \frac{1}{(2\pi)^{g(g+1)}} \prod_{i=1}^g (p^i + (-1)^i) \zeta(2i) (2i - 1)! = \prod_{i=1}^g \frac{|B_{2i}|}{4i} (p^i + (-1)^i),$$

where ζ is the Riemann zeta function and B_{2i} is the $2i$ -th Bernoulli number.

Proof. This is the classical mass formula of $\text{SS}(p, g)$: see Hashimoto–Ibukiyama [HI80, Prop. 9] and Equation (5.3) of Ekedahl [Eke87]. $\square$

Lemma 6.5. *Let $\text{SS}(p, g)$ be the set of $\overline{\mathbb{F}}_p$ -isomorphism classes of superspecial PPAVs of dimension g . Then we have that*

1. the mass $\mu(\text{SS}(p, g)) \leq e^{\text{poly}(g)} p^{g(g+1)/2}$,
2. defining a function $\alpha_g : \text{SS}(p, 1) \times \text{SS}(p, g-1) \rightarrow \text{SS}(p, g)$ by $(E, X) \mapsto E \times X$, we have

$$\frac{\mu(\text{im}(\alpha_g))}{\mu(\text{SS}(p, g))} = \Omega(p^{1-g}/|B_{2g}|),$$

where B_{2i} is the $2i$ -th Bernoulli number.

Proof. Using Stirling's approximation, we see that $\prod_{i=1}^g (2i-1)! \leq e^{2g^2 \log_e(2g)}$. The first point of the lemma then follows from Theorem 6.4.

Note that the function α_g is not surjective. We want to estimate the mass $\mu_{p,g}^{(1)}$ of the image of α_g . By following [Eke87, §7], we can see that for any given $g > 1$ it is a polynomial in p of degree $g(g-1)/2 + 1 = g(g+1)/2 - (g-1)$. Indeed, by inclusion-exclusion using [Eke87, Proposition 7.1] we get

$$\mu_{p,g}^{(1)} = \sum_{k=1}^g \frac{\mu(\text{SS}(p, 1))^k (-1)^{k-1}}{k!} \mu(\text{SS}(p, g-k)). \quad (24)$$

Since $\mu(\text{SS}(p, g))$ is a polynomial of degree $g(g+1)/2$, the k -th term in the sum is a polynomial in p of degree $f(k) = k + \frac{(g-k)(g-k+1)}{2}$. For $1 \leq k \leq g$, the maximum of $f(k)$ is achieved only at $k = 1$, so the mass $\mu_{p,g}^{(1)}$ has degree $1 + g(g-1)/2$. For $2 \leq g \leq 5$ this is enough to conclude (because the ratio has degree $g(g+1)/2 - (1 + g(g-1)/2) = 1 - g$).

For $g > 5$ we assume (as we may) that $p > 2$ and we use [Eke87, Proposition 7.5 (ii)]: denoting $M(p, g) = \mu(\text{SS}(p, g))$ and $M'(p, g)$ the mass of indecomposable PPAVs, we have

$$\mu(\text{im}(\alpha_g)) \geq M(p, 1)M'(p, g-1) \geq M(p, 1) \frac{M(p, g-1)}{g-1}.$$

Dividing both sides by the total mass $M(p, g)$ and observing that $\frac{M(p, g-1)}{M(p, g)} = \frac{4g}{|B_{2g}|(p^g + (-1)^g)}$, we obtain the desired ratio bound:

$$\frac{\mu(\text{im}(\alpha_g))}{M(p, g)} \geq \frac{M(p, 1)4g}{(g-1)|B_{2g}|(p^g + (-1)^g)} = \Omega(p^{1-g}/|B_{2g}|).$$

□

Remark 6.6. The masses stated in Lemma 6.5 are explicitly computable: $\mu(V) = \mu(\text{SS}(p, g))$ is given in Theorem 6.4, and $\mu(S) = \mu_{p,g}^{(1)}$ is given in Equation (24). By using the conditional bound in Theorem 4.1 for $\lambda(\mathcal{G})/d$, the walk length required in the following Algorithm 2 is therefore the explicit quantity

$$n = \frac{\log \left(2 \# \text{Aut}(X)^{1/2} \mu(V) \mu(S)^{-1/2} \right)}{\log \left(\ell^{g/2} + \ell^{-g/2} \right) - \log 2}. \quad (25)$$

We can now analyse the runtime of the algorithm `FindSplitting`.

Proposition 6.7. *There is an algorithm (Algorithm 2) which, on input a superspecial PPAV X of dimension g over $\mathbb{F}_p$ and a prime $\ell \neq p$, outputs a polarised isogeny $\varphi : X \rightarrow Y_1 \times Y_2$ where Y_1 and Y_2 is a superspecial PPAV of dimension 1 and $g - 1$, respectively, in expected time $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$. In the theta model, the complexity becomes $(\ell + \log p)^{O(g^2)} \cdot p^{g-1}$.*

Proof. Let $S = \text{im}(\alpha_g)$ denote the target set of vertices, which corresponds to the superspecial PPAVs that decompose a 1-dimensional factor (i.e., those of the form $Y_1 \times Y_2$). By Lemma 6.5, the mass ratio is $\frac{\mu(S)}{\mu(V)} = \Omega(p^{1-g})$.

We want to choose a walk length n large enough so that the probability of landing S is at least half of its mass ratio, i.e., $\sum_{x \in S} P_n(x) \geq \frac{1}{2} \frac{\mu(S)}{\mu(V)}$. By Lemma 2.10, we have that:

$$\left| \sum_{x \in S} P_n(x) - \frac{\mu(S)}{\mu(V)} \right| \leq \mu(X)^{-1/2} \mu(S)^{1/2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n.$$

Thus, it is sufficient to have the following inequality for choosing n :

$$\mu(X)^{-1/2} \mu(S)^{1/2} \left(\frac{\lambda(\mathcal{G})}{d} \right)^n \leq \frac{1}{2} \frac{\mu(S)}{\mu(V)},$$

i.e., we find the explicit minimum walk length required:

$$n \geq \frac{\log(2\mu(X)^{-1/2} \mu(S)^{-1/2} \mu(V))}{-\log\left(\frac{\lambda(\mathcal{G})}{d}\right)}.$$

Note that the total mass $\mu(V)$ and $\mu(S)$ are bounded by a polynomial in p by Lemma 6.5, and the denominator is a positive constant bounded away from zero (depending only on ℓ and g) by Theorem 4.1, we satisfy this inequality by choosing $n = O(\log p)$. Hence, the probability of that Y is of the form $Y_1 \times Y_2$, where $\dim(Y_1) = 1$ and $\dim(Y_2) = g - 1$, is $\sum_{x \in S} P_n(x) \geq \frac{1}{2} \Omega(p^{1-g}) = \Omega(p^{1-g})$. The expected number of iterations of the **while** loop is $(\sum_{x \in S} P_n(x))^{-1} = O(p^{g-1})$.

By Lemma 6.3, the cost of an iteration is $(\ell + \log p)^{\text{poly}(g)}$, and thus the total expected running time is $O(p^{g-1}) \cdot (\ell + \log p)^{\text{poly}(g)} = (\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$.

In the theta model, generating the random walk is achieved in time $n \cdot (\ell + \log p)^{O(g)}$ (cf. Lemma 6.3) while the (dominated) subroutine $\mathcal{A}.\text{decompose}$ operates in time $(\ell + \log p)^{O(g^2)}$. Thus, we conclude that the final expected complexity in the theta model of $(\ell + \log p)^{O(g^2)} \cdot p^{g-1}$ as stated. $\square$

Remark 6.8. The subroutine $\mathcal{A}.\text{decompose}$ works by first guessing the decomposition $X_1 \times \cdots \times X_g$ with $\sum \dim(X_i) = g$, and then determining if X decomposes in this way. For use in Algorithm 2, it could therefore be easily adapted to just check for the specific decomposition $Y_1 \times Y_2$ where $\dim(Y_1) = 1$ and $\dim(Y_2) = g - 1$. This adaptation would in fact be more efficient since we would not have to check for all possible decompositions of g .

Applying the algorithm `FindSplitting` to the $(g - 1)$ -dimensional factor and iterating g times, we descend all the way to a product of elliptic curves. The isogenies obtained at the successive stages can be assembled into a single polarised isogeny $X \rightarrow E_1 \times \cdots \times E_g$. We obtain the algorithm `ProductCurves`, given in Algorithm 3.

Algorithm 3 `ProductCurves` ^{$\mathcal{A}$}

Require: A PPAV X defined over $\mathbb{F}_q$ of dimension g , and a prime $\ell \neq p$.

Ensure: Supersingular elliptic curves $E_1, \dots, E_g$, and an ℓ^n -isogeny $\varphi : X \rightarrow E_1 \times \cdots \times E_g$.

```

1:  $E_1, Y, m, \varphi \leftarrow \text{FindSplitting}^{\mathcal{A}}(X, \ell)$ 
2: for  $k$  from 2 to  $g - 1$  do
3:    $E_k, Y, m_k, \varphi_k \leftarrow \text{FindSplitting}^{\mathcal{A}}(Y, \ell)$ 
4:    $[\ell^{m_k}] \leftarrow$  multiplication-by- $\ell^{m_k}$  map on  $E_1 \times \cdots \times E_{k-1}$ 
5:    $\pi \leftarrow$  product isogeny  $[\ell^{m_k}] \times \varphi_k$ 
6:    $\varphi \leftarrow \mathcal{A}.\text{composition}(\varphi, \pi)$ 
7: end for
8:  $E_g \leftarrow Y$ 
9: return  $E_1, \dots, E_g, \varphi$ 

```

Theorem 6.9. *There is an algorithm (Algorithm 3) which, on input a superspecial PPAV X of dimension g over $\mathbb{F}_p$ and a prime $\ell \neq p$, outputs a polarised isogeny $\varphi : X \rightarrow E_1 \times \cdots \times E_g$, where $E_1, \dots, E_g$ are supersingular elliptic curves, in expected time $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$. In the theta model, the complexity is $(\ell + \log p)^{O(g^2)} \cdot p^{g-1}$.*

Proof. We claim Algorithm 3 implements the desired algorithm. The result follows by observing that all subroutines of $\mathcal{A}$ run in $(\ell + \log p)^{\text{poly}(g)}$, `FindSplitting` ^{$\mathcal{A}$} runs in $(\ell + \log p)^{\text{poly}(g)} \cdot p^{g-1}$ by Proposition 6.7, and each of these routines runs g times. The strengthening in the theta model follows by observing that, in this case, all subroutines of $\mathcal{A}$ run in $(\ell + \log p)^{O(g)}$ and `FindSplitting` ^{$\mathcal{A}$} runs in $(\ell + \log p)^{O(g^2)} \cdot p^{g-1}$.

6.3 Solving the superspecial isogeny problem

The general *dimension g superspecial isogeny problem* asks us to find a polarised isogeny $\varphi : X \rightarrow X'$ between superspecial PPAVs X and X' of dimension g over $\mathbb{F}_p$. For $g \geq 2$, the best known (heuristic) algorithm for solving the superspecial isogeny problem is due to Costello and Smith [CS20]. At the high level, the algorithm runs in the following steps: (1) find isogenies $\phi : X \rightarrow E_1 \times \cdots \times E_g$ and $\phi' : X' \rightarrow E'_1 \times \cdots \times E'_g$, where $E_1, \dots, E_g, E'_1, \dots, E'_g$ are supersingular elliptic curves; and (2) for $i = 1, \dots, g$, find isogenies $\psi_i : E_i \rightarrow E'_i$ and compose them to form a product isogeny $\pi : E_1 \times \cdots \times E_g \rightarrow E'_1 \times \cdots \times E'_g$. Arguing that Step (2) has a failure probability of $1/2^g$, Costello and Smith conclude that the

algorithm runs in $\tilde{O}(p^{g-1})$ classical bit operations on average [CS20, Theorem 1] under the (false) hypothesis that the graph $\mathcal{G}_\ell(p, g)$ is Ramanujan. With our new understanding of the spectral properties of $\mathcal{G}_\ell(p, g)$, in Theorem 6.9 we introduced an algorithm that implements Step (1) and provably terminates. We use this to rigorously prove the termination of the Costello–Smith algorithm.

Remark 6.10. In contrast to the analysis in [CS20], our complexity analysis takes into account how the PPAVs are represented, and the dependency on g .

Theorem 6.11. *There exists an algorithm that, on input two dimension g superspecial principally polarised abelian varieties X and X' defined over $\mathbb{F}_p$, outputs a polarised isogeny $\varphi : X \rightarrow X'$ in expected time $(\log p)^{\text{poly}(g)} \cdot g \cdot p^{g-1}$. In the theta model, the complexity becomes $(\log p)^{O(g^2)} \cdot g \cdot p^{g-1}$.*

Proof. The algorithm runs as follows. First, run the algorithm from Theorem 6.9 on inputs X and $\ell = 2$ to obtain a polarised isogeny $\phi : X \rightarrow E_1 \times \cdots \times E_g$, where $E_1, \dots, E_g$ are supersingular elliptic curves. Repeat for X' to obtain $\phi' : X' \rightarrow E'_1 \times \cdots \times E'_g$. This terminates in expected time $(\log p)^{\text{poly}(g)} \cdot p^{g-1}$.

Then, run [PW23, Algorithm 7] to find isogenies $\psi_i : E_i \rightarrow E'_i$ of degree 2^{2a_i} for some $a_i \in \mathbb{Z}$. By [PW24b, Proposition 8.7], the algorithm terminates in expected time $\text{poly}(\log p) p^{1/2}$. As the 2-valuation of $\deg(\psi_i)$ is even for each i , we can construct a product isogeny $\pi : E_1 \times \cdots \times E_g \rightarrow E'_1 \times \cdots \times E'_g$ following [CS20, Lemma 3]. The algorithm outputs $(\phi')^\dagger \circ \pi \circ \phi$, a polarised isogeny $X \rightarrow X'$. $\square$

6.4 Computing endomorphism rings of superspecial PPAVs

In [ABI⁺25, Corollary 5.20], the authors show that assuming rapid mixing properties of the superspecial isogeny graph for $g = 2$, there is an algorithm that on input a superspecial Jacobian $X/\mathbb{F}_{p^2}$ outputs a basis of $\text{End}(X)$ in expected time $\tilde{O}(p)$. As in the previous section, we can use Theorem 4.1 to rigorously prove the runtime of this algorithm. We furthermore generalise this result to compute the endomorphism ring of any superspecial PPAV $X/\mathbb{F}_{p^2}$ of dimension $g \geq 2$.

The strategy proceeds in two steps. First, we compute the endomorphism ring of a product of supersingular elliptic curves, where everything can be done in the language of quaternion algebras. Second, we transport this ring back to X along an isogeny produced by Theorem 6.9. The second step is where our rapid mixing result enters, and it is also what fixes the overall complexity.

For the first step, note that every endomorphism of $E_1 \times \cdots \times E_g$ is given by a $g \times g$ matrix whose (i, j) entry lies in $\text{Hom}(E_i, E_j)$, so it suffices to compute a basis of each $\text{Hom}(E_i, E_j)$. This is done in Algorithm 4, and works as follows. The diagonal entries are the endomorphism rings $\text{End}(E_i)$, each of which can be computed in expected time $\text{poly}(\log p) \cdot \sqrt{p}$ by [PW24a, Theorem 8.8]. The off-diagonal entries then come for free: given bases of $\text{End}(E_i)$ and $\text{End}(E_j)$, we can embed both into a common quaternion algebra and compute the ideal connecting the two resulting maximal orders. This yields a basis of $\text{Hom}(E_i, E_j)$ in time $\text{poly}(\log p)$, as shown in Lemma 6.12.

Algorithm 4 Computing $\text{Hom}(E_1, E_2)$ from the endomorphism rings

Require: Efficient representations of two supersingular elliptic curves E_1 and E_2 , and of a basis of the rings $\text{End}(E_1)$ and $\text{End}(E_2)$.

Ensure: An efficient representation of a basis of $\text{Hom}(E_1, E_2)$.

- 1: Find integers a, b and an embedding $\iota_1 : \text{End}(E_1) \rightarrow \left(\frac{-a, -b}{\mathbb{Q}}\right)$. ▷ [HW25, Proposition 13]
 - 2: Find integers a', b' and an embedding $\iota_2 : \text{End}(E_2) \rightarrow \left(\frac{-a', -b'}{\mathbb{Q}}\right)$. ▷ [HW25, Proposition 13]
 - 3: Find an isomorphism $j : \left(\frac{-a', -b'}{\mathbb{Q}}\right) \rightarrow \left(\frac{-a, -b}{\mathbb{Q}}\right)$. ▷ [CKMZ22, Proposition 4.1]
 - 4: Write $\mathcal{O}_i$ the image of $\text{End}(E_i)$ into $\left(\frac{-a, -b}{\mathbb{Q}}\right)$ through the above maps, for $i = 1, 2$.
 - 5: Let $I = I(\mathcal{O}_1, \mathcal{O}_2)$ be the connecting ideal, with four generators $\alpha_1, \dots, \alpha_4 \in I$. ▷ [KV10, Algorithm 3.5]
 - 6: **for** $i = 1, 2, 3, 4$ **do**
 - 7: Let $I_i = I\bar{\alpha}_i / \text{Nrd}(I)$.
 - 8: Let $\varphi_i = \varphi_{\iota_1^{-1}(I_i)} : E_1 \rightarrow E_2$. ▷ Using the method of [PR23] for ideal-to-isogeny conversion
 - 9: **end for**
 - 10: Let B be a basis of the lattice generated by $\varphi_i \circ \gamma$ for all i and all $\gamma \in \text{Aut}(E_1)$. ▷ [BP87]
 - 11: **return** B .
-

Lemma 6.12. *There is an algorithm (Algorithm 4) which, on input two supersingular elliptic curves E_1 and E_2 over $\mathbb{F}_{p^2}$, together with bases of their endomorphism rings, outputs a basis of $\text{Hom}(E_1, E_2)$ in expected time $\text{poly}(\log p)$.*

Proof. Let us first recall some standard notation for the Deuring correspondence. Given any left-ideal J in $\text{End}(E_1)$, we write $\varphi_J : E_1 \rightarrow E$ the unique (up to post-composition by an isomorphism) isogeny from E_1 of degree $\text{Nrd}(J)$ such that $\bigcap_{\alpha \in J} \ker \alpha = \ker \varphi_J$. By the Deuring correspondence, the map $J \mapsto \varphi_J$ is a bijection from left-ideals in $\text{End}(E_1)$ to isomorphism classes of isogenies of domain E_1 . Its inverse is the map $(\varphi : E_1 \rightarrow E) \mapsto I_\varphi := \text{Hom}(E, E_1) \circ \varphi$.

Using the notation from Algorithm 4, let $J = \iota_1^{-1}(I)$ (a left-ideal of $\text{End}(E_1)$), and $\beta_i = \iota_1^{-1}(\alpha_i)$, which form a basis of J . We have $J = \text{Hom}(E_2, E_1) \circ \varphi_J$, so there exist isogenies $\psi_i : E_2 \rightarrow E_1$ such that $\beta_i = \psi_i \circ \varphi_J$. Since $(\beta_i)_i$ is a basis of $\text{Hom}(E_2, E_1) \circ \varphi_J$, we deduce that $(\hat{\psi}_i)_i$ is a basis of $\text{Hom}(E_1, E_2)$. Now,

$$\begin{aligned} \iota_1^{-1}(I_i) &= \iota_1^{-1}(I\bar{\alpha}_i / \text{Nrd}(I)) = J\hat{\beta}_i / \deg(\varphi_J) \\ &= \text{Hom}(E_2, E_1) \circ \varphi_J \circ \hat{\varphi}_J \circ \hat{\psi}_i / \deg(\varphi_J) \\ &= \text{Hom}(E_2, E_1) \circ \hat{\psi}_i. \end{aligned}$$

Hence, we have $\varphi_i = \hat{\psi}_i$ (up to post-composition by an isomorphism). This proves that the elements $\varphi_i \circ \gamma$ with $\gamma \in \text{Aut}(E_1)$ generate all of $\text{Hom}(E_1, E_2)$, hence Algorithm 4 is correct. The complexity follows from the results cited at each step of Algorithm 4. □

Remark 6.13. The output of Algorithm 4 is represented as $g \times g$ matrices of isogenies between elliptic curves. It can then be converted into any efficient model for PPAVs by a careful composition with the embedding and projection maps (in time $(\log p)^{\text{poly}(g)}$).

We can now compute the endomorphism ring of a product of supersingular elliptic curves, as detailed in Algorithm 5.

Algorithm 5 Computing $\text{End}(E_1 \times \cdots \times E_g)$

Require: Efficient representations of g supersingular elliptic curves $E_1, \dots, E_g$.

Ensure: An efficient representation of a basis of $\text{End}(E_1 \times \cdots \times E_g)$.

```

1: for  $i = 1, \dots, g$  do
2:   Compute a basis of  $\text{End}(E_i)$ . ▷ Using [PW24a, Theorem 8.8]
3: end for
4: for  $i = 1, \dots, g$  do
5:   for  $j = 1, \dots, g$  do
6:     Compute a basis of  $H_{i,j} = \text{Hom}(E_i, E_j)$ , using  $\text{End}(E_i)$  and  $\text{End}(E_j)$ . ▷
       Lemma 6.12
7:   end for
8: end for

```

9: **return** the basis of $\text{End}(E_1 \times \cdots \times E_g) = \begin{pmatrix} H_{1,1} & \cdots & H_{1,g} \\ \vdots & \ddots & \vdots \\ H_{g,1} & \cdots & H_{g,g} \end{pmatrix}$.

Lemma 6.14. *There is an algorithm (Algorithm 5) which, on input g supersingular elliptic curves $E_1, \dots, E_g$ over $\mathbb{F}_{p^2}$, outputs a basis of $\text{End}(E_1 \times \cdots \times E_g)$ in expected time $\text{poly}(\log q) \cdot (g^2 + g\sqrt{p})$.*

Proof. This follows from the theorems and algorithms referred to in Algorithm 5, and from the fact that any endomorphism of $E_1 \times \cdots \times E_g$ can be written in matrix form. The first loop has complexity $\text{poly}(\log q) \cdot g \cdot \sqrt{p}$, and the second has complexity $\text{poly}(\log q) \cdot g^2$. $\square$

For the second step, we proceed as follows. Let $\varphi : X \rightarrow E_1 \times \cdots \times E_g$ be a 2^a -isogeny obtained from Theorem 6.9 with $\ell = 2$, and write $\mathcal{O} = \text{End}(E_1 \times \cdots \times E_g)$. Conjugating by φ gives a subring $\varphi^\dagger \mathcal{O} \varphi \subseteq \text{End}(X)$, and since $\varphi^\dagger \varphi = [2^a]$ we also have $2^{2a} \text{End}(X) \subseteq \varphi^\dagger \mathcal{O} \varphi$. The index $[\text{End}(X) : \varphi^\dagger \mathcal{O} \varphi]$ is therefore a power of 2: the computed subring may fall short of $\text{End}(X)$, but only 2-adically. Running Theorem 6.9 a second time with $\ell = 3$ produces a second subring of index a power of 3. As the two indices are coprime, the ring generated by both is all of $\text{End}(X)$. By employing Theorem 6.9, we require no heuristic to prove the runtime, yielding our final result.

Theorem 6.15. *There is an algorithm (Algorithm 6) which, on input a superspecial PPAV X (in any efficient model) of dimension $g > 1$, outputs a basis of*

Algorithm 6 Computing $\text{End}(X)$

Require: An efficient representations of a superspecial PPAV X .

Ensure: An efficient representation of a basis of $\text{End}(X)$.

- 1: $(E_1, \dots, E_g, \varphi) \leftarrow \text{ProductCurves}^{\mathcal{A}}(X, 2)$ a 2-isogeny path $\varphi : X \rightarrow E_1 \times \dots \times E_g$.
▷ Theorem 6.9
 - 2: $(E'_1, \dots, E'_g, \varphi') \leftarrow \text{ProductCurves}^{\mathcal{A}}(X, 3)$ a 3-isogeny path $\varphi' : X \rightarrow E'_1 \times \dots \times E'_g$.
▷ Theorem 6.9
 - 3: Compute a basis $(\alpha_i)_i$ of $\text{End}(E_1 \times \dots \times E_g)$ and $(\alpha'_i)_i$ of $\text{End}(E'_1 \times \dots \times E'_g)$. ▷ Lemma 6.14
 - 4: $\beta_i \leftarrow \varphi^\dagger \circ \alpha_i \circ \varphi$ for all i . ▷ $(\beta_i)_i$ is a basis of $\varphi^\dagger \circ \text{End}(E_1 \times \dots \times E_g) \circ \varphi \subseteq \text{End}(X)$
 - 5: $\beta'_i \leftarrow \varphi'^\dagger \circ \alpha'_i \circ \varphi'$ for all i . ▷ $(\beta'_i)_i$ is a basis of $\varphi'^\dagger \circ \text{End}(E'_1 \times \dots \times E'_g) \circ \varphi' \subseteq \text{End}(X)$
 - 6: Compute a basis $(\gamma_i)_i$ of the ring generated by both $(\beta_i)_i$ and $(\beta'_i)_i$. ▷ [BP87]
 - 7: **return** $(\gamma_i)_i$.
-

$\text{End}(X)$ in expected time $(\log p)^{\text{poly}(g)} \cdot p^{g-1}$. In the theta model, the complexity becomes $(\log p)^{O(g^2)} \cdot p^{g-1}$.

Proof. The strategy we follow is similar to the case of dimension 2 described in [ABI⁺25, Corollary 5.20]. Write $\mathcal{O} = \text{End}(E_1 \times \dots \times E_g)$. The endomorphisms $(\beta_i)_i$ of X generate the subring $\varphi^\dagger \mathcal{O} \varphi \subseteq \text{End}(X)$. We also have $\varphi \text{End}(X) \varphi^\dagger \subseteq \mathcal{O}$, hence $2^{2a} \text{End}(X) \subseteq \varphi^\dagger \mathcal{O} \varphi$, since φ is a 2^a -isogeny. We deduce

$$2^{2a} \text{End}(X) \subseteq \varphi^\dagger \mathcal{O} \varphi \subseteq \text{End}(X),$$

hence $[\text{End}(X) : \varphi^\dagger \mathcal{O} \varphi]$ divides $[\text{End}(X) : 2^{2a} \text{End}(X)]$, a power of 2; write $[\text{End}(X) : \varphi^\dagger \mathcal{O} \varphi] = 2^A$. Similarly, the endomorphisms $(\beta'_i)_i$ of X generate the subring $\varphi'^\dagger \mathcal{O}' \varphi' \subseteq \text{End}(X)$ of index a power of 3, say 3^B . Therefore, the subring generated by all of $(\beta_i)_i$ and $(\beta'_i)_i$ has index dividing $\gcd(2^A, 3^B) = 1$, so they generate $\text{End}(X)$. This proves the correctness of the algorithm. Its running time follows from the results cited at each step of Algorithm 6. $\square$

The cost is dominated by the calls to Theorem 6.9; the computations of Lemma 6.14 contribute only $\text{poly}(\log p) \cdot (g^2 + g\sqrt{p})$, which is negligible for $g > 1$. For $g = 2$, we recover the complexity of [ABI⁺25, Corollary 5.20], now unconditionally.

References

- ABI⁺25. Samuele Anni, Gaetan Bisson, Annamaria Iezzi, Elisa Lorenzo García, and Benjamin Wesolowski, *On the computation of endomorphism rings of abelian surfaces over finite fields*, 2025.
- Arp24. Sarah Arpin, *Adding level structure to supersingular elliptic curve isogeny graphs*, Journal de théorie des nombres de Bordeaux **36** (2024), no. 2, 405–443.
- Art89. James Arthur, *Unipotent automorphic representations: conjectures*, Astérisque (1989), no. 171-172, 13–71, Orbites unipotentes et représentations, II.

- Art13. James Arthur, *The Endoscopic Classification of Representations: Orthogonal and Symplectic groups*, American Mathematical Society Colloquium Publications, vol. 61, American Mathematical Society, 2013.
- ATY24. Yusuke Aikawa, Ryokichi Tanaka, and Takuya Yamauchi, *Isogeny graphs on superspecial abelian varieties: eigenvalues and connection to Bruhat–Tits buildings*, Canadian Journal of Mathematics **76** (2024), no. 6, 1891–1916.
- 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*, Advances in Cryptology – EUROCRYPT 2023, Part II (Lyon, France) (Carmit Hazay and Martijn Stam, eds.), Lecture Notes in Computer Science, vol. 14005, Springer, Cham, Switzerland, April 23–27, 2023, pp. 405–437.
- BF23. Andrea Basso and Tako Boris Fouotsa, *New SIDH countermeasures for a more efficient key exchange*, Advances in Cryptology – ASIACRYPT 2023, Part VIII (Guangzhou, China) (Jian Guo and Ron Steinfeld, eds.), Lecture Notes in Computer Science, vol. 14445, Springer, Singapore, Singapore, December 4–8, 2023, pp. 208–233.
- BFLS26. Giacomo Borin, Luca De Feo, Guido Lido, and Sina Schaeffler, *The sqinstructor: a guide to sqsign and the deuring correspondence with level structures*, Advances in Cryptology – CRYPTO 2026 (Cham) (Nadia Heninger and Mike Rosulek, eds.), Springer Nature Switzerland, 2026, pp. 537–569.
- BFP⁺26. Andrea Basso, Luca De Feo, Sikhar Patranabis, Ilinca Radulescu, and Benjamin Wesolowski, *Forensic categories: a framework for SQIsign-like primitives*, Cryptology ePrint Archive, Paper 2026/1171, 2026.
- BP87. Johannes Buchmann and Michael Pohst, *Computing a lattice basis from a system of generating vectors*, European Conference on Computer Algebra, Springer, 1987, pp. 54–63.
- Car12. Ana Caraiani, *Local-global compatibility and the action of monodromy on nearby cycles*, Duke Math. J. **161** (2012), no. 12, 2311–2413.
- CD23. Wouter Castryck and Thomas Decru, *An efficient key recovery attack on SIDH*, Advances in Cryptology – EUROCRYPT 2023, Part V (Lyon, France) (Carmit Hazay and Martijn Stam, eds.), Lecture Notes in Computer Science, vol. 14008, Springer, Cham, Switzerland, April 23–27, 2023, pp. 423–447.
- CDS20. Wouter Castryck, Thomas Decru, and Benjamin Smith, *Hash functions from superspecial genus-2 curves using richelot isogenies*, Journal of Mathematical Cryptology **14** (2020), no. 1, 268–292.
- Che22. Gaëtan Chenevier, *Statistics for Kneser p -neighbors*, Bull. Soc. Math. Fr. **150** (2022), no. 3, 473–516 (English), with an appendix by Olivier Taïbi.
- CKMZ22. Tímea Csahók, Péter Kutas, Mickaël Montessinos, and Gergely Zábrádi, *Explicit isomorphisms of quaternion algebras over quadratic global fields*, Research in Number Theory **8** (2022), no. 4, 77.
- CL12. Pierre-Henri Chaudouard and Gérard Laumon, *Le lemme fondamental pondéré. II. Énoncés cohomologiques*, Ann. of Math. (2) **176** (2012), no. 3, 1647–1781.
- CL26. Giulio Codogni and Guido Maria Lido, *Spectral theory of isogeny graphs*, Journal of Number Theory **286** (2026), 131–184.
- CLG09. Denis Xavier Charles, Kristin E. Lauter, and Eyal Z. Goren, *Cryptographic hash functions from expander graphs*, Journal of Cryptology **22** (2009), no. 1, 93–113.

- Clo88. Laurent Clozel, *Motifs et Formes Automorphes: Applications du Principe de Fonctorialité*, Automorphic Forms, Shimura Varieties, and L -functions (Laurent Clozel and James S. Milne, eds.), vol. 1, 1988, Proceedings of a Conference held at the University of Michigan, Ann Arbor.
- Clo13. Laurent Clozel, *Purity reigns supreme*, Int. Math. Res. Not. IMRN (2013), 328–346.
- CPW26. Maria Corte-Real Santos, Etienne Piasecki, and Benjamin Wesolowski, *A computational framework for principally polarized abelian varieties and applications*, Cryptology ePrint Archive, Paper 2026/1142, 2026.
- CS20. Craig Costello and Benjamin Smith, *The supersingular isogeny problem in genus 2 and beyond*, Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020 (Paris, France) (Jintai Ding and Jean-Pierre Tillich, eds.), Springer, Cham, Switzerland, April 15–17, 2020, pp. 151–168.
- Dem72. Michel Demazure, *Lectures on p -divisible groups*, Lect. Notes Math., vol. 302, Springer, Cham, 1972 (English).
- DFP24. Luca De Feo, Tako Boris Fouotsa, and Lorenz Panny, *Isogeny problems with level structure*, Advances in Cryptology – EUROCRYPT 2024, Part VII (Zurich, Switzerland) (Marc Joye and Gregor Leander, eds.), Lecture Notes in Computer Science, vol. 14657, Springer, Cham, Switzerland, May 26–30, 2024, pp. 181–204.
- DG70. Michel Demazure and Alexander Grothendieck (eds.), *Schémas en groupes. III: Structure des schémas en groupes réductifs. Exposés XIX à XXVI. Séminaire de Géométrie Algébrique du Bois Marie 1962/64 (SGA 3), dirigé par Michel Demazure et Alexander Grothendieck. Revised reprint*, Lect. Notes Math., vol. 153, Springer, Cham, 1970 (French).
- DKL⁺20. Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski, *SQISign: Compact post-quantum signatures from quaternions and isogenies*, Advances in Cryptology – ASIACRYPT 2020, Part I (Daejeon, South Korea) (Shiho Moriai and Huaxiong Wang, eds.), Lecture Notes in Computer Science, vol. 12491, Springer, Cham, Switzerland, December 7–11, 2020, pp. 64–93.
- DR26. Thomas Decru and Krijn Reijnders, *Expander properties of superspecial isogeny digraphs with level structure*, Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part IV (Nadia Heninger and Mike Rosulek, eds.), Lecture Notes in Computer Science, vol. 16803, Springer, 2026, pp. 478–505.
- Eke87. Torsten Ekedahl, *On supersingular curves and abelian varieties*, Mathematica Scandinavica **60** (1987), 151–178.
- FJP14. Luca De Feo, David Jao, and Jérôme Plût, *Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies*, Journal of Mathematical Cryptology **8** (2014), no. 3, 209–247.
- Fon77. Jean-Marc Fontaine, *Groupes p -divisibles sur les corps locaux*, Astérisque, vol. 47-48, Société Mathématique de France (SMF), Paris, 1977 (French).
- Fon90. ———, *p -adic representations of local fields. I*, The Grothendieck Festschrift, Collect. Artic. in Honor of the 60th Birthday of A. Grothendieck. Vol. II, Prog. Math. 87, 249-309 (1990)., 1990.
- FS22. E. Florit and B. Smith, *Automorphisms and isogeny graphs of abelian varieties, with applications to the superspecial Richelot isogeny graph*, Arithmetic, geometry, cryptography, and coding theory 2021, Contemp. Math.,

- vol. 779, Amer. Math. Soc., [Providence], RI, 2022, pp. 103–132. MR 4445773
- FT19. E. Victor Flynn and Yan Bo Ti, *Genus two isogeny cryptography*, Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019 (Chongqing, China) (Jintai Ding and Rainer Steinwandt, eds.), Springer, Cham, Switzerland, May 8–10, 2019, pp. 286–306.
- God64. Roger Godement, *Domaines fondamentaux de groupes arithmétiques*, Sémin. Bourbaki 15 (1962/63), Exp. No. 257, 25 p. (1964), 1964.
- Gro98. Benedict H. Gross, *On the Satake isomorphism*, Galois representations in arithmetic algebraic geometry. Proceedings of the symposium, Durham, UK, July 9–18, 1996, Cambridge: Cambridge University Press, 1998, pp. 223–237 (English).
- HI80. Ki-ichiro Hashimoto and Tomoyoshi Ibukiyama, *On class numbers of positive definite binary quaternion Hermitian forms*, J. Fac. Sci. Univ. Tokyo Sect. IA Math. **27** (1980), no. 3, 549–601. MR 603952
- HW25. Arthur Herlédan Le Merdy and Benjamin Wesolowski, *Unconditional foundations for supersingular isogeny-based cryptography*, TCC 2025: 23rd Theory of Cryptography Conference, Part III (Aarhus, Denmark) (Benny Applebaum and Huijia (Rachel) Lin, eds.), Lecture Notes in Computer Science, vol. 16270, Springer, Cham, Switzerland, December 1–5, 2025, pp. 266–297.
- JZ23. Bruce W. Jordan and Yevgeny Zaytman, *Isogeny graphs of superspecial abelian varieties and Brandt matrices*, 2023.
- Kan97. Ernst Kani, *The number of curves of genus two with elliptic differentials*, J. Reine Angew. Math. **485** (1997), 93–121. MR 1442190
- KMM⁺25. Sabrina Kunzweiler, Luciano Maino, Tomoki Moriya, Christophe Petit, Giacomo Pope, Damien Robert, Miha Stopar, and Yan Bo Ti, *Radical 2-isogenies and cryptographic hash functions in dimensions 1, 2 and 3*, PKC 2025: 28th International Conference on Theory and Practice of Public Key Cryptography, Part III (Røros, Norway) (Tibor Jager and Jiaxin Pan, eds.), Lecture Notes in Computer Science, vol. 15676, Springer, Cham, Switzerland, May 12–15, 2025, pp. 265–299.
- Kot84. Robert E. Kottwitz, *Stable trace formula: Cuspidal tempered terms*, Duke Math. J. **51** (1984), 611–650 (English).
- Kot92. ———, *Points on some Shimura varieties over finite fields*, J. Amer. Math. Soc. **5** (1992), no. 2, 373–444.
- KV10. Markus Kirschmer and John Voight, *Algorithmic enumeration of ideal classes for quaternion orders*, SIAM Journal on Computing **39** (2010), no. 5, 1714–1747.
- LO98. Ke-Zheng Li and Frans Oort, *Moduli of supersingular abelian varieties*, Lect. Notes Math., vol. 1680, Berlin: Springer, 1998 (English).
- LPW09. David A. Levin, Yuval Peres, and Elizabeth L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009, With a chapter by James G. Propp and David B. Wilson. MR 2466937
- Mil79. J. S. Milne, *Points on Shimura varieties mod p* , Automorphic forms, representations and L-functions, Proc. Symp. Pure Math. Am. Math. Soc., Corvallis/Oregon 1977, Proc. Symp. Pure Math. 33, No. 2, 165–184 (1979), 1979.
- Mil86. ———, *Abelian varieties*, Arithmetic geometry (Storrs, Conn., 1984), Springer, New York, 1986, pp. 103–150. MR 861974

- ML98. Saunders Mac Lane, *Categories for the working mathematician*, second ed., Graduate Texts in Mathematics, vol. 5, Springer-Verlag, New York, 1998. MR 1712872
- MMP⁺23. Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski, *A direct key recovery attack on SIDH*, Advances in Cryptology – EUROCRYPT 2023, Part V (Lyon, France) (Carmit Hazay and Martijn Stam, eds.), Lecture Notes in Computer Science, vol. 14008, Springer, Cham, Switzerland, April 23–27, 2023, pp. 448–471.
- Mum66. David Mumford, *On the equations defining abelian varieties 1*, *Inventiones mathematicae* **1** (1966), no. 4, 287–354.
- Mum70. D. Mumford, *Abelian varieties*, 2nd edn. Oxford University Press, London, 1970. MR 282985
- MW16a. Colette Mœglin and Jean-Loup Waldspurger, *Stabilisation de la formule des traces tordue*, vol. 1, Springer International Publishing, 2016.
- MW16b. ———, *Stabilisation de la formule des traces tordue*, vol. 2, Springer International Publishing, 2016.
- Oh02. Hee Oh, *Uniform pointwise bounds for matrix coefficients of unitary representations and applications to Kazhdan constants*, *Duke Math. J.* **113** (2002), no. 1, 133–192 (English).
- Piz90. Arnold K Pizer, *Ramanujan graphs and Hecke operators*, *Bulletin of the American Mathematical Society* **23** (1990), no. 1, 127–137.
- PR94. Vladimir Platonov and Andrei Rapinchuk, *Algebraic groups and number theory. Transl. from the Russian by Rachel Rowen*, Pure Appl. Math., Academic Press, vol. 139, Boston, MA: Academic Press, 1994 (English).
- PR23. Aurel Page and Damien Robert, *Introducing clapoti(s): Evaluating the isogeny class group action in polynomial time*, *Cryptology ePrint Archive*, Paper 2023/1766, 2023.
- PW23. Aurel Page and Benjamin Wesolowski, *The supersingular endomorphism ring and one endomorphism problems are equivalent*, *Cryptology ePrint Archive*, Paper 2023/1399, 2023.
- PW24a. ———, *The supersingular endomorphism ring and one endomorphism problems are equivalent*, Advances in Cryptology – EUROCRYPT 2024, Part VI (Zurich, Switzerland) (Marc Joye and Gregor Leander, eds.), Lecture Notes in Computer Science, vol. 14656, Springer, Cham, Switzerland, May 26–30, 2024, pp. 388–417.
- PW24b. Aurel Page and Benjamin Wesolowski, *The supersingular endomorphism ring and one endomorphism problems are equivalent*, *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Springer, 2024, pp. 388–417.
- Rob23. Damien Robert, *Breaking SIDH in polynomial time*, Advances in Cryptology – EUROCRYPT 2023, Part V (Lyon, France) (Carmit Hazay and Martijn Stam, eds.), Lecture Notes in Computer Science, vol. 14008, Springer, Cham, Switzerland, April 23–27, 2023, pp. 472–503.
- Sat63. Ichirô Satake, *Theory of spherical functions on reductive algebraic groups over p-adic fields*, *Publ. Math., Inst. Hautes Étud. Sci.* **18** (1963), 229–293 (English).
- Shi24. Sug Woo Shin, *Weak transfer from classical groups to general linear groups*, *Essent. Number Theory* **3** (2024), no. 1, 19–62 (English).
- Sil92. A. Silverberg, *Fields of definition for homomorphisms of abelian varieties*, *J. Pure Appl. Algebra* **77** (1992), no. 3, 253–262. MR 1154704

- Tak17. Katsuyuki Takashima, *Efficient algorithms for isogeny sequences and their cryptographic applications*, Mathematical Modelling for Next-Generation Cryptography: CREST Crypto-Math Project, Springer, 2017, pp. 97–114.
- Tat66. J. Tate, *Endomorphisms of Abelian varieties over finite fields*, Invent. Math. **2** (1966), 134–144 (English).
- WM71. W. C. Waterhouse and J. S. Milne, *Abelian varieties over finite fields*, 1969 Number Theory Institute, Proc. Symp. Pure Math. 20, 53-64 (1971)., 1971.

A A brief introduction to categories

A general reference for this section is [ML98]. A category $\mathcal{C}$ consists of a class of *objects* and sets of *morphisms* between those objects. For every objects $x, y \in \mathcal{C}$, there is a set $\text{Hom}_{\mathcal{C}}(x, y)$ of morphisms, sometimes denoted $f : x \rightarrow y$, and there is a law of associative composition law for morphisms with compatible source and target, and an identity morphism $1_x \in \text{Hom}_{\mathcal{C}}(x, x)$ 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 $\text{End}_{\mathcal{C}}(x) = \text{Hom}_{\mathcal{C}}(x, x)$ of *endomorphisms* of x , the group $\text{Aut}_{\mathcal{C}}(x)$ of *automorphisms* of x .

Let $\mathcal{C}, \mathcal{D}$ be categories. A (*covariant*) *functor* $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ is an association of an object $\mathcal{F}(x) \in \mathcal{D}$ for every object $x \in \mathcal{C}$, and of a morphism $\mathcal{F}(f) \in \text{Hom}_{\mathcal{D}}(\mathcal{F}(x), \mathcal{F}(y))$ for every morphism $f \in \text{Hom}_{\mathcal{C}}(x, y)$, that respects composition and identities. Functors can be composed in the obvious way. For example, the identity functor $\text{id}_{\mathcal{C}}$ maps any category $\mathcal{C}$ to itself by sending objects and morphisms to themselves.

Let $\mathcal{F}, \mathcal{F}' : \mathcal{C} \rightarrow \mathcal{D}$ be functors. A *morphism of functors* $\psi : \mathcal{F} \rightarrow \mathcal{F}'$ is a collection $\psi = (\psi_x)_{x \in \mathcal{C}}$ of morphisms $\psi_x : \mathcal{F}(x) \rightarrow \mathcal{F}'(x)$ in $\mathcal{D}$ such that for every morphism $f : x \rightarrow y$ 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 functor $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ is an *equivalence of categories* if there exists a functor $\mathcal{H} : \mathcal{D} \rightarrow \mathcal{C}$ and isomorphisms of functors $\mathcal{H} \circ \mathcal{F} \cong \text{id}_{\mathcal{C}}$ and $\mathcal{F} \circ \mathcal{H} \cong \text{id}_{\mathcal{D}}$.