KLPT^2: Algebraic pathfinding in dimension two and applications

Wouter Castryck[1], Thomas Decru[1], Péter Kutas[2,4], Abel Laval[3], Christophe Petit[3,4], Yan Bo Ti[5,6]

[1] COSIC, KU Leuven, Belgium
[2] Faculty of Informatics, Eötvös Loránd University, Hungary
[3] Computer Science Department, Université Libre de Bruxelles, Belgium
[4] School of Computer Science, University of Birmingham, United Kingdom
[5] DSO National Laboratories, Singapore
[6] Temasek Laboratories, National University of Singapore, Singapore

Abstract. Following Ibukiyama, Katsura and Oort, all principally polarized superspecial abelian surfaces over $\overline{\mathbb{F}}_p$ can be represented by a certain type of 2×2 matrix g, having entries in the quaternion algebra $B_{p,\infty}$. We present a heuristic polynomial-time algorithm which, upon input of two such matrices g_1, g_2, finds a "connecting matrix" representing a polarized isogeny of smooth degree between the corresponding surfaces. Our algorithm should be thought of as a two-dimensional analog of the KLPT algorithm from 2014 due to Kohel, Lauter, Petit and Tignol for finding a connecting ideal of smooth norm between two given maximal orders in $B_{p,\infty}$.

The KLPT algorithm has proven to be a versatile tool in isogeny-based cryptography, and our analog has similar applications; we discuss two of them in detail. First, we show that it yields a polynomial-time solution to a two-dimensional analog of the so-called constructive Deuring correspondence: given a matrix g representing a superspecial principally polarized abelian surface, realize the latter as the Jacobian of a genus-2 curve (or, exceptionally, as the product of two elliptic curves if it concerns a product polarization). Second, we show that, modulo a plausible assumption, Charles–Goren–Lauter style hash functions from superspecial principally polarized abelian surfaces require a trusted set-up. Concretely, if the matrix g associated with the starting surface is known then collisions can be produced in polynomial time. We deem it plausible that all currently known methods for generating a starting surface indeed reveal the corresponding matrix. As an auxiliary tool, we present an explicit table for converting $(2,2)$-isogenies into the corresponding connecting matrix, a step for which a previous method by Chu required super-polynomial (but sub-exponential) time.

1 Introduction

In isogeny-based cryptography, the core problem is that of finding an explicit isogeny between two isogenous elliptic curves over a finite field. Here, "explicit" often implicates that the degree of the isogeny is powersmooth, or a power of

some small prescribed prime number ℓ. For reasons of both security and efficiency, almost all cryptographic constructions restrict their focus to supersingular elliptic curves. Famously, Deuring [23] proved that such curves are (essentially) in one-to-one correspondence with maximal orders in the quaternion algebra $B_{p,\infty}$ ramified at p and ∞; here p denotes the field characteristic. Under this correspondence, isogenies correspond to ideals, and the isogeny-finding problem translates into finding a connecting ideal between two given maximal orders $\mathcal{O}_0, \mathcal{O}_1 \subset B_{p,\infty}$, where one then aims for integral ideals I whose norm $\mathsf{n}(I)$ is powersmooth or a power of ℓ. Interestingly, this quaternion version of the isogeny-finding problem can be dealt with efficiently: in 2014, Kohel, Lauter, Petit, and Tignol [38] proposed a polynomial-time algorithm, now commonly known as the KLPT algorithm, for solving exactly this problem.

This result has had an amplitude of consequences, both constructive and destructive. For example, it breaks the second pre-image resistance of the Charles–Goren–Lauter (CGL) hash function [26] when using an untrusted set-up. A more recent cryptanalytic example is the break of pSIDH [15].[†] More fundamentally, it has led to a key insight in isogeny-based cryptography. Namely, on one hand, given a maximal order in $B_{p,\infty}$, one can use the KLPT algorithm to compute a corresponding supersingular elliptic curve $E/\overline{\mathbb{F}}_p$ in polynomial time: this is called the constructive Deuring correspondence and it is practical for cryptographically sized values of p [27]. On the other hand, the converse problem, namely computing the endomorphism ring of a given supersingular elliptic curve, is believed to be very hard. By now, we understand, in a heuristic-free way,[‡] that this is in fact the central hard problem in (supersingular) isogeny-based cryptography [53]. That is, the Deuring correspondence is a one-way function, and it allows for trapdoors, e.g., in the form of secret isogenies to an easy base curve. This has sparked many important constructions, where we highlight the Galbraith–Petit–Silva (GPS) signature scheme [29] and SQIsign [21].

Recently, the field of isogeny-based cryptography was shaken up by the use of higher-dimensional principally polarized abelian varieties. Earlier works such as [10, 16, 28, 50] studied these objects in their own right, but the real catalysts were the attacks on SIDH [9, 41, 46] which revealed a very powerful interplay between higher dimension and dimension one, i.e., the world of elliptic curves. Constructive applications followed soon, especially because the machinery allows for efficient representations of isogenies of arbitrary degree [47]. This has culminated in various new schemes, including SQIsign variants [3, 20, 25, 43] improving over their ancestor in terms of speed, compactness, and security foundations.

In view of these current trends, a higher-dimensional analog of the KLPT algorithm is an important lacking tool. The direct provocation for this research is the PhD thesis by Chu [16], who mentions this as a missing ingredient in a GPS-style signature scheme from superspecial principally polarized abelian surfaces. Here, the Deuring correspondence is to be replaced with a correspondence

[†]The attack from [15] does not invoke the KLPT algorithm directly; rather, it uses and adapts several of its subroutines.

[‡]Modulo a reliance on the generalized Riemann hypothesis.

due to Ibukiyama, Katsura and Oort [34] describing principal polarizations and polarized isogenies in terms of 2×2 matrices with entries in $B_{p,\infty}$. This missing analog of KLPT is exactly the central result of our paper.

Main contributions:

- *KLPT^2.* We provide the desired two-dimensional analog of the KLPT algorithm: upon input of two matrices g_1, g_2 representing two principally polarized superspecial abelian surfaces, it finds in heuristic polynomial time a "connecting matrix", representing a polarized isogeny of reduced degree N. We provide versions both for $N = \ell^e$ (where ℓ is a small prescribed prime number) and for N powersmooth. In both cases, the value of N achieved is in $O(p^{25+\varepsilon})$. The main techniques are the following. First we observe that if the matrices are in a very special form, then finding a connecting matrix is easy (Lemma 3.2). This turns the problem into a transformation problem: instead of connecting two matrices, try to transform one matrix into a standard form. An important challenge is to realize an output upper bound that only depends on p and not on the sizes of the matrices g_i. This is handled using certain size reductions and solving certain Diophantine equations. One noteworthy ingredient is an algorithm that given $a, c \in \mathcal{O}$ (where $\mathcal{O}$ is a maximal order in $B_{p,\infty}$) that have coprime norm, finds $b, d \in \mathcal{O}$ such that the reduced norm of the matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is a power of ℓ (or powersmooth) in $O(p^{3+\varepsilon})$. Very surprisingly, this problem is essentially equivalent to 1-dimensional KLPT (Section 3.2). We deem it very likely that the exponent $25 + \varepsilon$ can be improved, but leave such sharpenings for future work (and we note that future improvements on one-dimensional KLPT also improve our results).
- *Constructive Ibukiyama–Katsura–Oort (IKO) correspondence.* In Section 4.1, we describe an efficient algorithm for matrix-to-isogeny conversion for powersmooth degrees (thereby ticking off an unsurprising but missing ingredient in Chu's aforementioned signature scheme). Combined with our KLPT^2 algorithm, this yields a heuristic polynomial-time method for an analog of the constructive Deuring correspondence, described in Section 5.1: given a matrix g representing a principally polarized superspecial abelian surface, we explicitly realize this surface as either the Jacobian of a genus-2 curve, or as a product of elliptic curves if it concerns a product polarization.
- *Polynomial-time isogeny-to-matrix conversion for chains of $(2,2)$-isogenies.* In order to transfer more advanced applications of KLPT to dimension two, one also needs an efficient solution to the converse problem: given a principally polarized superspecial abelian surface A_1 with known matrix g_1, along with a polarized isogeny $\varphi : A_1 \to A_2$, find a matrix g_2 corresponding to A_2 along with a connecting matrix corresponding to φ. In the special case of $(2,2)$-isogenies, Appendix A provides an explicit look-up table from which the connecting matrix can simply be read off (and then g_2 follows right away). By mimicking a method due to Eisenträger, Hallgren, Lauter, Morrison and Petit [26, Algorithm 9], in Section 14 we lift this to

a polynomial-time solution for arbitrary chains of $(2,2)$-isogenies, through a repeated application of KLPT^2. In doing so, we by-pass the need for invoking Chu's super-polynomial (but sub-exponential) time algorithm for the principal ideal problem (PIP) in quaternionic matrix rings [16, Chapter 2].
– *Attacks on CGL style hash functions.* Section 5.4 describes our main application: an attack on CGL-type hash functions in dimension two, which were explicitly proposed in [10, 28, 39, 50], in the case of an untrusted set-up. This is very similar to the KLPT-based attacks [26] on the original CGL hash function. Concretely, if the starting surface comes with a known matrix g (which seems a fair assumption to make in all untrusted instantiations) then we can use the KLPT^2 algorithm to find collisions. Very similarly, this also breaks a verifiable delay function based on genus-2 curves [14], again if no trusted set-up is used.

We conclude this introduction by noting that our KLPT^2 algorithm can be seen as a constructive proof (modulo heuristic assumptions) of the dimension-two case of Jordan and Zaytman's recent result that the graph of $(\ell,\dots,\ell)$-isogenies between superspecial principally polarized abelian varieties is connected [36].

Outline

We provide some background on the KLPT algorithm, principal polarizations, the Ibukiyama–Katsura–Oort correspondence, and quaternionic matrices in Section 2. We describe our generalization of KLPT to dimension two in Section 3 and discuss basic matrix-to-isogeny and isogeny-to-matrix conversions in Section 4. We describe our applications of KLPT^2 in Section 5. In Section 6 we discuss some natural directions for further research. Finally, Appendix A contains our look-up table for converting $(2,2)$-isogenies to matrices.

Acknowledgments and support

We thank Jonathan Komada Eriksen, Aurel Page, Damien Robert and Gábor Ivanyos for sharing some helpful insights. Castryck and Decru are supported in part by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT – No. 101020788), as well as by CyberSecurity Research Flanders with reference number VR20192203. Together with Kutas, they are also supported by the CELSA alliance through the MaCro project. Decru is partly supported by Fonds de la Recherche Scientifique (FRS-FNRS) and by Fonds voor Wetenschappelijk Onderzoek (FWO) with reference number 1245025N. Kutas and Petit are partly supported by EPSRC through grant number EP/V011324/1. Kutas is supported by the Hungarian Ministry of Innovation and Technology NRDI Office within the framework of the Quantum Information National Laboratory Program and by the grant "EXCELLENCE-151343". Kutas is also supported by János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

2 Preliminaries

2.1 Deuring correspondence and the KLPT algorithm

For general background on quaternion algebras, the reader is referred to [51], for now let us briefly recall that a (rational) quaternion algebra B is a central simple algebra of dimension 4 over $\mathbb{Q}$. An order in B is a subring $\mathcal{O} \subset B$ containing 1 which has rank 4 as a $\mathbb{Z}$-module. An order is called maximal if it is maximal with respect to inclusion. The isomorphism class of a quaternion algebra is determined by its local behaviour: for which completions $\mathbb{Q}_v$ do we have that $B \otimes_{\mathbb{Q}} \mathbb{Q}_v$ is a division algebra? Such places v are called ramified.[†] In this paper we will be concerned with $B_{p,\infty}$, the unique quaternion algebra up to isomorphism which is ramified at ∞ and at a fixed prime number p (typically of cryptographic size). The reason is that the endomorphism ring of every supersingular elliptic curve over $\overline{\mathbb{F}}_p$ is isomorphic to a maximal order $\mathcal{O} \subset B_{p,\infty}$. Under this isomorphism, the degree of an endomorphism corresponds to the (reduced) norm $\mathsf{n}(u)$ of the corresponding quaternion u.

Example 2.1. If $p \equiv 3 \bmod 4$ then one can realize the quaternion algebra $B_{p,\infty}$ as $\mathbb{Q}\langle 1, i, j, k\rangle$ with $i^2 = -1$, $j^2 = -p$ and $k = ij = -ji$. The elliptic curve $E_0 : y^2 = x^3 + x$ with $j(E_0) = 1728$ is supersingular. Here $\mathrm{End}(E_0) \cong \mathcal{O}_0$ with

$$\mathcal{O}_0 = \left\langle 1, i, \frac{i+j}{2}, \frac{1+k}{2} \right\rangle.$$

One isomorphism $\tau : \mathcal{O}_0 \xrightarrow{\cong} \mathrm{End}(E_0)$ arises by letting $\tau(i) : (x, y) \mapsto (-x, \sqrt{-1}y)$ and $\tau(j) : (x, y) \mapsto (x^p, y^p)$. As mentioned: $\mathsf{n}(u) = \deg(\tau(u))$ for all $u \in \mathcal{O}_0$.

The Deuring correspondence [23] asserts that this turns into a categorical equivalence between supersingular elliptic curves defined over $\overline{\mathbb{F}}_p$ (up to Galois conjugation) and maximal orders in $B_{p,\infty}$ (up to isomorphism or, equivalently, up to conjugation). On the elliptic curve side, the non-zero morphisms are isogenies $\varphi : E_0 \to E_1$. On the quaternion side, such an isogeny φ corresponds to a rank-4 sub-$\mathbb{Z}$-module $I \subset B_{p,\infty}$ which is a left, resp. right, ideal of a maximal order $\mathcal{O}_0 \cong \mathrm{End}(E_0)$, resp. $\mathcal{O}_1 \cong \mathrm{End}(E_1)$. This ideal is then referred to as a connecting ideal of $\mathcal{O}_0$ and $\mathcal{O}_1$. Note that endomorphism rings can be embedded into $B_{p,\infty}$ in many ways: any embedding can be post-composed with conjugation. This warrants the notion of equivalent ideals: a left ideal $J \subset \mathcal{O}_0$ will be a right ideal of an order $\mathcal{O}_1' \cong \mathcal{O}_1$ if and only if there exists $\beta \in B_{p,\infty} \setminus \{0\}$ such that $J = I\beta$. On the geometry side, this corresponds to different isogenies $E_0 \to E_1$ connecting the same curves. The left ideals $I, J \subset \mathcal{O}_0$ are then said to be equivalent.

There is an explicit geometric view on Deuring's construction of the ideal I: it can be seen as the subset of $\mathrm{End}(E_0)$ that is obtained by post-composing φ with all elements of $\mathrm{Hom}(E_1, E_0)$. Thus I encodes the set of all isogenies $E_1 \to E_0$, and the norm of every element of I is divisible by the degree of φ. More precisely, it can be shown that $\deg(\varphi)$ equals $\mathsf{n}(I) = \gcd\{\mathsf{n}(u) \mid u \in I\}$, the norm of I.

[†] In the non-ramified cases we have $B \otimes_{\mathbb{Q}} \mathbb{Q}_v \cong \mathrm{M}_2(\mathbb{Q}_v)$.

The Deuring correspondence implies that there is a natural quaternion analog of the ℓ-isogeny pathfinding problem. Indeed, upon input of two maximal orders $\mathcal{O}_0, \mathcal{O}_1 \subset B_{p,\infty}$ connected by an ideal I, it amounts to finding an equivalent left ideal $J \subset \mathcal{O}_0$ of norm ℓ^e for some $e \geq 1$. An alternative viewpoint taking the geometric interpretation into account is as follows: when given one connecting ideal I, it is enough to find $\sigma \in I$ such that $\mathsf{n}(\sigma) = \mathsf{n}(I)\ell^e$. This is exactly the problem that is addressed by the KLPT algorithm [38].[†] It is then easy to check that $J = I\beta$ with $\beta = \bar{\sigma}/\mathsf{n}(I)$ is an equivalent ideal with norm ℓ^e. Geometrically, under the above identification of I with $\mathrm{Hom}(E_1, E_0)\varphi$, we can write $\sigma = \tau\varphi$ for a degree-ℓ^e isogeny $\tau : E_1 \to E_0$, and then J corresponds to $\mathrm{Hom}(E_1, E_0)\varphi\hat{\sigma}/\deg(\varphi) = \mathrm{Hom}(E_1, E_0)\varphi\hat{\varphi}\hat{\tau}/\deg(\varphi) = \mathrm{Hom}(E_1, E_0)\hat{\tau}$.

Remark 2.2. This is an important view on KLPT as we will need it in exactly the version where it finds an element of prescribed norm in a certain ideal.

The way KLPT proceeds is as follows. Using a simple trick one can assume knowledge of a left ideal $I \subset \mathcal{O}_0$ of prime norm N, so we would like to find an element $\sigma \in I$ of norm $N\ell^e$. First one finds $\gamma \in \mathcal{O}_0$ whose norm is $N\ell^{e_0}$. Now the ideal $J = \mathcal{O}_0 N + \mathcal{O}_0\gamma$ is going to have norm N. Locally, i.e., modulo N, I and J reduce to proper left ideals of the matrix ring $\mathrm{M}_2(\mathbb{Z}/N\mathbb{Z})$ and such ideals only differ by right-multiplication by an invertible element δ. Such a δ can be computed locally and lifted to $\mathcal{O}_0$ (using an explicit isomorphism between $\mathcal{O}_0/N\mathcal{O}_0$ and $\mathrm{M}_2(\mathbb{Z}/N\mathbb{Z})$) which implies that $\gamma\delta \in I$. Now the key is that δ is determined modulo N only, so what is left to do is choose an appropriate lifting such that $\mathsf{n}(\delta) = \ell^{e_1}$. This step is called *strong approximation* and in [38] it is carried out for special extremal orders $\mathcal{O}_0$, i.e., maximal orders containing an imaginary quadratic order with small discriminant (it can be modified to work for arbitrary orders, see [21, Section 5] and [15, Section 5]). Now it is clear that $\gamma\delta$ will fit our criteria with $e = e_0 + e_1$. The KLPT algorithm can guarantee that $\ell^e \in O(p^{3+\varepsilon})$.[‡]

2.2 Principally polarized abelian varieties

For detailed background, we refer to [6, 16, 34, 42]. An abelian variety over an algebraically closed field is a projective algebraic variety which is also an algebraic group. The notion generalizes that of an elliptic curve, which is an abelian variety of dimension one. However, for most uses (including in cryptography), the more relevant generalization is that of an abelian variety equipped with a *principal polarization*. Unfortunately, this notion does not admit a down-to-earth definition. Luckily, the exact construction is not really important for this paper, which is mostly algebraic in nature. Therefore, the reader who is unfamiliar with the notation and terminology below can just think of a polarization as a certain

[†] In other applications of KLPT one searches for connecting ideals having powersmooth norm, but the approach is entirely the same.

[‡] The original bound from KLPT is in the order of $p^{3.5}$, but an improvement due to Petit and Smith [45], reported in [13, Algorithm 13], reduces this to $p^{3+\varepsilon}$ as stated.

kind of isogeny (i.e. a finite surjective homomorphism) between A and a companion abelian variety $\hat{A}$ called its dual.[†] We include a formal definition, e.g., to allow the reader to verify the proof of Theorem 2.7 further down:

Definition 2.3. *A* polarization *on a g-dimensional abelian variety A is an isogeny of the form*

$$\begin{aligned} \lambda : A &\to \hat{A} = \mathrm{Pic}^0(A) \\ P &\mapsto [t_{-P}(D) - D] \end{aligned}$$

with D an ample divisor on A, where t_{-P} denotes point-wise translation by $-P$. It can be shown that $\deg(\lambda) = (D^g/g!)^2$ *with D^g the self-intersection number of D. If this degree is equal to one then the polarization is called* principal. *Write* $\mathrm{PPol}(A)$ *for the set of principal polarizations on A.*

The reason why the notion of a principally polarized abelian variety still generalizes that of an elliptic curve is that, in the latter case, there is a unique principal polarization, called the canonical polarization. It is given by the negated Abel–Jacobi map: $P \mapsto [(\infty) - (P)]$. The uniqueness typically no longer holds in higher dimension. This is notoriously true for *superspecial* abelian varieties, which are our main objects of interest. A g-dimensional superspecial abelian variety is a variety which — as an unpolarized variety — is isomorphic to a product of g supersingular elliptic curves. It can be shown that, for a fixed characteristic p, all such products are pairwise isomorphic as soon as $g \geq 2$ [48, Theorem 3.5]. However, this unique isomorphism class carries $\Theta(p^{g(g+1)/2})$ inequivalent principal polarizations (in the sense of Definition 2.6 below).

Definition 2.4. *A (polarized) isogeny between two principally polarized abelian varieties (A, λ_A) and (B, λ_B) is an isogeny $\varphi : A \to B$ that respects the polarizations, i.e., there exists a positive integer N for which the following diagram commutes*

$$\begin{array}{ccc} A & \xrightarrow{\varphi} & B \\ {\scriptstyle [N]\lambda_A}\downarrow & & \downarrow{\scriptstyle \lambda_B} \\ \hat{A} & \xleftarrow[\hat{\varphi}]{} & \hat{B} \end{array}$$

Here, $\hat{\varphi}$ is the dual isogeny, defined by taking inverse image divisors under φ. One has $\deg(\varphi) = N^g$*, and we call* $N = \mathrm{degrd}(\varphi)$ *the* reduced degree *of φ. If $N = 1$ then φ is called a (polarized) isomorphism; we write* $(A, \lambda_A) \cong (B, \lambda_B)$.

Remark 2.5. Given a principally polarized abelian variety (A, λ_A), an abelian variety B and an isogeny $\varphi : A \to B$, in general there does not exist a principal polarization $\lambda_B : B \to \hat{B}$ such that φ is polarized. If it does exist, then λ_B is unique and called the induced principal polarization. Assuming that $\deg(\varphi) = N^g$ for some integer N coprime with the field characteristic, a necessary and sufficient condition for existence [42, Proposition 13.8 and Remark 13.9] is that

[†]Not all isogenies $A \to \hat{A}$ are polarizations: a certain positivity condition should be satisfied. E.g., if λ is a polarization, then $-\lambda$ is not. See [17, pp. 6–7] for a discussion.

$\ker(\varphi)$ is a maximal isotropic subgroup of $A[N]$, where isotropic means that $e_{N,\lambda_A}(P,Q) = 1$ for all $P, Q \in A[N]$, with $e_{N,\lambda_A} : A[N] \times A[N] \to \mu_N$ the N-Weil pairing with respect to the principal polarization λ_A. In this case $\deg\mathrm{rd}(\varphi) = N$.

For any isogeny $\varphi : A \to B$ and any choice of principal polarizations λ_A, λ_B, it is natural to consider the *adjoint* isogeny

$$\tilde{\varphi} = \lambda_A^{-1} \hat{\varphi} \lambda_B : B \to A$$

with respect to λ_A, λ_B. If φ is polarized, then so is $\tilde{\varphi}$ and we have $\tilde{\varphi}\varphi = [\deg\mathrm{rd}(\varphi)]$ and $\varphi\tilde{\varphi} = [\deg\mathrm{rd}(\varphi)]$. In the context of elliptic curves, the adjoint isogeny can be identified with the dual isogeny $\hat{\varphi} : \hat{B} \to \hat{A}$ via the canonical polarization, with which any isogeny is compatible. This is not the case in higher dimensions. This forces us to make a clear distinction between the dual isogeny, which is independent from any polarization, and the adjoint isogeny, which depends on a choice of principal polarizations and which exhibits the common properties we are familiar with from the elliptic curve case.

If λ is a principal polarization on an abelian variety A, then the adjoint operator $\mathrm{Ros}_\lambda : \alpha \mapsto \tilde{\alpha} = \lambda^{-1}\hat{\alpha}\lambda$, defines an involution of $\mathrm{End}(A)$ called the Rosati involution (with respect to λ).

Definition 2.6. *Two principal polarizations λ_1 and λ_2 on an abelian variety A are said to be* equivalent *if $(A, \lambda_1) \cong (A, \lambda_2)$, i.e., there exists an automorphism α of A such that $\hat{\alpha}\lambda_1\alpha = \lambda_2$. We write $\mathrm{PPol}^0(A)$ for the set of principal polarizations on A up to equivalence.*

Turning our focus to dimension $g = 2$, we recall that principally polarized abelian surfaces can be classified as follows: they are isomorphic to either

- a product $E_1 \times E_2$ of two elliptic curves, equipped with the *product polarization*, coming from $D = (E_1 \times \{\infty\}) + (\{\infty\} \times E_2)$, or
- the Jacobian $\mathrm{Jac}(C)$ of a genus-2 curve, equipped with the canonical polarization, coming from $D = (u(C))$ with $u : C \hookrightarrow \mathrm{Pic}^0(C) \cong \mathrm{Jac}(C) : P \mapsto [(P) - (\infty)]$ the Abel–Jacobi map (where $\infty \in C$ denotes any base point).

With respect to product polarizations, the adjoint admits a very explicit description which follows, e.g., along the lines of [5, Proposition 4.10]. Consider four elliptic curves E_1, E_2, E_3, E_4 and assume we have an isogeny $\varphi : E_1 \times E_2 \to E_3 \times E_4$, not necessarily polarized. We can write this isogeny in matrix form:

$$\varphi : \begin{pmatrix} P \\ Q \end{pmatrix} \mapsto \begin{pmatrix} \alpha_{13} & \alpha_{23} \\ \alpha_{14} & \alpha_{24} \end{pmatrix} \begin{pmatrix} P \\ Q \end{pmatrix}$$

where each $\alpha_{ij} : E_i \to E_j$ is a homomorphism (an isogeny or the zero map) of elliptic curves. With respect to the product polarizations on $E_1 \times E_2$ and $E_3 \times E_4$, we have

$$\tilde{\varphi} : \begin{pmatrix} P \\ Q \end{pmatrix} \mapsto \begin{pmatrix} \hat{\alpha}_{13} & \hat{\alpha}_{14} \\ \hat{\alpha}_{23} & \hat{\alpha}_{24} \end{pmatrix} \begin{pmatrix} P \\ Q \end{pmatrix},$$

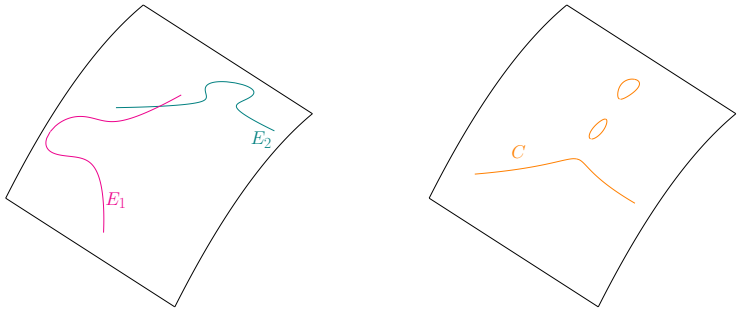

Fig. 1. A product of two elliptic curves on the left, and a genus-2 curve embedded in its Jacobian on the right.

so in this case the map $\varphi \mapsto \tilde{\varphi}$ can be thought of as a conjugate-transpose. The isogeny φ is polarized if and only if

$$\begin{pmatrix} \hat{\alpha}_{13} & \hat{\alpha}_{14} \\ \hat{\alpha}_{23} & \hat{\alpha}_{24} \end{pmatrix} \begin{pmatrix} \alpha_{13} & \alpha_{23} \\ \alpha_{14} & \alpha_{24} \end{pmatrix} = \begin{pmatrix} [N] & 0 \\ 0 & [N] \end{pmatrix}$$

for some positive integer N. This integer necessarily equals $\mathrm{degrd}(\varphi)$, so that $\deg(\varphi) = N^2$. In general, by [37, Corollary 64] and [30, Proposition 3.9], we have

$$\deg(\varphi) = (\deg \alpha_{13} + \deg \alpha_{14})(\deg \alpha_{23} + \deg \alpha_{24}) - \deg(\hat{\alpha}_{23}\alpha_{13} + \hat{\alpha}_{24}\alpha_{14}). \quad (1)$$

2.3 Ibukiyama–Katsura–Oort correspondence

In the remainder of the paper, we fix a prime $p \notin \{2, 3\}$ and a supersingular elliptic curve $E_0/\overline{\mathbb{F}}_p$. Let $B_{p,\infty}$ be the unique quaternion algebra (up to isomorphism) ramified exactly at p and infinity. Then, as mentioned, $\mathrm{End}(E_0)$ is isomorphic through the Deuring correspondence to a maximal order $\mathcal{O}_0$ of $B_{p,\infty}$. Define $A_0 = E_0 \times E_0$ and consider the product polarization λ_0. By our previous discussion, the endomorphism ring of A_0 is isomorphic to $\mathrm{M}_2(\mathcal{O}_0)$ and under this isomorphism the Rosati involution (i.e., the adjoint operator) with respect to λ_0 corresponds to the conjugate-transpose.

Recall that, considered without their polarizations, all superspecial abelian surfaces in characteristic p are isomorphic. Consequently, every principally polarized superspecial abelian surface (A, λ_A) is isomorphic to (A_0, λ) for some principal polarization λ on A_0. Explicitly, if $\varphi : A_0 \to A$ is an (unpolarized) isomorphism, then we can take $\lambda = \hat{\varphi}\lambda_A\varphi$. The following method due to Ibukiyama, Katsura and Oort can be used to represent a principal polarization λ on A_0 as a matrix with coefficients in $\mathcal{O}_0$. One considers the map

$$\begin{aligned} \mu : \mathrm{PPol}(A_0) &\to \mathrm{End}(A_0) \\ \lambda &\mapsto \lambda_0^{-1}\lambda, \end{aligned}$$

noting that the image $\lambda_0^{-1}\lambda$ can be identified with an element of $\mathrm{M}_2(\mathcal{O}_0)$.

Theorem 2.7. *The map μ is injective and its image, once transferred to the quaternion world through the Deuring correspondence, corresponds to*

$$\mathrm{Mat}(A_0) := \left\{ \begin{pmatrix} s & r \\ \bar{r} & t \end{pmatrix}, \quad s,t \in \mathbb{Z}_{>0}, r \in \mathcal{O}_0, st - r\bar{r} = 1 \right\} \quad \subset \ \mathrm{GL}_2(\mathcal{O}_0),$$

i.e., μ determines a bijection between $\mathrm{PPol}(A_0)$ *and* $\mathrm{Mat}(A_0)$.

Proof. This is [34, Corollary 2.9] specialized to principal polarizations (i.e., to ample divisors with self-intersection 2). □

Remark 2.8. An alternative way of specifying a principal polarization λ on A_0 is through the Rosati involution it induces on $\mathrm{M}_2(\mathcal{O}_0)$. This datum is very explicitly encoded in the matrix $g = \mu(\lambda)$:

$$\mathrm{Ros}_\lambda(\alpha) = \lambda^{-1}\hat{\alpha}\lambda = (\lambda_0^{-1}\lambda)^{-1}(\lambda_0^{-1}\hat{\alpha}\lambda_0)(\lambda_0^{-1}\lambda) = g^{-1}\,\mathrm{Ros}_{\lambda_0}(\alpha)g = g^{-1}\alpha^* g,$$

where we recall that the Rosati involution with respect to the product polarization λ_0 indeed amounts to the conjugate-transpose $-^*$.[†] Conversely, given black-box access to Ros_λ, one can reconstruct the matrix g via linear system solving, by considering $g\,\mathrm{Ros}_\lambda(b_i) = b_i^* g$ for a $\mathbb{Z}$-basis $b_1, \ldots, b_{16}$ of $\mathrm{M}_2(\mathcal{O}_0)$.

In a natural way, the matrix representation extends to polarized isogenies. Let $\lambda_1, \lambda_2 \in \mathrm{PPol}(A_0)$ be represented by matrices $g_1, g_2 \in \mathrm{Mat}(A_0)$ and let $\varphi : (A_0, \lambda_1) \to (A_0, \lambda_2)$ be a polarized isogeny of reduced degree N. Being an endomorphism of A_0, we can identify φ with a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$, and the property $\hat{\varphi}\lambda_2\varphi = N\lambda_1$ readily translates into

$$(\lambda_0^{-1}\hat{\varphi}\lambda_0)\lambda_0^{-1}\lambda_2\varphi = N\lambda_0^{-1}\lambda_1.$$

Using $\lambda_0^{-1}\lambda_i = g_i$ and identifying φ with γ, this can be rewritten as

$$\gamma^* g_2 \gamma = N g_1, \tag{2}$$

which is the chief equation of this entire paper. Conversely, whenever a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ satisfies (2), it determines a polarized isogeny $\varphi : (A_0, \lambda_1) \to (A_0, \lambda_2)$ of reduced degree N. In Section 4 we will discuss methods for converting polarized isogenies into matrices and vice versa.

We conclude with two remarks:

1. The equivalence relation for principal polarizations from Definition 2.6 naturally translates to the language of matrices as well: given $g_1, g_2 \in \mathrm{Mat}(A_0)$ encoding principal polarizations λ_1, λ_2 on A_0, we have

   $$\lambda_1 \sim \lambda_2 \quad \iff \quad \exists u \in \mathrm{GL}_2(\mathcal{O}_0), \quad u^* g_1 u = g_2.$$

   In this case, we say that the matrices are *congruent*; this terminology is taken from [32]. We then define $\mathrm{Mat}^0(A_0)$ as the set $\mathrm{Mat}(A_0)$ considered modulo congruence. Figure 2 summarizes the bijections that allow us to manipulate (isomorphism classes of) principally polarized superspecial abelian surfaces using only matrices with entries in $\mathcal{O}_0$.

[†] More generally, the adjoint of α with respect to $g_1 = \mu(\lambda_1)$, $g_2 = \mu(\lambda_2)$ is $g_1^{-1}\alpha^* g_2$.

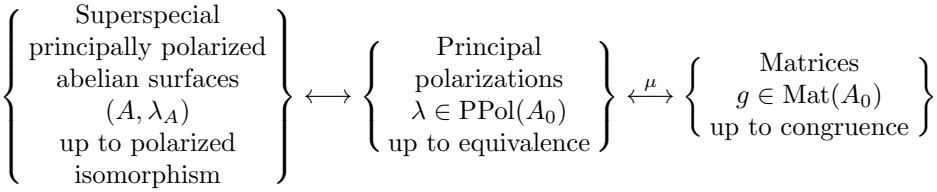

Fig. 2. Classification of principally polarized superspecial abelian surfaces

2. Every supersingular elliptic curve in characteristic p admits a model over $\mathbb{F}_{p^2}$, therefore the same is true for A_0 and the product polarization λ_0. When working with a model such that $\#A_0(\mathbb{F}_{p^2}) = (p \pm 1)^4$, as will be the case in practice, we know that all endomorphisms of A_0 are defined over $\mathbb{F}_{p^2}$ as well. Consequently, *every* principal polarization $\lambda = \lambda_0(\lambda_0^{-1}\lambda)$ is defined over $\mathbb{F}_{p^2}$. If (A, λ_A) is a superspecial principally polarized abelian surface defined over $\mathbb{F}_{p^2}$ such that $\#A(\mathbb{F}_{p^2}) = (p \pm 1)^4$, then it is $\mathbb{F}_{p^2}$-isomorphic to (A_0, λ) for some principal polarization λ on A_0. See [6] for an extended discussion.

2.4 Quaternionic matrices and determinants

When trying to define the determinant of a matrix

$$u = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{M}_2(B_{p,\infty}),$$

care is needed in view of the non-commutativity. Note that there is no ambiguity in the Hermitian case (i.e., for matrices that are invariant under taking the conjugate-transpose), which are always defined over a quadratic, hence commutative, subfield of $B_{p,\infty}$. In particular, it makes sense to consider $\det(uu^*)$ instead. Alternatively, one can consider the reduced norm $\mathcal{N}(u)$, defined as $\det(\iota(u \otimes 1))$, where

$$\iota : \mathrm{M}_2(B_{p,\infty}) \otimes_{\mathbb{Q}} \mathbb{C} \to \mathrm{M}_4(\mathbb{C})$$

is any isomorphism of $\mathbb{C}$-algebras. As the following lemma shows, this leads to the same result.

Lemma 2.9. $\det(uu^*) = \det(u^*u) = \mathsf{n}(a)\,\mathsf{n}(d) + \mathsf{n}(b)\,\mathsf{n}(c) - \mathsf{tr}(\bar{a}b\bar{d}c) = \mathcal{N}(u)$.

Proof. The first two equalities follow by explicit calculation. For the third equality we use that $\mathcal{N}(u) = \mathsf{n}(\Delta(u))$ by [24, Theorem 1, p. 146], where

$$\Delta(u) = \begin{cases} -bc & \text{if } a = 0, \\ ad - aca^{-1}b & \text{if } a \neq 0 \end{cases}$$

is the so-called Dieudonné determinant [24, Example 1, p. 133]. The statement follows by explicit calculation. (See [2, Example 2.5] for a related discussion.) □

One notable consequence of the above lemma is that the map $u \mapsto \det(uu^*)$ is multiplicative; indeed this property is immediate for $\mathcal{N}(-)$. Another interesting corollary, for which we could not find an explicit reference, is the following:

Corollary 2.10. *Let $E/\overline{\mathbb{F}}_p$ be a supersingular elliptic curve and let $\mathrm{End}(E) \cong \mathcal{O} \subset B_{p,\infty}$. Let $u \in \mathrm{End}(E^2)$, which via this isomorphism can be identified with an element of $\mathrm{M}_2(\mathcal{O})$. Then $\deg u = \mathcal{N}(u)$.*

Proof. This follows from (1) and an explicit calculation, using the identity

$$(\mathsf{n}(a)+\mathsf{n}(c))(\mathsf{n}(b)+\mathsf{n}(d)) - \mathsf{n}(\bar{a}b+\bar{c}d) = \mathsf{n}(a)\,\mathsf{n}(d) + \mathsf{n}(c)\,\mathsf{n}(b) - \mathsf{tr}(\bar{a}b\bar{d}c),$$

which in turn relies on the identity $\mathsf{n}(x+y) = \mathsf{n}(x) + \mathsf{n}(y) + \mathsf{tr}(x\bar{y})$. □

The multiplicativity also applies to the usual determinant when applied to Hermitian matrices. Up to sign, this is easy to see using that $\mathcal{N}(g) = \det(g)^2$ for any Hermitian matrix g. But the signs match as well:

Lemma 2.11. *Let $u, g, h \in \mathrm{M}_2(B_{p,\infty})$ where g, h are assumed Hermitian. Then*

- $\det(gh) = \det(g)\det(h)$,
- $\det(u^* g u) = \mathcal{N}(u)\det(g)$.

Proof. Using $\mathcal{N}(gh) = \mathcal{N}(g)\mathcal{N}(h)$, we know that either $\det(gh) = \det(g)\det(h)$ for all Hermitian g, h, or $\det(gh) = -\det(g)\det(h)$ for all Hermitian g, h (this can be seen, for instance, by working with indeterminate entries). But then we must be in the first case, since it applies whenever $g, h \in \mathrm{M}_2(\mathbb{Q})$. The second claim follows along similar lines. □

We end this section by showing that the "adjugate"[†] with respect to $\mathcal{N}(-)$ of an invertible matrix with entries in some subring $\mathcal{O} \subset B_{p,\infty}$ again has entries in $\mathcal{O}$.

Lemma 2.12. *If $u \in \mathrm{M}_2(\mathcal{O})$ is invertible in $\mathrm{M}_2(B_{p,\infty})$ then $u^{-1}\mathcal{N}(u) \in \mathrm{M}_2(\mathcal{O})$.*

Proof. Let $g \in \mathrm{M}_2(B_{p,\infty})$ be a Hermitian matrix, i.e., symmetric with respect to conjugate transpose. Then

$$g = \begin{pmatrix} s & r \\ \bar{r} & t \end{pmatrix}$$

where $s, t \in \mathbb{Q}$. If $\det(g) = st - \mathsf{n}(r) \neq 0$ then it is easy to see that g is invertible with inverse

$$\frac{1}{st-\mathsf{n}(r)} \begin{pmatrix} t & -r \\ -\bar{r} & s \end{pmatrix}$$

In particular, if $g \in \mathrm{M}_2(\mathcal{O})$ then also $g^{-1}\det(g) \in \mathrm{M}_2(\mathcal{O})$. Applying this to $g = uu^*$ yields $u^{*-1}u^{-1}\mathcal{N}(u) \in \mathrm{M}_2(\mathcal{O})$. Multiplying on the left with u^*, we get the desired result. □

[†]We intentionally avoid the word "adjoint", because the matrix $u^{-1}\mathcal{N}(u)$ should not be confused with the adjoint $\tilde{u}$ of u in the sense of Section 2.2. Firstly, the latter notion only makes sense when u describes a polarized isogeny with respect to certain principal polarizations. Secondly, in case it does make sense, we have $\tilde{u}u = u\tilde{u} = \mathrm{degrd}(u)\mathbb{I}_2$, whereas $u^{-1}\mathcal{N}(u)u = uu^{-1}\mathcal{N}(u) = \deg(u)\mathbb{I}_2 = \mathrm{degrd}(u)^2\mathbb{I}_2$ in view of Corollary 2.10. Recall that $\tilde{u} = g_1^{-1}u^*g_2$ when working with respect to principal polarizations associated with $g_1, g_2 \in \mathrm{Mat}(A_0)$.

3 Pathfinding in dimension 2

The goal of this section (and the main goal of the paper) is the description of an algorithm which solves the *algebraic pathfinding problem* in dimension 2. That is, upon input of $g_1, g_2 \in \mathrm{Mat}(A_0)$, the goal is to find a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ and a smooth integer N such that (2) holds. More precisely, we fix any small prime number ℓ and present the following solution, where N is a power of ℓ.

Theorem 3.1 (KLPT^2). *There exists a (heuristic) polynomial-time algorithm which upon input $g_1, g_2 \in \mathrm{Mat}(A_0)$ and a prime number $\ell \neq p$ returns $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that*

$$\gamma^* g_2 \gamma = \ell^e g_1$$

where $\ell^e \in O(p^{25+\varepsilon})$.

A proof-of-concept implementation of the algorithm can be found in:

`https://github.com/KLPT2/KLPT2`

Further down, in Theorem 3.14, we will present a variant for powersmooth N, which is often better-suited for (theoretical) applications.

Just as in the original KLPT algorithm, our algorithm relies on $\mathcal{O}_0$ being a special extremal order, i.e., we want it to contain a quadratic order whose discriminant has a very small absolute value [38, §2.3]. In fact, for simplicity, we just restrict to $p \equiv 3 \bmod 4$ and use the base curve $E_0 : y^2 = x^3 + x$ and maximal order $\mathcal{O}_0$ from Example 2.1. Note that the Gaussian integers $\mathbb{Z}[i]$ are contained in $\mathcal{O}_0$, so this order is of the desired kind.

3.1 Finding connecting matrices

Our proof strategy for Theorem 3.1 is based on the following lemma.

Lemma 3.2. *Let $h_1, h_2 \in \mathrm{M}_2(\mathcal{O}_0)$ be Hermitian matrices with equal upper-left entries and equal determinants, i.e., we have*

$$h_1 = \begin{pmatrix} D & r_1 \\ \bar{r}_1 & t_1 \end{pmatrix}, \quad h_2 = \begin{pmatrix} D & r_2 \\ \bar{r}_2 & t_2 \end{pmatrix}$$

for $D, t_1, t_2 \in \mathbb{Z}, r_1, r_2 \in \mathcal{O}_0$ such that $Dt_1 - \mathsf{n}(r_1) = Dt_2 - \mathsf{n}(r_2)$. Then for

$$\tau = \begin{pmatrix} D & r_1 - r_2 \\ 0 & D \end{pmatrix}$$

we have $\tau^* h_2 \tau = h_1$.

Proof. One calculates that

$$\begin{pmatrix} D & 0 \\ \bar{r}_1 - \bar{r}_2 & D \end{pmatrix} \begin{pmatrix} D & r_2 \\ \bar{r}_2 & t_2 \end{pmatrix} \begin{pmatrix} D & r_1 - r_2 \\ 0 & D \end{pmatrix} = \begin{pmatrix} D^3 & D^2 r_1 \\ D^2 \bar{r}_1 & D(\mathsf{n}(r_1) - \mathsf{n}(r_2) + Dt_2) \end{pmatrix},$$

so the only thing left to show is that $D(\mathsf{n}(r_1) - \mathsf{n}(r_2) + Dt_2) = D^2 t_1$. But this is true exactly because of the condition $Dt_1 - \mathsf{n}(r_1) = Dt_2 - \mathsf{n}(r_2)$. □

Note that in the above lemma we do not impose $\det(h_1) = \det(h_2) = 1$, so this is not always a special case of (2). We only want the two determinants to be equal, for reasons that will become apparent soon.

When given g_1, g_2, our goal is to transform them in a fashion such that Lemma 3.2 becomes applicable. This is aided by the following lemma:

Lemma 3.3. *Assume that $\delta^* g_2 \delta = N u^* g_1 u$ with $N \in \mathbb{Z}$, $u, \delta \in \mathrm{M}_2(\mathcal{O}_0)$. Then there exists $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\gamma^* g_2 \gamma = N\,\mathcal{N}(u)^2 g_1$.*

Proof. One can choose $\gamma = \delta u^{-1}\,\mathcal{N}(u)$. The equality $\gamma^* g_2 \gamma = N\,\mathcal{N}(u)^2 g_1$ is clearly satisfied and Lemma 2.12 implies that $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$. □

This naturally leads to the following plan for solving the problem $\gamma^* g_2 \gamma = \ell^e g_1$. Namely, given $g \in \mathrm{Mat}(A_0)$ we want to find $u \in \mathrm{M}_2(\mathcal{O}_0)$ with the following properties:

- $\mathcal{N}(u) = \ell^{e_1}$ where e_1 does not depend on g (but u does).
- The top left entry of $u^* g u$ is ℓ^{e_2}, where e_2 does not depend on g.

How does this solve our initial problem? First we transform g_1 and g_2 with an appropriate u_1 and u_2 in the above fashion. Then we invoke Lemma 3.2 as by design the two sides have the same top left entry and the same determinant, by Lemma 2.11. This yields a matrix $\tau \in \mathrm{M}_2(\mathcal{O}_0)$ such that

$$\tau^* u_2^* g_2 u_2 \tau = \ell^{2e_2} u_1^* g_1 u_1. \tag{3}$$

We can then apply Lemma 3.3 with $\delta = u_2\tau$ to return

$$\gamma = u_2 \tau u_1^{-1}\,\mathcal{N}(u_1), \tag{4}$$

which has reduced degree $\ell^{2(e_1+e_2)}$.

Remark 3.4. Although our approach is purely algebraic, it is instructive to understand what happens on the geometry side. For $i = 1, 2$, let λ_i be the principal polarization on A_0 associated with g_i under the IKO correspondence. In general, the Hermitian matrix $u_i^* g_i u_i$ does not correspond to a principal polarization on A_0 (as its determinant is $\mathcal{N}(u_i) = \ell^{e_1}$), yet it still corresponds to a polarization λ_i', which is the pull-back of λ_i under the endomorphism u_i. This leads to the outer squares in the following diagram:

$$\begin{array}{ccccccc}
A_0 & \xleftarrow{u_1} & A_0 & \xrightarrow{\tau} & A_0 & \xrightarrow{u_2} & A_0 \\
{\scriptstyle D^2\lambda_1}\downarrow & & \downarrow{\scriptstyle D^2\lambda_1'} & & {\scriptstyle \lambda_2'}\downarrow & & \downarrow{\scriptstyle \lambda_2} \\
\hat{A}_0 & \xrightarrow[\hat{u}_1]{} & \hat{A}_0 & \xleftarrow[\hat{\tau}]{} & \hat{A}_0 & \xleftarrow[\hat{u}_2]{} & \hat{A}_0
\end{array} \tag{5}$$

The polarizations of the left square were scaled by $D^2 = \ell^{2e_2}$ in order to be compatible with the middle square, which refers to our application of Lemma 3.2: the matrices g_i are crafted such that $D^2\lambda_1'$ is the pull-back of λ_2' along the easy

endomorphism τ, depicted in blue. Finally, Lemma 3.3 explains how to "flip" the left square. Indeed, we can naturally extend it to

$$\begin{array}{ccccc} A_0 & \xleftarrow{u_1} & A_0 & \xleftarrow{u_1^{-1}\mathcal{N}(u_1)} & A_0 \\ {\scriptstyle D^2\lambda_1}\downarrow & & \downarrow{\scriptstyle D^2\lambda_1'} & & \downarrow{\scriptstyle D^2\mathcal{N}(u_1)^2\lambda_1} \\ \hat{A}_0 & \xrightarrow[\hat{u}_1]{} & \hat{A}_0 & \xrightarrow[\hat{u}_1^{-1}\mathcal{N}(u_1)]{} & \hat{A}_0 \end{array}$$

Thus, by substituting the (flipped version of the) block on the right for the left square in (5), we obtain a diagram as in Definition 2.4, showing that λ_2 pulls back to $D^2\mathcal{N}(u_1)^2\lambda_1$ under $\gamma = u_2\tau u_1^{-1}\mathcal{N}(u_1)$. That is, γ is a polarized isogeny of reduced degree $D^2\mathcal{N}(u_1)^2$, as wanted.

So now our focus is on a single

$$g = \begin{pmatrix} s & r \\ \bar{r} & t \end{pmatrix} \in \mathrm{Mat}(A_0),$$

where along the way we will explicitly bound ℓ^{e_1}, ℓ^{e_2} by appropriate constants. First let us calculate what the top left entry of u^*gu is.

Lemma 3.5. *Let $u = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. Then the top left corner of u^*gu is given by*

$$s' := s \cdot \mathsf{n}(a) + t \cdot \mathsf{n}(c) + \mathsf{tr}(\bar{c}\bar{r}a)$$

and the bottom right corner is given by

$$t' := s \cdot \mathsf{n}(b) + t \cdot \mathsf{n}(d) + \mathsf{tr}(\bar{b}\bar{r}d).$$

Proof. This follows from a simple calculation. □

Note in particular that the top left corner s' only depends on a and c (likewise, the bottom right corner t' only depends on b and d). This motivates the following rough strategy:

1. Find $a, c \in \mathcal{O}_0$ such that s' is a fixed power of ℓ.
2. Given a, c, find values for $b, d \in \mathcal{O}_0$ such that the reduced norm $\mathcal{N}(u)$ is another fixed power of ℓ.

We first concentrate on Step 2, then come back to Step 1 in Section 3.4.

3.2 Controlling the reduced norm

Let us be given non-zero $a, c \in \mathcal{O}_0$, where we assume that $\mathsf{n}(a)$ and $\mathsf{n}(c)$ are coprime; this will indeed be ensured. The goal of this section is to find $x, y \in \mathcal{O}_0$ such that

$$\mathcal{N}\begin{pmatrix} a & x \\ c & y \end{pmatrix} = \mathsf{n}(a)\,\mathsf{n}(y) + \mathsf{n}(c)\,\mathsf{n}(x) - \mathsf{tr}(\bar{a}x\bar{y}c) = \ell^{e_0}$$

for some fixed power ℓ^{e_0} (we will eventually have $e_1 = 2e_0$). We will not be solving this Diophantine equation directly. Instead, we will show that it amounts to a pathfinding problem in dimension 1 and invoke the standard KLPT algorithm.

We view $\mathcal{O}_0^2$ as a free right $\mathcal{O}_0$-module of rank 2, together with a quadratic module structure given by $Q((x,y)) = \mathsf{n}(a)\,\mathsf{n}(y) + \mathsf{n}(c)\,\mathsf{n}(x) - \mathsf{tr}(\bar{a}x\bar{y}c)$. Since $\mathcal{N}$ is a norm, it is clear that Q is either positive or semi-positive definite, and we will informally refer to it as a norm on $\mathcal{O}_0^2$ (in fact, it is a semi-norm). The first observation is that Q is identically zero on the free rank-1 submodule $(a,c)\mathcal{O}_0$. Actually, a simple calculation shows that every element of $(a,c)\mathcal{O}_0$ is orthogonal to every element in $\mathcal{O}_0^2$. The following lemma reveals a complementary submodule.

Lemma 3.6. *Let $M_1 = (a,c)\mathcal{O}_0$. Furthermore, let α, β be integers such that $\alpha\,\mathsf{n}(a) + \beta\,\mathsf{n}(c) = 1$. Let $M_2 = (\beta\,\mathsf{n}(c)a, -\alpha\,\mathsf{n}(a)c)B_{p,\infty} \cap \mathcal{O}_0^2$. Then M_2 is a right $\mathcal{O}_0$-module and $M_1 \oplus M_2 = \mathcal{O}_0^2$.*

Proof. M_2 is a right $\mathcal{O}_0$-module as it is the intersection of two right $\mathcal{O}_0$-modules. Any element $u \in M_2$ can be written as $(\beta\,\mathsf{n}(c)a, -\alpha\,\mathsf{n}(a)c)z$ where $z \in B_{p,\infty}$. It is easy to check that

$$Q(u) = \mathsf{n}(ac)\,\mathsf{n}(z), \tag{6}$$

so that only the 0 vector has 0 norm, hence its intersection with M_1 is trivial.

Now we show why $M_1 + M_2 = \mathcal{O}_0^2$. It is enough to show that $M_1 + M_2$ contains $(1,0)$ and $(0,1)$. One has that

$$(a,c)\alpha\bar{a} + (\beta\,\mathsf{n}(c)a, -\alpha\,\mathsf{n}(a)c)\frac{1}{\mathsf{n}(a)}\bar{a} = (1,0)$$

because $\alpha\,\mathsf{n}(a) + \beta\,\mathsf{n}(c) = 1$. It is easy to see that the second term is indeed in M_2. Our claim follows from similarly noting that

$$(a,c)\beta\bar{c} - (\beta\,\mathsf{n}(c)a, -\alpha\,\mathsf{n}(a)c)\frac{1}{\mathsf{n}(c)}\bar{c} = (0,1). \qquad \square$$

We are now ready to prove the main result of this subsection:

Proposition 3.7. *The module M_2 is $\mathsf{n}(c)$-homothetic to the right $\mathcal{O}_0$-ideal $I = \mathsf{n}(c)\mathcal{O}_0 + a\bar{c}\mathcal{O}_0$. More precisely, the map*

$$\begin{gathered}\tau : M_2 \to I \\ (\beta\,\mathsf{n}(c), -\alpha c\bar{a})o_1 + (\beta a\bar{c}, -\alpha\,\mathsf{n}(a))o_2 \mapsto \mathsf{n}(c)o_1 + a\bar{c}o_2, \quad o_1, o_2 \in \mathcal{O}_0\end{gathered}$$

is a well-defined isomorphism of right $\mathcal{O}_0$-modules such that $\mathsf{n}(\tau(m)) = \mathsf{n}(c)Q(m)$ for all $m \in M_2$.

Proof. First note that $(\beta\,\mathsf{n}(c), -\alpha c\bar{a}) \in M_2$ because

$$(\beta\,\mathsf{n}(c), -\alpha c\bar{a}) = (\beta\,\mathsf{n}(c)a, -\alpha\,\mathsf{n}(a)c)a^{-1}.$$

Likewise, we find that $(\beta a\bar{c}, -\alpha\,\mathsf{n}(a)) \in M_2$. Next, observe that

$$Q((\beta\,\mathsf{n}(c), -\alpha c\bar{a})o_1 + (\beta a\bar{c}, -\alpha\,\mathsf{n}(a))o_2)$$

can be rewritten as

$$\begin{aligned} Q\left((\beta\,\mathsf{n}(c)a, -\alpha\,\mathsf{n}(a)c)(a^{-1}o_1 + c^{-1}o_2)\right) &= \mathsf{n}(ac)\,\mathsf{n}(a^{-1}o_1 + c^{-1}o_2) \\ &= \frac{\mathsf{n}(a\,\mathsf{n}(c))}{\mathsf{n}(c)}\,\mathsf{n}(a^{-1}o_1 + c^{-1}o_2) \\ &= (1/\,\mathsf{n}(c))\,\mathsf{n}(a\,\mathsf{n}(c)a^{-1}o_1 + a\,\mathsf{n}(c)c^{-1}o_2) \\ &= (1/\,\mathsf{n}(c))\,\mathsf{n}(\mathsf{n}(c)o_1 + a\bar{c}o_2), \end{aligned}$$

where we have used (6) in the first step. This almost proves the proposition. Namely, it shows that τ defines an $\mathsf{n}(c)$-homothetic isomorphism between the module $M_2' \subset M_2$ generated by $(\beta\,\mathsf{n}(c), -\alpha c\bar{a}), (\beta a\bar{c}, -\alpha\,\mathsf{n}(a))$ and I. Note that it is a priori unclear that τ is a well-defined map, let alone an isomorphism, because the decomposition with respect to these generators may not be unique. However, this again follows from the homothetic property: the element $(0, 0)$, however decomposed, must map to an element of norm 0, hence it must map to 0. A similar argument also proves injectivity, while surjectivity comes for free.

So it remains to argue that $M_2' = M_2$. Assume $M_2' \subsetneq M_2$ and note that, by allowing for $o_1, o_2 \in B_{p,\infty}$, we can extend the domain of τ to M_2, still ending up with a well-defined injective morphism. The image $\tau(M_2)$ is a fractional right $\mathcal{O}_0$-ideal strictly containing I. Since $\mathsf{n}(I) = \gcd(\mathsf{n}(c)^2, \mathsf{n}(a\bar{c})) = \mathsf{n}(c)$, this means that $\tau(M_2)$ must contain an element whose norm is not an integer multiple of $\mathsf{n}(c)$. But this means that M_2 contains an element at which Q takes a value outside the integers: a contradiction. □

Why is Proposition 3.7 important? The KLPT algorithm can find an element $\omega \in I$ such that $\mathsf{n}(\omega) = \mathsf{n}(c)\ell^{e_0}$ where $\ell^{e_0} \in O(p^{3+\varepsilon})$. This element can be written as $\mathsf{n}(c)o_1 + a\bar{c}o_2$ (in polynomial time) because $\mathsf{n}(c)$ and $a\bar{c}$ are generators of I as a right $\mathcal{O}_0$-module. Proposition 3.7 implies that the norm with respect to Q of the vector $(\beta\,\mathsf{n}(c), -\alpha c\bar{a})o_1 + (\beta a\bar{c}, -\alpha\,\mathsf{n}(a))o_2$ is exactly ℓ^{e_0}, and this is what we wanted to achieve. Turning back to the language of matrices:

$$\begin{pmatrix} a & \beta\,\mathsf{n}(c)o_1 + \beta a\bar{c}o_2 \\ c & -\alpha c\bar{a}o_1 - \alpha\,\mathsf{n}(a)o_2 \end{pmatrix}$$

is a way of completing the first column $\begin{pmatrix} a & c \end{pmatrix}^T$ to a 2×2 matrix with reduced norm ℓ^{e_0}.

Remark 3.8. This can be simplified: from the proof of Lemma 3.6 it follows that the above matrix can be rewritten as

$$\begin{pmatrix} a & o_1 \\ c & -o_2 \end{pmatrix} \cdot \begin{pmatrix} 1 & -\alpha\bar{a}o_1 + \beta\bar{c}o_2 \\ 0 & 1 \end{pmatrix}.$$

Since the second factor is an element of $\mathrm{GL}_2(\mathcal{O}_0)$, it is equally fine to work with the matrix on the left: its reduced norm is also equal to ℓ^{e_0}.

Remark 3.9. Under the Deuring correspondence, the right ideal I corresponds to an incoming isogeny $\psi : E \to E_0$. The dual of this isogeny is the degree-$\mathsf{n}(c)$ factor of $c\bar{a}$ emanating from E_0, which can also be described as a push-forward:

$$\begin{array}{ccc} E_0 & \xrightarrow{a} & E_0 \\ \downarrow{\scriptstyle c} & \swarrow{\scriptstyle c\bar{a}} & \downarrow{\scriptstyle [a]_* c = \hat{\psi}} \\ E_0 & & E \end{array}$$

In other words: $\psi = \widehat{[a]_* c}$.

3.3 Reduction of the matrix g

Thanks to the previous subsection, our task has been (essentially) reduced to finding $a, c \in \mathcal{O}_0$ in such a way that the top-left entry

$$s' = K((a,c)) := s \cdot \mathsf{n}(a) + t \cdot \mathsf{n}(c) + \mathsf{tr}(\bar{c}\bar{r}a) \tag{7}$$

of $u^* g u$ is some fixed power of ℓ, only depending on p. Moreover, we want to make sure that $\mathsf{n}(a)$ and $\mathsf{n}(c)$ are non-zero and coprime. Again, we see that K is a quadratic form on $\mathcal{O}_0^2$, but now we will just view the latter as a free $\mathbb{Z}$-module of rank 8, and analyze it as such:

Proposition 3.10. *The quadratic form K is positive definite and has determinant $(p/4)^4$.*

Proof. First we prove that K is positive definite. To see that it is definite, let $(a,c) \in \mathcal{O}_0^2 \setminus \{(0,0)\}$. It is easy to check that this can be seen as the first column of a matrix u with non-zero reduced norm. Assume that $u^* g u$ has the form $\left(\begin{smallmatrix} 0 & r' \\ \bar{r}' & t' \end{smallmatrix}\right)$, then $\det(u^* g u) = -\mathsf{n}(r') \leq 0$. But from Lemma 2.11 we find that $\det(u^* g u) = \mathcal{N}(u) \det(g) > 0$: a contradiction. This proves that K does not have a nontrivial zero. Since K is an 8-dimensional integer quadratic form this implies that K is not indefinite as every indefinite quadratic form in dimension at least 5 is isotropic, as wanted. Furthermore, K cannot be negative definite since $s > 0$.[†]

Writing $r = r_1 + r_2 i + r_3 j + r_4 k$, an explicit calculation shows that the matrix of the quadratic form K with respect to the basis $(1,0), (i,0), \ldots, (0,k)$ of $B_{p,\infty}^2$ is as follows:

$$\begin{pmatrix} s & 0 & 0 & 0 & r_1 & -r_2 & -pr_3 & -pr_4 \\ 0 & s & 0 & 0 & r_2 & r_1 & -pr_4 & pr_3 \\ 0 & 0 & sp & 0 & pr_3 & pr_4 & pr_1 & -pr_2 \\ 0 & 0 & 0 & sp & pr_4 & -pr_3 & pr_2 & pr_1 \\ r_1 & r_2 & pr_3 & pr_4 & t & 0 & 0 & 0 \\ -r_2 & r_1 & pr_4 & -pr_3 & 0 & t & 0 & 0 \\ -pr_3 & -pr_4 & pr_1 & pr_2 & 0 & 0 & tp & 0 \\ -pr_4 & pr_3 & -pr_2 & pr_1 & 0 & 0 & 0 & tp \end{pmatrix}$$

[†]Alternatively, the positive-definiteness follows from [34, Proposition 2.8] when applied to the pull-back polarization of the principal polarization λ corresponding to g (under the isogeny corresponding to u).

One can check that the determinant of this matrix is $p^4(st - \mathsf{n}(r))^4 = p^4$. Any matrix of base change between a $\mathbb{Z}$-basis of $\mathcal{O}_0^2$ and the above basis has determinant $1/16$, leading to the desired result. □

The goal of this subsection is to describe an intermediate step, where we wish to find a transformation matrix u making s' as small as possible. This can be achieved through lattice reduction: using Proposition 3.10, we see that (7) expresses s' as the squared-Euclidean length of a vector in a lattice in $\mathbb{R}^8$ having volume $(p/4)^2$.[†] Using the usual Minkowski bound we get that there exists one vector with

$$s' < 4\left(\frac{(p/4)^2}{\nu_8}\right)^{1/4} < \frac{3}{2}\sqrt{p}$$

where $\nu_8 = \pi^4/24$ denotes the volume of an 8-dimensional unit ball. In practice we can find corresponding a, c by using the Hermite–Korkine–Zolotarev lattice reduction algorithm (HKZ).

Once a, c realizing a small value of s' are found, we can complement them with b, d using the KLPT algorithm, as described in the previous subsection. However, remember that we want $\mathsf{n}(a)$ and $n(c)$ to be coprime for this. Furthermore, to simplify the analysis in Theorem 3.12 below, we will want s' to be a prime different from 2 and ℓ. This forces us to slightly enlarge the above bound. Assuming a sufficiently random behavior of the integers $K((a, c))$, we should be able to find such an $s' < R$ as soon as

$$\#\left\{\ (a,c) \in \mathcal{O}_0^2 \;\middle|\; K((a,c)) < R\ \right\} \geq \frac{\pi^2}{6} \cdot \frac{2\ell}{\varphi(2\ell)} \cdot \ln R,$$

where the factors on the right account for the conditions $\gcd(\mathsf{n}(a), \mathsf{n}(c)) = 1$, $2\ell \nmid s'$, and s' prime, respectively. Applying the Gaussian heuristic, we can approximate the left-hand side by $\nu_8 R^4/(p/4)^2$. Thus we see that

$$R = \sqrt{p}(\ln p)^{1/4}$$

should be good enough, where we took some margin, for the simplicity of this expression, to leave room for retrial, as well as to back-up for potential biases in the Gaussian heuristic and in the distribution of quaternionic norms, e.g., of the kind [7, Conjecture 6].

Despite the smallness of s', the other entries of $u^* g u$ may become quite large. Thus, we use an extra transformation to keep these values contained. For this we can use a matrix of the form $\left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right)$. To enlighten notation, let us explain this step directly on g, rather than on $u^* g u$:

$$\left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right)^* g \left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right) = \begin{pmatrix} s & \alpha s + r \\ \bar{\alpha} s + \bar{r} & \mathsf{n}(\alpha) s + \mathsf{tr}(\bar{\alpha} r) + t \end{pmatrix}$$

[†] E.g., this follows via the Cholesky decomposition of the matrix in the proof of Proposition 3.10.

The main observation here is twofold. First s does not change. Second r changes to $\alpha s + r$, thus we can attain anything in $\mathcal{O}_0$ that is congruent to r modulo s. In particular we can ensure that the coordinates $r_i, i = 1, \ldots, 4$ of r with respect to the basis from Example 2.1 satisfy $|r_i| \leq s/2$. This implies that $\mathsf{n}(r) \leq s^2(p+5)/8$. Note that $\left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right) \in \mathrm{GL}_2(\mathcal{O}_0)$, hence we do not have to worry about the reduced norm in this step.

Applying this to

$$u^* g u = \begin{pmatrix} s' & r' \\ \bar{r}' & t' \end{pmatrix},$$

first note that

$$s't' - \mathsf{n}(r') = \mathcal{N}(u) = \ell^{e_0} \in O(p^{3+\varepsilon})$$

by the KLPT step, where we have used Lemma 2.11. Thus from $\mathsf{n}(r') \leq s'^2(p+5)/8$ and $s' \leq \sqrt{p}(\ln p)^{1/4}$, one finds that $t' \leq \ell^{e_0}/s' + s'(p+5)/8 \in O(p^{3+\varepsilon}/s')$.

Finally, we will also want that $\ell \nmid t'$, or equivalently $\ell \nmid \mathsf{n}(r')$. This is easy to achieve by slightly tweaking α if needed. Indeed, one easily checks that, by relaxing the bounds $|r_i| \leq s/2$ to $|r_i| \leq s$, it can be ensured that $\mathsf{tr}(r') \not\equiv -1 \bmod \ell$. Then, if it so happens that $\ell \mid \mathsf{n}(r')$, an extra transformation using $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ will fix this issue.

In summary, by applying a suitable transformation $g \leftarrow u^* g u$, we can reduce to the case where g has bounded entries satisfying some non-divisibility conditions, at the cost of increasing the determinant from 1 to a power of ℓ. For clarity we give a definition for this case, while adding in another heuristic assumption, which should be satisfied with overwhelming probablity:

Definition 3.11. *A matrix*

$$g = \begin{pmatrix} s & r \\ \bar{r} & t \end{pmatrix}, \quad s, t \in \mathbb{Z}_{>0}, r \in \mathcal{O}_0, st - r\bar{r} > 0$$

is called ℓ-reduced *if*

- $\det(g) = st - \mathsf{n}(r) = \ell^{e_0}$ *for some* $e_0 \geq 0$,
- $s \leq \sqrt{p}(\ln p)^{1/4}$ *is a prime number not dividing* $2\ell t$,
- $\mathsf{n}(r) \leq s^2 p$ *is not a multiple of* ℓ.

The extra assumption is $s \nmid t$. Note that the conditions imply $s \nmid \mathsf{n}(r)$ and $\ell \nmid t$.

3.4 Controlling the top-left entry and finalizing the algorithm

Starting from a reduced matrix g as in Definition 3.11, with determinant $\ell^{e_0} \in O(p^{3+\varepsilon})$, we now show how to find a matrix u, of ℓ-power reduced norm, such that $u^* g u$ has a top left corner equal to ℓ^{e_2} for some $e_2 \geq 0$. As discussed before, this amounts to solving the Diophantine equation

$$\ell^{e_2} = s\,\mathsf{n}(a) + t\,\mathsf{n}(c) + \mathsf{tr}(\bar{c}\bar{r}a) \tag{8}$$

in such a way that $\gcd(\mathsf{n}(a), \mathsf{n}(c)) = 1$, and complementing with appropriate b, d via the KLPT algorithm.

Theorem 3.12. *Let $g \in \mathrm{M}_2(\mathcal{O}_0)$ be a reduced matrix as in Definition 3.11. There exists a (heuristic) polynomial-time algorithm that finds a solution to* (8) *with $\mathsf{n}(a)$ and $\mathsf{n}(c)$ coprime, provided that $\ell^{e_2} \in \Theta(p^{6.5+\varepsilon})$.*

Proof. We make the following restrictions: we take a of the form $a_1 + a_2 i \in \mathbb{Z}[i]$ and we take c of the form $c_1\bar{r}j + c_2\bar{r}k \in \bar{r}j\mathbb{Z}[i]$. Since $\mathsf{tr}(\bar{c}\bar{r}a)$ is zero for every such choice of a, c, equation (8) simplifies to $\ell^{e_2} = s\,\mathsf{n}(a) + t\,\mathsf{n}(c)$. We then solve the quadratic equation

$$t\,\mathsf{n}(c) = tp\,\mathsf{n}(r)(c_1^2 + c_2^2) \equiv \ell^{e_2} \bmod s.$$

Since s is an odd prime and $s \nmid t, p, \mathsf{n}(r), \ell$, this provides us with an irreducible conic equation over $\mathbb{F}_s$ which always has a solution: this gives us c, with $c_1, c_2 \in \{0, \ldots, s-1\}$. Now we have that $\ell^{e_2} - t\,\mathsf{n}(c)$ is divisible by s, so we are left with the equation

$$\frac{\ell^{e_2} - t\,\mathsf{n}(c)}{s} = \mathsf{n}(a). \tag{9}$$

Since $a \in \mathbb{Z}[i]$ this can be solved using Cornacchia's algorithm, provided we know the factorization of $\ell^{e_2} - t\,\mathsf{n}(c)$. Thus we iterate until (9) has a solution and we can factor $\ell^{e_2} - t\,\mathsf{n}(c)$ efficiently. Here one expects a polylogarithmic number of iterations. The reason for the size constraints on ℓ^{e_2} is that one needs $\ell^{e_2} - t\,\mathsf{n}(c)$ to be positive, as otherwise it cannot be the sum of two squares. From

$$t\,\mathsf{n}(c) = t \cdot p \cdot \mathsf{n}(r) \cdot (c_1^2 + c_2^2) \in O\left(\frac{p^{3+\varepsilon}}{s} \cdot p \cdot s^2 p \cdot 2s^2\right) \subset O(p^{6.5+\varepsilon})$$

the bound follows. Note that (9) does not guarantee that $\gcd(\mathsf{n}(a), \mathsf{n}(c)) = 1$, so a number of retries, each time choosing different representants of $c_1, c_2 \bmod s$ (or choosing a genuinely different solution to the above quadratic equation over $\mathbb{F}_s$), may be needed. This does not affect the above asymptotic estimate. □

Remark 3.13. In particular, from (9) it is clear that c should be chosen such that $\ell \nmid \mathsf{n}(c)$, for otherwise $\ell \mid \gcd(\mathsf{n}(a), \mathsf{n}(c))$. This is the reason for the condition $\ell \nmid \mathsf{n}(r)$ in Definition 3.11. It is interesting to specialize this to our main case of interest $\ell = 2$: both $\mathsf{n}(r)$ and $c_1^2 + c_2^2$ should be odd. Then, assuming $e_0, e_2 \geq 2$, equation (9) implies that

$$-t\,\mathsf{n}(c) = -tp\,\mathsf{n}(r)(c_1^2 + c_2^2) \equiv s\,\mathsf{n}(a) \bmod 4 \quad \Rightarrow \quad -t^2 p(c_1^2 + c_2^2) \equiv \mathsf{n}(a) \bmod 4,$$

showing that $\mathsf{n}(a) \equiv -1 \cdot 3 \cdot 1 \equiv 1 \bmod 4$. This is a necessary condition for the Cornacchia-step to succeed.

We are now ready to prove our main result:

Proof of Theorem 3.1: The algorithm to find $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ when given

$$g_1 = \begin{pmatrix} s_1 & r_1 \\ \bar{r}_1 & t_1 \end{pmatrix}, \quad g_2 = \begin{pmatrix} s_2 & r_2 \\ \bar{r}_2 & t_2 \end{pmatrix} \quad \in \mathrm{Mat}(A_0)$$

is summarized in Algorithm 1. From the preceding discussions, it should be clear that all steps are heuristically polynomial-time. As for the output length, note that the matrices u_1, u_2 produced in Step 7 have reduced norm ℓ^{e_1} with $e_1 = 2e_0$, and for $i = 1, 2$ the upper-left entry of $u_i^* g_i u_i$ equals ℓ^{e_2}. Thus, from (4) we find that γ has reduced degree

$$\ell^e = \ell^{2(e_1+e_2)} = (\ell^{e_0})^4 \cdot (\ell^{e_2})^2 \in O(p^{12+\varepsilon} \cdot p^{13+\varepsilon})$$

in view of the KLPT bound and Theorem 3.12. □

Algorithm 1: KLPT^2: An algorithm to solve the quaternion ℓ-isogeny path problem in dimension 2

Input : $g_1, g_2 \in \mathrm{Mat}(A_0)$
Output: $\gamma \in \mathrm{M}_2(\mathcal{O})$ such that $\gamma^* g_2 \gamma = \ell^e g_1$ with $\ell^e \in O(p^{25+\varepsilon})$

1 **For** *i=1,2* **do**
2 Find a, c using lattice reduction such that $\gcd(\mathsf{n}(a), \mathsf{n}(c)) = 1$, and $s_i\,\mathsf{n}(a) + t_i\,\mathsf{n}(c) + \mathsf{tr}(\bar{c}\bar{r}_i a) < \sqrt{p}(\ln p)^{1/4}$ is prime (not $2, \ell$).
3 Find b, d using KLPT as described in Section 3.2 such that the reduced norm of $u := \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is ℓ^{e_0}.
4 Find α such that $g' = \left(\begin{smallmatrix} s' & r' \\ \bar{r}' & t' \end{smallmatrix}\right) := \left(\begin{smallmatrix} 1 & 0 \\ \bar{\alpha} & 1 \end{smallmatrix}\right) u^* g_i u \left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right)$ is reduced.
5 Find a', c' using lattice reduction such that $\gcd(\mathsf{n}(a'), n(c')) = 1$ and $s'\,\mathsf{n}(a') + t'\,\mathsf{n}(c') + \mathsf{tr}(\bar{c}'\bar{r}'a') = \ell^{e_2}$ using Theorem 3.12.
6 Find b', d' using KLPT as described in Section 3.2 such that the reduced norm of $u' = \left(\begin{smallmatrix} a' & b' \\ c' & d' \end{smallmatrix}\right)$ is ℓ^{e_0}.
7 Let $u_i = u\left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right) u'$.
8 Compute τ connecting $u_1^* g_1 u_1$ and $u_2^* g_2 u_2$ as in Lemma 3.2.
9 **Return** $\gamma := u_2 \tau u_1^{-1} \mathcal{N}(u_1)$ as in (4).

The algebraic pathfinding problem was studied here for $N = \ell^e$ similarly to the original KLPT algorithm. However, it is clear that both KLPT and Theorem 3.12 can be adjusted to any number that is big enough; see also [29]. Now by invoking powersmooth versions of KLPT and Theorem 3.12 we get a powersmooth degree isogeny, since the product of powersmooth numbers is still powersmooth. This implies the following version of Theorem 3.1:

Theorem 3.14. *There exists a (heuristic) polynomial-time algorithm which upon input $g_1, g_2 \in \mathrm{Mat}(A_0)$ and a smoothness bound B returns $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that*

$$\gamma^* g_2 \gamma = N g_1$$

where $N \in O(p^{25+\varepsilon})$ is B-powersmooth.

3.5 Finding short isogenies

For certain applications one might be interested in a version of the algebraic isogeny problem $\gamma^* g_2 \gamma = N g_1$ where N is as small as possible. For elliptic curves this can be achieved via lattice reduction, but in higher dimension the set of polarized isogenies between two principally polarized abelian varieties no longer forms a lattice. Heuristically, one expects that $N \in O(p^{3/4+\varepsilon})$ should be feasible.[†] Here we briefly sketch a method which realizes $N \in O(p^{3+\varepsilon})$ under the assumption that one of the surfaces, say the codomain, concerns A_0 (equipped with the product polarization λ_0). In this case the matrix g_2 is simply the 2×2 identity matrix $\mathbb{I}_2 \in \mathrm{M}_2(\mathcal{O}_0)$. Observe the following:

Lemma 3.15. *Consider a matrix*

$$g = \begin{pmatrix} s & r \\ \bar{r} & t \end{pmatrix}, \quad s, t \in \mathbb{Z}_{>0},\ r \in \mathcal{O}_0$$

and assume that s is a prime congruent to 1 mod 4 and that $\det(g) = st - \mathsf{n}(r)$ is a square. Then there exists $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\gamma^* \gamma = s^2 g$.

Proof. We look first look for an upper-triangular $\gamma_0 = \left(\begin{smallmatrix} a & b \\ 0 & d \end{smallmatrix}\right) \in \mathrm{M}_2(B_{p,\infty})$ such that $\gamma_0^* \gamma_0 = g$. We get the following equations:

$$\mathsf{n}(b) + \mathsf{n}(d) = t,\ \bar{a}b = r,\ \ \mathsf{n}(a) = s$$

We solve $\mathsf{n}(a) = s$ using Cornacchia's algorithm and then choose $b = ra/s$. Now what remains is to choose d such that $\mathsf{n}(b) + \mathsf{n}(d) = t$. One finds that

$$\mathsf{n}(d) = t - \mathsf{n}(b) = t - \frac{\mathsf{n}(r)}{s} = \frac{\det(g)}{s} = \frac{A^2}{s}$$

for some $A \in \mathbb{Z}$. Thus we can choose $d = Aa/s$. Now $\gamma := s\gamma_0 \in \mathrm{M}_2(\mathcal{O}_0)$ satisfies $\gamma^* \gamma = s^2 g$. □

Now we can recycle the work done in the previous sections. As explained in Section 3.3, using lattice reduction with respect to the quadratic form $K((a, c))$ we can find $u = \left(\begin{smallmatrix} a & * \\ c & * \end{smallmatrix}\right)$ such that the top-left corner s of $u^* g u$ is a prime of size $O(p^{1/2+\varepsilon})$, such that $\mathsf{n}(a)$ and $\mathsf{n}(c)$ are coprime, and such that $s \equiv 1 \bmod 4$ (the latter congruence was not included in Section 3.3, but since this is a very mild assumption, it is not expected to cause a noticeable increase in size). Along the lines of Section 3.2, we can then complete $u = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ such that $\mathcal{N}(u) = \det(u^* g u)$ is a square, where we claim that this square can be chosen of order $O(p^{1+\varepsilon})$. Indeed, using lattice reduction in the ideal I from Proposition 3.7 it is expected that we can find an element $\omega \in I$ of norm $\mathsf{n}(c)q$ where $q \in O(p^{1/2+\varepsilon})$ is a prime

[†]For any prime ℓ there are about ℓ^3 emanating polarized (ℓ, ℓ)-isogenies [8, Lemma 2]. This gives about $B^{4+\varepsilon}$ emanating isogenies of reduced degree at most B, while the total number of principally polarized superspecial abelian surfaces is about $p^3/2880$, see [6]. Thus we can heuristically expect that $B \in O(p^{3/4+\varepsilon})$ should suffice.

congruent to 1 mod 4. Using Cornacchia's algorithm we can find $\alpha \in \mathbb{Z}[i] \subset \mathcal{O}_0$ such that $\mathsf{n}(\alpha) = q$. Then also $\omega\alpha \in I$ and it has norm $\mathsf{n}(c)q^2$. An application of Proposition 3.7 then yields the claim. Putting this all together, we get the following result:

Theorem 3.16. *Let $g \in \mathrm{Mat}(A_0)$. There exists a (heuristic) polynomial-time algorithm that finds $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\gamma^*\gamma = Ng$ and $N \in O(p^{3+\varepsilon})$.*

Proof. By Lemma 3.3 we can take $N = s^2\,\mathcal{N}(u)^2 \in O(p^{1+\varepsilon} \cdot p^{2+\varepsilon})$. □

4 Translating between matrices and isogenies

The main applications of the standard KLPT algorithm go hand in hand with efficient methods for converting left (non-zero) ideals of $\mathcal{O}_0$ into isogenies emanating from E_0 and vice versa. Likewise, in order to put KLPT^2 to practical use, we need methods for translating appropriately chosen 2×2 matrices with entries in $\mathcal{O}_0$ to polarized isogenies emerging from A_0 and conversely.

The analogy with the elliptic curve case becomes more apparent when noting that $\mathrm{M}_2(\mathcal{O}_0)$ is a principal ideal ring. Consequently, we have a natural identification of left ideals $I \subset \mathrm{M}_2(\mathcal{O}_0)$ with their generating matrices $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$, up to left-multiplication with elements of $\mathrm{GL}_2(\mathcal{O}_0)$. On a high level, the known approaches for translating between ideals and isogenies in dimension one carry over to dimension two. But in the case of isogeny-to-ideal conversion there is an important caveat: the ideal returned by the standard isogeny-to-ideal approaches is described in terms of multiple generators, and extracting a single generating matrix from this description is not a trivial task. Indeed, a large part of Chu's thesis [16, Chapter 2] is devoted to the design of a sub-exponential time algorithm for solving this instance of the principal ideal problem (PIP).

In this section, we describe some first routines for converting matrices to isogenies and vice versa; our main result is presented in Section 4.3, where we show how to by-pass the PIP for chains of $(2,2)$-isogenies, which for applications is the main case of interest. Then, in Section 5, we will enhance these basic routines through the use of KLPT^2.

4.1 Matrices to polarized isogenies from A_0

This is stated in [16, Section A.2] as a "required routine", but no details are given, even though the method is not too surprising. The input is a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ of reduced norm N^2, where $N = N_1N_2\cdots N_r$ is assumed powersmooth, i.e., the factors N_i are pairwise coprime and bounded by B for some constant $B = \mathrm{poly}(\log p)$. In view of Remark 2.5, we also assume that the kernel of γ, when identified with an endomorphism of A_0, is a maximal isotropic subgroup of $A_0[N]$ with respect to the N-Weil pairing for the product polarization λ_0. (If γ fails to meet this condition, then our method will detect this along the way.) The desired output is a chain of polarized isogenies

$$A_0 \xrightarrow{\varphi_1} A_1 \xrightarrow{\varphi_2} A_2 \xrightarrow{\varphi_3} \dots \xrightarrow{\varphi_r} A_r \tag{10}$$

of respective reduced degrees N_i, such that $\ker(\gamma) = \ker(\varphi_r \circ \cdots \circ \varphi_2 \circ \varphi_1)$, where each A_i is either a product $E_1 \times E_2$ of two elliptic curves equipped with the product polarization, or the Jacobian $\mathrm{Jac}(C)$ of a curve of genus 2 equipped with the canonical polarization.

The method starts off by computing a set of generators

$$S_i = \{(U_{ij}, V_{ij})\}_j, \qquad U_{ij}, V_{ij} \in E_0$$

of $(\ker\gamma)[N_i]$ for each $i = 1, \ldots, r$.[†] This can be done by first picking a basis $P_i, Q_i \in E_0[N_i]$. Such points can be found over a field extension of degree at most B^2. Then

$$(P_i, 0),\ (0, P_i),\ (Q_i, 0),\ (0, Q_i) \tag{11}$$

is a basis of $A_0[N_i]$, and the requested generators can be found by expressing that $x(P_i,0)+y(0,P_i)+z(Q_i,0)+w(0,Q_i) = (xP_i+zQ_i, yP_i+wQ_i)$ is annihilated by γ and solving a system of four homogeneous linear equations in the unknowns $x, y, z, w \in \mathbb{Z}/N_i\mathbb{Z}$. Explicitly writing down these equations involves discrete logarithm computations in groups of size N_i, so it is actually simpler to evaluate the adjoint isogeny[‡]

$$\tilde{\gamma} = N\gamma^{-1} \in \mathrm{M}_2(\mathcal{O}_0)$$

in the four points (11): their images generate $(\ker\gamma)[N_i]$.

Remark 4.1. As a sanity check, one can verify that these generators span a group with N_i^2 elements and that

$$e_{N_i}(U_{ij_1}, U_{ij_2}) \cdot e_{N_i}(V_{ij_1}, V_{ij_2}) = 1$$

for each pair $(U_{ij_1}, V_{ij_1}), (U_{ij_2}, V_{ij_2}) \in S_i$; this checks that the group is maximal isotropic with respect to the N_i-Weil pairing for the product polarization on A_0.

After gathering this data for $i = 1, \ldots, r$, we first compute a polarized isogeny

$$\varphi_1 : A_0 \to A_1$$

with kernel $(\ker\gamma)[N_1]$. For this step, various methods are available. E.g., if $N_1 = 2^e$ for some small exponent e, then this can be done through an e-fold application of the classical formulae due to Richelot [49]; the occasional gluing and splitting steps can be handled using [33]. For the general case, we refer to [18]. Even though the elements of $(\ker\gamma)[N_i]$ may live over an extension field only, the isogeny φ_1 itself is $\mathbb{F}_{p^2}$-rational. We then push the generators of $(\ker\gamma)[N_i]$ for $i = 2, \ldots, r$ through this isogeny and repeat, starting from A_1. Eventually we arrive at A_r in polynomial time, as wanted. The method is summarized in Algorithm 2.

[†]If N_i is prime then it is possible to use 2 generators, but in general one may need 3 generators (or even 4 generators in case γ factors through scalar multiplication).

[‡]Note: $N\gamma^{-1} \in \mathrm{M}_2(\mathcal{O}_0)$ relies on γ being polarized and is a stronger statement than Lemma 2.12 (which says that $N^2\gamma^{-1} \in \mathrm{M}_2(\mathcal{O}_0)$).

Algorithm 2: MatrixToIsogeny: powersmooth degree

Input : $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ with $\mathrm{degrd}(\gamma) = N_1 \cdots N_r$ powersmooth
Output: polarized isogenies $\varphi_r \circ \cdots \circ \varphi_1 : A_0 \to A_r$ with $\deg \varphi_i = N_i$

1 $\tilde{\gamma} \leftarrow N\gamma^{-1}$, $\varphi_0 = \mathrm{id}$.
2 **For** $i = 1, \ldots, r$ **do**
3 $\quad P_i, Q_i \leftarrow$ basis of $E_0[N_i]$.
4 $\quad S_i \leftarrow \tilde{\gamma}\left(\{(P_i, 0), (0, P_i), (Q_i, 0), (0, Q_i)\}\right)$.
5 $\quad$ `// Generators of` $(\ker \gamma)[N_i]$`.`
6 **For** $i = 1, \ldots, r$ **do**
7 $\quad S_i \leftarrow (\varphi_{i-1} \circ \cdots \circ \varphi_0)(S_i)$.
8 $\quad \varphi_i \leftarrow$ isogeny $A_{i-1} \to A_i$ with kernel $\langle S_i \rangle$.
9 **Return** $\varphi_r \circ \cdots \circ \varphi_1 : A_0 \to A_r$.

4.2 Polarized isogenies from A_0 to matrices

Conversely, given a chain of polarized isogenies emanating from A_0 as in (10), where the degrees $N_i = \deg \varphi_i$ are pairwise coprime and bounded by B, here the goal is to produce a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\ker(\gamma) = \ker(\varphi_r \cdots \varphi_2 \circ \varphi_1)$. Such a matrix is uniquely determined up to left-multiplication with an element of $\mathrm{GL}_2(\mathcal{O}_0)$ and will automatically satisfy

$$\gamma^* g_r \gamma = N \cdot \mathbb{I}_2$$

with $N = N_1 N_2 \cdots N_r$, for some representant $g_r \in \mathrm{Mat}(A_0)$ of the class in $\mathrm{Mat}^0(A_0)$ corresponding to the principally polarized abelian surface A_r. Note that g_r can then be computed as $N\gamma^{*-1}\gamma^{-1} = N(\gamma\gamma^*)^{-1}$.

At a high level, this conversion can be done as in Algorithm 3, which is just a slight variation on Chu's method from [16, Algorithm A.2.2]. Concerning Step 3, recall that $N_i \leq B$ implies that the elements of G_i are defined over an extension field of degree $O(B^2)$. In Step 6, it is possible to narrow the search space by taking into account the Weil pairing. This is done in [16, Algorithm A.2.2], but here we content ourselves with a naive search, since the pith of the method lies in Step 8 anyway; this is the aforementioned instance of the principal ideal problem, for which we rely on Chu's sub-exponential time subroutine, described in [16, Chapter 2].

4.3 Isogeny-to-matrix conversion for chains of $(2, 2)$-isogenies

We now discuss the main result of Section 4, namely an extremely efficient method for isogeny-to-matrix conversion for $(2, 2)$-isogenies. This is simply based on an explicit list of matrices in $\mathrm{M}_2(\mathcal{O}_0)$ whose kernels cover every possible subgroup of $A_0[2]$ that is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^2$. The number of such subgroups

Algorithm 3: IsogenyToMatrix: powersmooth degree

Input : chain $\varphi_r \circ \cdots \circ \varphi_1 : A_0 \to A_r$ of polarized isogenies with $N_i := \mathrm{degrd}(\varphi_i) \leq B$ pairwise coprime

Output: $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\ker(\gamma) = \ker(\varphi_r \circ \cdots \circ \varphi_1)$, $g_r \in \mathrm{Mat}(A_0)$ corresponding to A_r

$\gamma \leftarrow \mathbb{I}_2$.
For $i = 1, \ldots, r$ **do**
  $G_i \leftarrow (\tilde{\varphi}_{i-1} \circ \cdots \circ \tilde{\varphi}_2 \circ \tilde{\varphi}_1)(\ker \varphi_i) \subset A_0[N_i]$.
    `// Pulling back the kernel of` φ_i `to` A_0`.`
  $K_i \leftarrow \gamma(G_i)$. `// Pushing the kernel forward under` γ`.`
  Find matrix $\Gamma_i \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\ker(\Gamma_i) \cap A_0[N_i] = K_i$.
    `// Exhaustive search over` $\mathrm{M}_2(\mathcal{O}_0)/\mathrm{M}_2(\mathcal{O}_0)N_i$`.`
  $\gamma_i \leftarrow$ generator of left ideal $\mathrm{M}_2(\mathcal{O}_0)\Gamma_i + \mathrm{M}_2(\mathcal{O}_0)N_i$. `// PIP.`
  $\gamma \leftarrow \gamma_i\gamma$
Return γ, $N_1 \cdots N_r(\gamma\gamma^*)^{-1}$.

is given by the Gaussian binomial coefficient

$$\begin{bmatrix} 4 \\ 2 \end{bmatrix}_2 = 35.$$

The list can be found in Appendix A and was produced through a combination of naive search and symbolic verification (observe that, except for the matrices corresponding to the trivial groups $E_0[2] \times \{\infty\}$ and $\{\infty\} \times E_0[2]$, all matrix entries only take coordinates $-1, 0, 1$ with respect to the basis from Example 2.1).

While this may seem very specific, an iterated application allows for isogeny-to-matrix conversion in one of the main cases of interest, namely where the input isogeny $\varphi : A_0 \to A$ is a $(2^e, 2^e)$-isogeny for some $e \geq 1$, i.e., with kernel

$$K \cong \frac{\mathbb{Z}}{2^e\mathbb{Z}} \times \frac{\mathbb{Z}}{2^e\mathbb{Z}},$$

and we moreover assume that K is generated by points defined over $\mathbb{F}_{p^2}$ (or a small-degree extension thereof). The method is totally straightforward and can be found in Algorithm 4. In contrast with the previous section, we do not assume that 2^e is polynomially bounded, i.e., we drop the powersmoothness assumption.[†]

[†]Of course, there is still an implicit bound coming from the rationality assumption, i.e., of the kind $2^e \mid p + 1$. In Section 5.2 we will use the KLPT^2 algorithm to get rid of this assumption.

Algorithm 4: IsogenyToMatrix22

Input : subgroup $K \cong (\mathbb{Z}/2^e\mathbb{Z})^2$ of A_0 generated by points over $\mathbb{F}_{p^2}$
Output: $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\ker(\gamma) = K$

1 $\gamma \leftarrow \mathbb{I}_2$, $K_1 \leftarrow K$.
2 **For** $i = 1, \ldots, e$ **do**
3 $\quad G_i \leftarrow 2^{e-i}K_i$.
4 $\quad \gamma_i \leftarrow$ matrix with kernel G_i. `// Look up in Appendix A.`
5 $\quad \gamma \leftarrow \gamma_i\gamma$, $K_{i+1} \leftarrow \gamma_i(K_i)$.
6 **Return** γ.

The method works equally well for more general polarized isogenies $\varphi : A_0 \to A$ of reduced degree 2^e, say with kernel

$$K \cong \frac{\mathbb{Z}}{2^e\mathbb{Z}} \times \frac{\mathbb{Z}}{2^{e-f}\mathbb{Z}} \times \frac{\mathbb{Z}}{2^f\mathbb{Z}}$$

for some $f \in \{1 \ldots, \lfloor e/2 \rfloor\}$, as long as this kernel is generated by rational points. The main caveat lies in Step 3, where one should be more careful: indeed, in this case $2^{e-1}K \ncong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. A clean workaround, which serves as a warm-up for Section 5.4, is to define the subgroup

$$K' = \langle 2^{e-f}P, R \rangle \cong \frac{\mathbb{Z}}{2^f\mathbb{Z}} \times \frac{\mathbb{Z}}{2^f\mathbb{Z}}$$

with $P \in K$ any point of order 2^e and $R \in K$ any point of order 2^f that is not halvable in K. Since $e_{2^f,\lambda_0}(2^{e-f}P, R) = e_{2^e,\lambda_0}(P, R) = 1$, this concerns a maximal isotropic subgroup of $A_0[2^f]$. We can now run Algorithm 4 on input K', returning a matrix γ', and then rerun the algorithm on input

$$\gamma'(K') \cong K/K' \cong \frac{\mathbb{Z}}{2^{e-f}\mathbb{Z}} \times \frac{\mathbb{Z}}{2^{e-f}\mathbb{Z}},$$

after initializing $\gamma \leftarrow \gamma'$ rather than $\gamma \leftarrow \mathbb{I}_2$ in Step 1.

Remark 4.2. We did not generate similar lists for (ℓ, ℓ)-isogenies with $\ell > 2$, but the simple shape of the matrices in Appendix A makes it reasonable to assume that this should be doable for the first few primes. Observe that the search can be sped up (and the lists can be shortened) by considering kernels up to base change by matrices from $\mathrm{GL}_2(\mathcal{O}_0)$; for $\ell = 2$ there was no need for implementing this. In the worst case, such lists can be generated using Chu's algorithm from [16]. By doing this for all ℓ up to the powersmoothness bound B, this approach shifts the sub-exponential portion of the isogeny-to-matrix conversion algorithm (Steps 6–8 in Algorithm 3) to a one-time precomputation step.

5 Applications of KLPT^2

We are ready to discuss a number of applications of the KLPT^2 algorithm.

5.1 Constructive IKO correspondence

For elliptic curves, the *constructive Deuring correspondence* asks to solve the following problem: upon input of a maximal order $\mathcal{O} \subset B_{p,\infty}$, return a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ such that $\mathrm{End}(E) \cong \mathcal{O}$. The KLPT algorithm can be turned into a heuristic polynomial-time algorithm for solving this problem. At a high level, the method works as follows. One starts from an elliptic curve $E_0/\mathbb{F}_{p^2}$ having a known, special extremal endomorphism ring $\mathrm{End}(E_0) \cong \mathcal{O}_0$. Using the KLPT algorithm, one computes a left ideal $I \subset \mathcal{O}_0$ of powersmooth norm N connecting $\mathcal{O}_0$ and $\mathcal{O}$. This ideal can then be converted into an isogeny emerging from E_0 using Algorithm 2. The codomain of this isogeny is a valid output for the constructive Deuring correspondence.

For *unpolarized* superspecial abelian surfaces, the direct analog of the constructive Deuring correspondence is void: all such surfaces are pairwise isomorphic and therefore share the same endomorphism ring, namely $\mathrm{M}_2(\mathcal{O}_0)$. However, in the principally polarized case, the endomorphism ring comes equipped with an extra datum: the Rosati involution, which as explained in Remark 2.8 is completely encoded in the matrix $g \in \mathrm{Mat}(A_0)$ corresponding to λ. Therefore, a more meaningful counterpart of the constructive Deuring correspondence reads:

Theorem 5.1 (constructive IKO correspondence). *There exists a (heuristic) polynomial-time algorithm which upon input $g \in \mathrm{Mat}(A_0)$, either finds two elliptic curves E_1, E_2 or finds a genus-2 curve C such that for*

$$(A, \lambda) = (E_1 \times E_2, \textit{product polarization}), \quad \textit{resp.}$$
$$(A, \lambda) = (\mathrm{Jac}(C), \textit{canonical polarization}),$$

we have $(A, \lambda) \cong (A_0, \mu^{-1}(g))$, with μ the map from Theorem 2.7.

Proof. Using our pathfinding algorithm from Theorem 3.14 we can find $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that

$$\gamma^* g \gamma = N\mathbb{I}_2,$$

with N powersmooth. To produce the desired output, one then simply converts γ into a polarized isogeny emanating from A_0 using Algorithm 2. If the codomain of this polarized isogeny is a product $E_1 \times E_2$, we output E_1, E_2; when landing on a Jacobian $\mathrm{Jac}(C)$, output C. □

5.2 Relaxing powersmoothness assumptions when translating between matrices and isogenies

(i) Matrices to isogenies from A_0 in arbitrary degree. Let us be given a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ as in Section 4.1, but we drop the assumption that N is

powersmooth. We claim that, using KLPT^2, we can nevertheless convert γ into a polarized isogeny φ emanating from A_0. This mimicks well-known techniques from the elliptic curve case [26, 53]. First, recall that a matrix $g \in \mathrm{Mat}(A_0)$ representing the codomain can be computed as

$$g = N(\gamma\gamma^*)^{-1}.$$

Then, using Theorem 3.14, we can find a matrix γ' and a powersmooth integer N' such that

$$\gamma'^* g \gamma' = N' \cdot \mathbb{I}_2,$$

and we know that γ, γ' correspond to polarized isogenies φ, φ' with the same codomain:

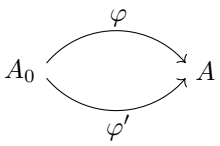

where $\mathrm{degrd}(\varphi) = N$ and $\mathrm{degrd}(\varphi') = N'$. We can compute φ' as a composition of small-degree isogenies using Algorithm 2, which also reveals A. We have $N'\varphi = \varphi'\tilde{\varphi}'\varphi$ where we note that

$$\tilde{\varphi}'\varphi = \lambda_0^{-1}\hat{\varphi}'\lambda\varphi = \lambda_0^{-1}\hat{\varphi}'\lambda_0\,\lambda_0^{-1}\lambda\,\varphi \ \in \mathrm{End}(A_0)$$

can be identified with $\gamma'^* g \gamma \in \mathrm{M}_2(\mathcal{O}_0)$. Thus we can evaluate

$$\varphi(P) = \frac{1}{N'}\varphi'(\gamma'^* g \gamma P)$$

on any input point P whose order is coprime with N'. This is enough for considering φ as being known, e.g., in view of [46, 47].

Remark 5.2 (matrices to isogenies from A_0 in smooth degree). If N is smooth (but not powersmooth, so that Algorithm 2 may not be applicable) then the above "evaluation representation" of φ may not be the preferred format. Rather, one may want an explicit decomposition $\varphi = \varphi_r \circ \cdots \circ \varphi_1$ into isogenies of small degree. A polynomial-time conversion between these formats is possible through a repeated use (of the two-dimensional analogue) of [47, Corollary 6.8], but this is not practical. Unfortunately, more direct methods such as [53, Algorithm 4] come with PIP-style challenges. However, in our main case of interest $N = 2^e$ these issues can be by-passed using our list of matrices from Appendix A, leading to the method described in Algorithm 5.

(ii) Isogenies from A_0 to matrices in smooth degree. The KLPT^2 algorithm can also be used to convert polarized isogenies from A_0 into matrices when the degree is smooth, rather than powersmooth. For this we recycle a trick due

Algorithm 5: MatrixToIsogeny22

Input : $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ with $\mathrm{degrd}(\gamma) = 2^e$

Output: polarized isogenies $\varphi_e \circ \cdots \circ \varphi_1 : A_0 \to A_e$ with $\mathrm{degrd}\,\varphi_i = 2$

1 **For** $i = 1, \ldots, e-1$ **do**
2 $\quad G_i \leftarrow (\ker \gamma)[2]$.
3 $\quad \gamma_i \leftarrow$ matrix with kernel G_i. `// Look up in Appendix A.`
4 $\quad \gamma \leftarrow \gamma\gamma_i^{-1}$.
5 $\gamma_e \leftarrow \gamma$, $\gamma \leftarrow \mathbb{I}_2$.
6 `// Input` γ `decomposed as` $\gamma_e \cdots \gamma_1$`; then reinitialize` γ.
7 **For** $i = 1, \ldots, e$ **do**
8 $\quad \gamma \leftarrow \gamma_i\gamma$.
9 $\quad g_i \leftarrow 2^i(\gamma\gamma^*)^{-1}$. `// Codomain matrix of` φ_i.
10 $\quad$ Find $\gamma' \in \mathrm{M}_2(\mathcal{O}_0)$ and odd powersmooth N' such that $\gamma'^* g_i \gamma' = N' \cdot \mathbb{I}_2$. `// Mild strengthening of Theorem 3.14.`
11 $\quad$ Using Algorithm 2, convert γ' to polarized isogeny $\varphi' : A_0 \to A_i$.
12 $\quad G_i \leftarrow \ker(\tilde{\gamma}\gamma')[2]$. `// Note` $\tilde{\gamma} = 2^i\gamma^{-1}$.
13 $\quad \varphi_i \leftarrow$ adjoint of isogeny $A_i \to A_{i-1}$ with kernel $\varphi'(G_i)$.
14 **Return** $\varphi_e \circ \cdots \circ \varphi_1 : A_0 \to A_e$.

to Eisenträger, Hallgren, Lauter, Morrison and Petit [26, Algorithm 9]; see also Wesolowski [53, Algorithm 3]. The method is detailed in Algorithm 6, where we note that Steps 3–9 are trivial at iteration $i = 1$. The overall runtime remains sub-exponential, in view of Chu's subroutine for the principal ideal problem (PIP) invoked in Step 12.

However, thanks to our explicit list of matrices from Appendix A, for chains of $(2,2)$-isogenies this can be done in polynomial time, even when the kernel is not generated by rational points, as was assumed in Section 4.3. This is done by factoring $N = 2^e = N_1 N_2 \cdots N_r$ into smaller powers and running a slightly modified version of Algorithm 6, where each iteration of Steps 10–12 is replaced with a short series of look-ups, of the kind described in the for-loop in Algorithm 4.

Remark 5.3. As a continuation of Remark 4.2, let us note that for $N = \ell^e$ with $\ell > 2$ the sub-exponential part of Algorithm 6 can again be handled using a one-time precomputation.

5.3 Translating between matrices and isogenies from other starting surfaces

(i) Matrices to polarized isogenies. Next, let us be given a matrix $g_1 \in \mathrm{Mat}(A_0)$ and a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ of reduced norm N^2 (for arbitrary N), with the promise that γ defines a polarized isogeny emanating from (A_0, λ_1), where

Algorithm 6: IsogenyToMatrix: smooth degree

Input : chain $\varphi_r \circ \cdots \circ \varphi_1 : A_0 \to A_r$ of polarized isogenies with $N_i := \mathrm{degrd}(\varphi_i) \leq B$

Output: $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\ker(\gamma) = \ker(\varphi_r \circ \cdots \circ \varphi_1)$, $g_r \in \mathrm{Mat}(A_0)$ corresponding to A_r

$\gamma \leftarrow \mathbb{I}_2$.
For $i = 1, \ldots, r$ **do**
  $g_i \leftarrow N_1 \cdots N_{i-1}(\gamma\gamma^*)^{-1}$. `// Codomain matrix of` φ_{i-1}.
  Find $\gamma' \in \mathrm{M}_2(\mathcal{O}_0)$ and powersmooth N' with $\gcd(N, N_i) = 1$ and $\gamma'^* g_i \gamma' = N' \cdot \mathbb{I}_2$. `// Mild strengthening of Theorem 3.14.`
  Using Algorithm 2, convert γ' to polarized isogeny $\varphi' : A_0 \to A_{i-1}$.
    `// Domain of` φ_i.
  $G_i \leftarrow \tilde{\varphi}'(\ker \varphi_i) \subset A_0[N_i]$.
    `// Pulling back the kernel of` φ_i `to` A_0.
  $K_i \leftarrow \gamma'(G_i)$. `// Pushing the kernel forward under` γ'.
  Find matrix $\Gamma_i \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\ker(\Gamma_i) \cap A_0[N_i] = K_i$.
    `// Exhaustive search over` $\mathrm{M}_2(\mathcal{O}_0)/\mathrm{M}_2(\mathcal{O}_0)N_i$.
  $\gamma_i \leftarrow$ generator of left ideal $\mathrm{M}_2(\mathcal{O}_0)\Gamma_i + \mathrm{M}_2(\mathcal{O}_0)N_i$. `// PIP.`
  $\gamma \leftarrow \gamma_i \gamma$
Return γ, $N_1 \cdots N_r(\gamma\gamma^*)^{-1}$.

$\lambda_1 = \mu^{-1}(g_1)$ is the principal polarization corresponding to g_1. Our goal is to tackle the following enhanced version of the constructive IKO correspondence: return the top row in a commutative diagram of the form

$$\begin{array}{ccc} \mathrm{Jac}(C_1) \text{ or } E_{11} \times E_{12} & \xrightarrow{\varphi} & \mathrm{Jac}(C_2) \text{ or } E_{21} \times E_{22} \\ \cong\downarrow & & \downarrow\cong \\ (A_0, \lambda_1) & \xrightarrow{\gamma} & (A_0, \lambda_2) \end{array}$$

That is, for $i = 1, 2$ one should return the underlying genus-2 curve C_i or elliptic curves E_{i1}, E_{i2}, along with an efficient representation of φ. This can be done as follows. If N is powersmooth, then using a mild strengthening of Theorem 3.14 we can compute a matrix $\kappa \in \mathrm{M}_2(\mathcal{O}_0)$ and a powersmooth integer K such that $\gcd(K, N) = 1$ and $\kappa^* g_1 \kappa = K \cdot \mathbb{I}_2$. This implies that

$$(\gamma\kappa)^* g_2 (\gamma\kappa) = NK \cdot \mathbb{I}_2$$

with $g_2 = \mu(\lambda_2) = N\gamma^{*-1} g_1 \gamma^{-1}$. We can then run Algorithm 2 on input $\gamma\kappa$, first processing the factors of K, to end up with an isogeny that naturally factors as

$$(A_0, \lambda_0) \longrightarrow \begin{matrix}\mathrm{Jac}(C_1) \text{ or} \\ E_{11} \times E_{12}\end{matrix} \xrightarrow{\varphi} \begin{matrix}\mathrm{Jac}(C_2) \text{ or} \\ E_{21} \times E_{22}\end{matrix} :$$

hence the desired output. If N is not powersmooth then we first apply Theorem 3.14 to replace γ with a matrix γ' of powersmooth reduced norm N'^2 (i.e., satisfying $\gamma'^* g_2 \gamma' = N' \cdot g_1$) and proceed as in Section 5.2.

(ii) Polarized isogenies to matrices. We can easily extend the foregoing methods for isogeny-to-matrix conversion from (A_0, λ_0) to any principally polarized starting surface (A, λ), say given as a Jacobian $\mathrm{Jac}(C)$ or a product of elliptic curves $E_1 \times E_2$, as soon as a corresponding matrix $g \in \mathrm{Mat}(A_0)$ is known. As explained in Section 4.2, the output of isogeny-to-matrix conversion is determined up to left-multiplication with an element of $\mathrm{GL}_2(\mathcal{O}_0)$ only. Here, there is an extra ambiguity: taking a different representant of g in $\mathrm{Mat}^0(A_0)$ amounts to right-multiplication with an element of $\mathrm{GL}_2(\mathcal{O}_0)$. Apart from this subtlety, the problem easily reduces to the case $(A, \lambda) = (A_0, \lambda_0)$ and therefore the runtimes are alike:

- sub-exponential time if the degree is smooth, but
- polynomial time if this degree is a power of 2.

Indeed, by means of Theorem 3.1 we can find a matrix γ' with reduced degree 2^e connecting $\mathbb{I}_2$ and g. Using Algorithm 5 this matrix can be converted into a polarized isogeny $\varphi' : (A_0, \lambda_0) \to (A, \lambda)$, represented as a chain of $(2,2)$-isogenies. Now if φ is a polarized isogeny emanating from (A, λ), then $\varphi \circ \varphi'$ is a polarized isogeny emanating from (A, λ_0), and if γ is a matrix corresponding to $\varphi \circ \varphi'$, then $\gamma\gamma'^{-1}$ is a matrix corresponding to φ.

5.4 Attacks on CGL hash functions

We now arrive at our main cryptographic application: finding collisions for two-dimensional variants of the Charles–Goren–Lauter (CGL) hash function [12], in the case of an untrusted set-up.

CGL hash functions in dimension two. In 2018, Takashima [50] proposed the first such variant, using random non-backtracking walks in the $(2,2)$-isogeny graph of superspecial principally polarized abelian surfaces. It was observed by Flynn and Ti [28] that such hash functions admit trivial collisions, coming from the fact that every $(4,2,2)$-isogeny admits three different decompositions into two $(2,2)$-isogenies. Therefore, starting with [10], all subsequent proposals restrict to "good" chains of $(2,2)$-isogenies, i.e., composing to a $(2^e, 2^e)$-isogeny

for some $e \geq 1$. This means that the kernel of every outgoing $(2,2)$-isogeny trivially intersects the kernel of the dual of the previous, incoming $(2,2)$-isogeny.[†]

Let us briefly detail how CGL hash functions in dimension two are currently constructed, generalizing from $(2,2)$-isogenies to (ℓ,ℓ)-isogenies for any small prime ℓ. During set-up, an initial node in the graph is chosen, corresponding to some superspecial principally polarized abelian surface A_1, as well as ℓ^3 "allowed" outgoing edges, corresponding to a subset of the set of $(\ell^2+1)(\ell+1)$ outgoing polarized (ℓ,ℓ)-isogenies from this initial node. At each node, the outgoing edges are sorted in some deterministic way; e.g. by comparing the invariants of all the neighbor nodes. The input message mess is mapped deterministically to $(m_1, m_2, \ldots, m_k) \in \{0,1,\ldots,\ell^3-1\}^*$, with some padding if necessary. To hash, one of the ℓ^3 allowed edges is chosen according to the value of m_1, and we compute the corresponding neighbor node, yielding a new principally polarized abelian surface A_2. Using the value m_2, we choose one of the ℓ^3 outgoing edges corresponding to an (ℓ,ℓ)-isogeny whose kernel trivially intersects the kernel of the dual of the previous (ℓ,ℓ)-isogeny. This results in a node corresponding to a surface A_3 and we repeat this process until we have landed on a node corresponding to a surface A_{k+1}. We deterministically map suitable invariants of A_{k+1} to $\{0,1\}^n$, where $n \approx 3\log p$, and use this as the output of our hash function.

KLPT^2 produces "bad" chains. In dimension one, the KLPT algorithm can be used to compute second pre-images for the CGL hash function as soon as the endomorphism ring of the starting curve is known [26]. In dimension two, a very similar reasoning applies as soon as a matrix $g_1 \in \mathrm{Mat}(A_0)$ corresponding to the initial node is known. But the conclusion is more subtle because, with overwhelming probability, our KLPT^2 algorithm does not return (ℓ^e,ℓ^e)-isogenies, in view of Corollary 5.5 below:

Lemma 5.4. *Let $a, c \in \mathcal{O}_0$ be non-zero elements such that $\gcd(\mathsf{n}(a),\mathsf{n}(c)) = 1$ and assume that $\ell \nmid \mathsf{n}(a), \mathsf{n}(c)$. Let $b, d \in \mathcal{O}_0$ be such that*

$$\mathcal{N}(u) = \ell^{e_0}, \qquad u = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

computed by finding an element of norm $\mathsf{n}(c)\ell^{e_0}$ in $\mathsf{n}(c)\mathcal{O}_0 + a\bar{c}\mathcal{O}_0$ via the KLPT algorithm, as explained in Section 3.2. Assume that the degree-ℓ^{e_0} component of this element is cyclic (this is generically expected). Then $\ker(u) \cong \mathbb{Z}/\ell^{e_0}\mathbb{Z}$.

Proof. As before, choose integers α, β such that $1 = \alpha\,\mathsf{n}(a) + \beta\,\mathsf{n}(c)$. Expressing that a point tuple $(P, P') \in A_0[\ell^{e_0}] = E_0[\ell^{e_0}]^2$ is contained in $\ker(u)$ and multiplying on the left with the row matrix $\begin{pmatrix}\alpha\bar{a} & \beta\bar{c}\end{pmatrix}$, we find

$$P = -(\alpha\bar{a}b + \beta\bar{c}d)P' \tag{12}$$

[†]Rather than merely not coinciding with it: this is how "non-backtracking" was understood in Takashima's proposal [50].

as a necessary condition. Next, multiplying with $\begin{pmatrix}1 & 0\end{pmatrix}$ and substituting P yields:

$$(-a(\alpha\bar{a}b + \beta\bar{c}d) + b)\, P' = 0 \quad \Rightarrow \quad (\mathsf{n}(c)b - a\bar{c}d)P' = 0, \tag{13}$$

where we have assumed $\gcd(\beta, \ell) = 1$; if $\gcd(\beta, \ell) > 1$ then $\gcd(\alpha, \ell) = 1$ and we can proceed similarly, instead multiplying with $\begin{pmatrix}0 & 1\end{pmatrix}$. Since $\mathsf{n}(\mathsf{n}(c)b - a\bar{c}d) = \mathsf{n}(c)\mathcal{N}(u) = \ell^{e_0}\,\mathsf{n}(c)$ and $\ell \nmid \mathsf{n}(c)$, we find from (12) and (13) that $\ker(u)$ is isomorphic to the kernel of the returning degree-ℓ^{e_0} component of the endomorphism $\mathsf{n}(c)b - a\bar{c}d$. By Remark 3.8, this is precisely the isogeny that is produced by our run of KLPT, which we assumed to be cyclic. □

Corollary 5.5. *Under the assumption from Lemma 5.4 that the KLPT algorithm produces cyclic isogenies, to describe the kernel of the matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ returned by Algorithm 1, one needs at least 3 generators.*

Proof. From (4) we know that $\gamma = u_2\tau u_1^{-1}\mathcal{N}(u_1)$, where we recall from Algorithm 1 that $u_1 = u\left(\begin{smallmatrix}1 & \alpha\\ 0 & 1\end{smallmatrix}\right)u'$ for matrices u, u' of reduced norm ℓ^{e_0}. The matrix u is computed as in the proof of Theorem 3.12. From (9) it follows that u satisfies the assumptions from Lemma 5.4, therefore its kernel is cyclic of size ℓ^{e_0}. But γ factors through $u^{-1}\mathcal{N}(u) = \ell^{e_0}u^{-1}$, and the kernel of this factor is isomorphic to $(\mathbb{Z}/\ell^{e_0}\mathbb{Z})^3$ because u and $\ell^{e_0}u^{-1}$ compose to multiplication-by-ℓ^{e_0}. □

Collision finding. While this complicates the construction of second preimages, it still lends itself to finding collisions, as we now discuss in detail. The first step is to run the KLPT^2 algorithm on input g_1, g_1, resulting in a matrix $\gamma \in \mathrm{M}_2(\mathcal{O}_0)$ defining a polarized endomorphism

$$(A_0, \lambda_1) \longrightarrow (A_0, \lambda_1), \quad A_1 \cong (A_0, \lambda_1),\ g_1 = \mu(\lambda_1)$$

of reduced degree ℓ^e, where we recall from (4) that $e = 2e_1 + 2e_2$ with $e_1 = 2e_0$. As we have just explained, we need at least 3 generators to describe the kernel of γ, but we can be more precise. Namely, in "typical" situations, we expect the following properties to hold; the notation below again follows the pseudo-code from Algorithm 1.

- In the proof of Corollary 5.5, also the reduction matrix u' meets the assumptions from Lemma 5.4, i.e., $\ell \nmid \mathsf{n}(a'), \mathsf{n}(c')$. As a consequence we have $\ker(u') \cong \mathbb{Z}/\ell^{e_0}\mathbb{Z}$.
- The cyclic kernel of $u\left(\begin{smallmatrix}1 & \alpha\\ 0 & 1\end{smallmatrix}\right)$ trivially intersects $\ker(\ell^{e_0}u'^{-1}) \cong (\mathbb{Z}/\ell^{e_0}\mathbb{Z})^3$. This implies that $\ker(u_1) \cong \mathbb{Z}/\ell^{2e_0}\mathbb{Z} = \mathbb{Z}/\ell^{e_1}\mathbb{Z}$.
- Likewise $\ker(u_2) \cong \mathbb{Z}/\ell^{e_1}\mathbb{Z}$.
- Upon writing
$$\tau = \begin{pmatrix}\ell^{e_2} & x\\ 0 & \ell^{e_2}\end{pmatrix}$$
for some $x \in \mathcal{O}_0$, we have $\ell \nmid \mathsf{n}(x)$, implying that $\ker(\tau) \cong (\mathbb{Z}/\ell^{2e_2}\mathbb{Z})^2$.

– Furthermore, $\ker(\tau)$ trivially intersects the cyclic kernel of u_1. This implies that
$$\ker(\tau\ell^{e_1}u_1^{-1}) \cong \frac{\mathbb{Z}}{\ell^{e_1+2e_2}\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell^{e_1+2e_2}\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell^{e_1}\mathbb{Z}}.$$
– Defining
$$\tilde{\tau} = \begin{pmatrix} \ell^{e_2} & -x \\ 0 & \ell^{e_2} \end{pmatrix}$$
so that $\tilde{\tau}\tau = \tau\tilde{\tau} = \ell^{2e_2}\mathbb{I}_2$,[†] the cyclic kernel of u_2 trivially intersects $\ker(\tilde{\tau}) \cong (\mathbb{Z}/\ell^{2e_2}\mathbb{Z})^2$, which finally implies
$$\ker(\gamma) = \ker(u_2\tau\ell^{e_1}u_1^{-1}) \cong \frac{\mathbb{Z}}{\ell^{2e_1+2e_2}\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell^{e_1+2e_2}\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell^{e_1}\mathbb{Z}}. \tag{14}$$

Recall that this concerns a maximal isotropic subgroup for the ℓ^e-Weil pairing with respect to λ_1. By "typical", we mean that we expect (14) to hold with a constant probability depending only on ℓ (rapidly converging to 1 as $\ell \to \infty$). This is motivated by Lemma B.1. In particular, heuristically, we should end up with the shape (14) after a constant number of reruns. We stress that, while this shape is convenient for expository purposes, the attack method below can be adapted to any other isomorphism type.

Then the idea for converting γ into a collision is inspired by the reasoning in Section 4.3, where it was argued that there exists a subgroup of $\ker(\gamma)$ determining a polarized (ℓ^{e_1}, ℓ^{e_1})-isogeny $\gamma_1 : (A_0, \lambda_1) \to (A_0, \lambda)$, through which γ factors, such that the remaining factor $\gamma_2 : (A_0, \lambda) \to (A_0, \lambda_1)$ is a polarized $(\ell^{e_1+2e_2}, \ell^{e_1+2e_2})$-isogeny. (In fact, in Section 4.3 this was only explained for $\ell = 2$ and the product polarization on A_0, but the argument carries over and will be revisited below, in any case.) Thus we have two "good" paths

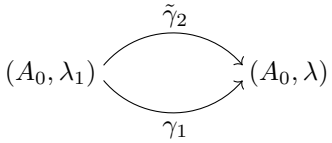

with the same codomain: this is the algebraic version of the desired collision. If we effectively succeed in finding γ_1, γ_2 then these matrices can be converted into two colliding isogenies emanating from A_1, by following the procedure described in Section 5.3.

In Section 4.3 we had an explicit description of the subgroup $\ker(\gamma_1)$, namely
$$\langle \ell^{e_1+2e_2}P, R \rangle \tag{15}$$
where $P \in \ker(\gamma)$ is any point of order $\ell^e = \ell^{2e_0+2e_2}$ and R is any point of order ℓ^{e_1} that cannot be divided by ℓ inside $\ker(\gamma)$. This explicit description is of lesser

[†] Hence the adjoint-like notation $\tilde{\tau}$, even though this property is considered regardless of any principal polarizations.

use to us, because the generators in (15) are in general defined over a huge-degree extension of $\mathbb{F}_{p^2}$ only. Before explaining our workaround, let us note that many other subgroups of $\ker(\gamma)$ are equally valid choices: the only crucial features are that the subgroup is isomorphic to $(\mathbb{Z}/\ell^{e_1}\mathbb{Z})^2$ (so that γ_1 is a "good" chain of isogenies), that it contains a point R that is not divisible by ℓ in $\ker(\gamma)$ (so that $\ker(\gamma)/\ker(\gamma_1) \cong (\mathbb{Z}/\ell^{e_1+2e_2}\mathbb{Z})^2$, i.e., also the cofactor γ_2 is a "good" chain of isogenies), and that it concerns a maximal isotropic subgroup with respect to the ℓ^{e_1}-Weil pairing for λ_1 (so that γ_1 is a polarized isogeny).

Lemma 5.6. *Consider any subgroup $G \cong (\mathbb{Z}/\ell^{e_1}\mathbb{Z})^3$ of $A_0[\ell^{e_1}]$ and let $K_1 \subset G[\ell]$ be maximal isotropic with respect to the ℓ-Weil pairing e_{ℓ,λ_1}. Consider the following iterative procedure for $i \geq 2$: let K_i be any subgroup of G for which*

$$K_i \supset K_{i-1}, \qquad K_i \cong \frac{\mathbb{Z}}{\ell^i\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell^i\mathbb{Z}}, \qquad e_{\ell^i,\lambda_1}|_{K_i \times K_i} = 1. \tag{16}$$

Then, regardless of the choices made, this procedure can be repeated up to $i = e_1$, and for every i we have that K_i is maximal isotropic with respect e_{ℓ^i,λ_1}.

Proof. First note that the last statement is immediate from the right-most property in (16): indeed, any isotropic subgroup for the ℓ^i-Weil pairing contains at most ℓ^{2i} elements. Also note that conditions (16) imply $\ell K_i = K_{i-1}$ for all $i = 2, \ldots, e_1$. Writing $K_{i-1} = \langle P_{i-1}, Q_{i-1}\rangle$, we necessarily have $K_i = \langle P_i, Q_i\rangle$ for certain points P_i, Q_i such that $\ell P_i = P_{i-1}$, $\ell Q_i = Q_{i-1}$. Such points can be found in G as long as $i \leq e_1$, so it remains to argue that they can be chosen with $e_{\ell^i,\lambda_1}(P_i, Q_i) = 1$. We know that it concerns some ℓ-th root of unity ζ since

$$e_{\ell^i,\lambda_1}(P_i, Q_i)^\ell = e_{\ell^i,\lambda_1}(P_i, \ell Q_i) = e_{\ell^{i-1},\lambda_1}(\ell P_i, \ell Q_i) = e_{\ell^{i-1},\lambda_1}(P_{i-1}, Q_{i-1}) = 1.$$

On the other hand, since $K_1 = \langle \ell^{i-1}P_i, \ell^{i-1}Q_i\rangle$ is maximal isotropic with respect to the ℓ-Weil pairing there must exist an ℓ-torsion point $X \in G$ such that $e_{\ell,\lambda_1}(X, \ell^{i-1}P_i)$ or $e_{\ell,\lambda_1}(X, \ell^{i-1}Q_i)$ is non-trivial, and by scaling X if needed we can assume that it concerns ζ. Let us assume w.l.o.g. that $e_{\ell,\lambda_1}(X, \ell^{i-1}P_i) = \zeta$. Then using $Q_i' = Q_i + X$ instead of Q_i fixes the issue:

$$e_{\ell^i,\lambda_1}(P_i, Q_i') = e_{\ell^i,\lambda_1}(P_i, Q_i)e_{\ell^i,\lambda_1}(P_i, X) = e_{\ell^i,\lambda_1}(P_i, Q_i)e_{\ell,\lambda_1}(\ell^{i-1}P_i, X) = 1. \square$$

Remark 5.7. Such a subgroup K_1 always exists by general facts from symplectic linear algebra [1, §1.2]. We will apply the lemma to $G = (\ker\gamma)[\ell^{e_1}]$ where this existence comes as no surprise: $K_1 = \ell^{e_1-1}K$, with K as in (15), is an example.

The lemma implies that γ_1 can be built by following a "greedy" approach. We start with any subgroup $K_1 \subset (\ker\gamma)[\ell]$ that

- is maximal isotropic with respect to e_{ℓ,λ_1},
- contains $\ell^{e_1-1}R \in K_1$, with R a point that is not divisible by ℓ in $\ker(\gamma)$.

Then the subgroup K_{e_1} produced by Lemma 5.6 will indeed be a suitable instance of $\ker(\gamma_1)$. A point of the form $\ell^{e_1-1}R$ can be found by taking any order-ℓ point

independent of $\ker(\tau)$ and $\ker(u_1)$, and taking its image under u_1. Once K_1 is fixed, we look for a matrix $\kappa_1 \in \mathrm{M}_2(\mathcal{O}_0)$ with kernel $H_1 = K_1$. We then know that γ factors through κ_1, and continue with the remaining factor $\gamma\kappa_1^{-1}$: we look for any maximal isotropic subgroup $H_2 \cong (\mathbb{Z}/\ell\mathbb{Z})^2$, with corresponding matrix κ_2, which forms a "good" extension of κ_1. We then continue with $\gamma\kappa_1^{-1}\kappa_2^{-1}$, and so on. In this way we implicitly build a tower of subgroups $K_i = (\ker \kappa_i \circ \cdots \circ \kappa_1)$ as in Lemma 5.6. More details can be found in Algorithm 7, which runs in sub-exponential time in view of Steps 5 and 12, but by replacing these steps with a look-up in Appendix A, the attack becomes polynomial-time for $\ell = 2$. We therefore conclude:

Proposition 5.8. *Assuming knowledge of a matrix $g_1 \in \mathrm{Mat}(A_0)$ corresponding to the initial node, the two-dimensional variant of the CGL hash function is not collision-resistant under plausible heuristic assumptions.*

On untrusted set-ups. We deem it likely that all currently known ways for constructing a superspecial principally polarized abelian surface A_1 implicitly reveal an isogeny to E_0^2. This would be analogous to the current situation for supersingular elliptic curves [4].

Up to our knowledge, the candidate ways for construction are:

- either letting $A_1 = E_1 \times E_2$ for supersingular elliptic curves E_1, E_2 (equipped with the product polarization),
- or letting A_1 be the mod-p reduction of a suitable principally polarized abelian surface $\tilde{A}_1/\mathbb{C}$ with complex multiplication (CM) by an order with small discriminant, using results of the type [31, Theorem 1],
- or obtaining A_1 by combining one of the previous constructions with a polarized isogeny walk.

In order to make our suspicion precise, we would need results on reductions of CM curves that are analogous to [11, §5] or [40]. If it is indeed true that A_1 always comes with a path to A_0, then using KLPT^2 and isogeny-to-matrix conversion, an attacker can compute a matrix $g_1 \in \mathrm{Mat}(A_0)$ corresponding to A_1, allowing for an application of Proposition 5.8. In other words, without a trusted set-up, CGL-type hash functions from superspecial principally polarized abelian surfaces can always be broken, at least in sub-exponential time (and in polynomial time if $\ell = 2$).

Example 5.9. In [10, §7] the starting surface is obtained from the Jacobian of $C : y^2 = x^6 + 1$ by means of a $(2,2)$-walk of length 10; this Jacobian is superspecial if and only if $p \equiv 5 \bmod 6$. However, there exists a $(2,2)$-isogeny

$$\begin{aligned} \Phi : \quad \mathrm{Jac}(C) &\to E_1^2, \\ [(x_1, y_1) + (x_2, y_2) - 2\infty] &\mapsto \left((x_1^2, y_1) + (x_2^2, y_2), \left(\frac{1}{x_1^2}, \frac{y_1}{x_1^3}\right) + \left(\frac{1}{x_2^2}, \frac{y_2}{x_2^3}\right)\right), \end{aligned}$$

Algorithm 7: CGLCollision

Input : principally polarized superspecial abelian surface A_1
$g_1 \in \mathrm{Mat}(A_0)$ such that $A_1 \cong (A_0, \mu^{-1}(g_1))$, small prime ℓ
Output: two "good" chains of (ℓ,ℓ)-isogenies $A_1 \to A$

Find $\gamma = u_2 \tau \ell^{e_1} u_1^{-1} \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\gamma^* g_1 \gamma = \ell^e g_1$. `// KLPT`^{2}`.`
$P_1 \leftarrow$ generator of $\ker(u_2)[\ell]$, $P_1 \leftarrow u_1\tilde{\tau}(P_1)$.
$R_1 \leftarrow$ point of $A_0[\ell] \setminus \langle \ker(\tau)[\ell], \ker(u_1)[\ell] \rangle$, $R_1 \leftarrow u_1(R_1)$.
$H_1 \leftarrow \langle P_1, R_1 \rangle$.
$\gamma_1 \leftarrow$ matrix with kernel H_1.
`// Via PIP as in Alg. 3 (steps 6-8) or App. A.`
For $i = 2, \ldots, e_1$ **do**
  $G_i \leftarrow \ker(\gamma\gamma_1^{-1})[\ell]$. `// Note has 3 generators.`
  For $H_i \in \{\text{rank-2 subgroups of } G_i\}$ **do**
    If $H_i \cap \ker \tilde{\gamma}_1 = \{0\}$ **then**
      `// Checks if chain is "good". Note` $\tilde{\gamma}_1 = 2^{i-1}\gamma_1^{-1}$.
      $\kappa \leftarrow$ matrix with kernel H_i.
      `// Via PIP as in Alg. 3 (steps 6-8) or App. A.`
      If $2^i(\kappa\gamma_1)^{*-1} g_1 (\kappa\gamma_1)^{-1} \in \mathrm{Mat}(A_0)$ **then**
        `// Checks if group is maximal isotropic.`
        Break "**For** H_i"-loop.
  $\gamma_1 \leftarrow \kappa\gamma_1$.
$\gamma_2 \leftarrow \gamma\gamma_1^{-1}$.
Convert γ_1 into chain of (ℓ,ℓ)-isogenies $\varphi_{1e_1} \circ \cdots \circ \varphi_{11} : A_1 \to A$.
Convert γ_2 into chain of (ℓ,ℓ)-isogenies $\varphi_{2,e_1+2e_2} \circ \cdots \circ \varphi_{21} : A_1 \to A$.
`// Using method from Section 5.2(ii).`
Return $\varphi_{1e_1} \circ \cdots \circ \varphi_{11}$, $\varphi_{2,e_1+2e_2} \circ \cdots \circ \varphi_{21}$.

where $E_1 : y^2 = x^3 + 1$, which is indeed supersingular if and only if $p \equiv 5 \bmod 6$. Under our assumption $p \equiv 3 \bmod 4$, we can then connect E_1^2 to E_0^2 with a product isogeny, stemming from a known isogeny from E_1 to E_0 (e.g., as detailed in [11, Example 20]). In a very recent paper [44], the authors find collisions for the hash function from [10] using knowledge of a *short* isogeny to E_1^2, but our results show that *any* known isogeny will do, in fact.

5.5 Attacks on verifiable delay functions.

In [14] the authors propose a 2-dimensional analog of the isogeny-based VDF from [22]. One of the (multiple) drawbacks of [22] is that it needs a trusted set-up because the KLPT algorithm provides a solution to the isogeny short-cut problem if the endomorphism ring of the starting surface is known. The selling

point of a genus-2 version of essentially the same VDF was that no trusted set-up would be necessary. Our results invalidate this selling point: a trusted set-up is necessary in the genus-2 case as well, as otherwise a similar attack applies.

6 Further research directions

We provide several future research directions that could be investigated further.

6.1 Improving the length of the path

At the moment, the bound on the reduced degree N produced by Theorems 3.1 and 3.14 is relatively large. One source for this is that γ is built from two matrices u_1, u_2, glued together by means of the matrix τ from Lemma 3.2 whose reduced norm scales quartically with the common top-left entry ℓ^{e_2} of $u_1^* g_1 u_1$ and $u_2^* g_2 u_2$. However, the main room for improvement seems to lie in the fact that each u_i itself is built in two steps:

- first we need to reduce the matrix g_i, in the sense of Definition 3.11, in order to bound the entries in a way that only depends on p,
- this is then used to construct u_i via Theorem 3.12, which moreover produces a rather large value of ℓ^{e_2}.

We conjecture that this two-step procedure could be avoided through a better understanding of the quadratic form $K((a,c)) = s \cdot \mathsf{n}(a) + t \cdot \mathsf{n}(c) + \mathsf{tr}(\bar{c}ra)$. Namely, we proved in Proposition 3.10 that the determinant of K is $(p/4)^4$, thus in particular independent of s, t, r. Thus it seems plausible that a more direct method would provide a better output.

6.2 Ensuring that $\ker(\gamma)$ is free of rank 2

Recall that in several applications, notably CGL-style hash functions, one is mainly interested in polarized isogenies γ for which $\ker(\gamma) \cong (\mathbb{Z}/N\mathbb{Z})^2$. As argued in Corollary 5.5, the paths returned in Theorems 3.1 and 3.14 are never of this type. Changing our KLPT^2 algorithm such that it always outputs isogenies whose kernels are free of rank 2 is an interesting future research goal. Positive results in this direction may also lead to proofs of the connectedness of the (ℓ, ℓ)-isogeny graph of superspecial principally polarized abelian surfaces by means of "good" extensions [10, Conjecture 3], an open problem that was not settled by the recent work by Jordan and Zaytman [36].

6.3 Endomorphism ring representations

Recall that the endomorphism ring computation problem is the central hard problem in (supersingular) isogeny-based elliptic curve cryptography. As explained in Section 5.1, the natural analog in dimension 2 is about computing the matrix g associated with a given principally polarized superspecial abelian

variety.[†] However, even in dimension one, two versions of this "endomorphism ring computation problem" exist: given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$,

- find a maximal order $\mathcal{O} \subset B_{p,\infty}$ such that $\mathrm{End}(E) \cong \mathcal{O}$,
- find four endomorphisms that generate $\mathrm{End}(E)$ as a $\mathbb{Z}$-module and that one can evaluate efficiently.

In [52] the first problem is called MaxOrder and the second is called EndRing. As it turns out, these problems are polynomial-time equivalent.

Computing the matrix g is the natural two-dimensional generalization of the MaxOrder problem. The natural analogue of the EndRing problem is that one requires to find 16 endomorphisms of a given principally polarized superspecial abelian variety A, with the following properties:

- the endomorphisms generate $\mathrm{End}(A)$ as a $\mathbb{Z}$-module and can be evaluated efficiently,
- the Rosati involution can be efficiently evaluated on these endomorphisms.

Here we outline a potential strategy for proving that the two problems are polynomial-time equivalent. First, if one is given a matrix g, then using our algebraic pathfinding algorithm with powersmooth degrees, by connecting the surface to $A_0 = E_0^2$ we can get an efficient representation of $\mathrm{End}(A)$ via lollipopping.

We sketch a potential proof for the converse direction, leaving a more precise version for future work. First, we compute an explicit isomorphism between $\mathrm{End}^0(A) = \mathrm{End}(A) \otimes_{\mathbb{Z}} \mathbb{Q}$ and $\mathrm{M}_2(B_{p,\infty})$ using [35]. A priori, this algorithm requires factoring, but this can be avoided. The way [35] works is that it computes a maximal order in $\mathrm{End}^0(A)^{\mathrm{op}} \otimes_{\mathbb{Q}} \mathrm{M}_2(B_{p,\infty})$, where $\cdot^{\mathrm{op}}$ refers to the opposite algebra, and then uses lattice reduction to find a zero divisor. Computing a maximal order usually requires factoring the discriminant of a starting order. However, starting from maximal orders in $\mathrm{End}^0(A)$ and $\mathrm{M}_2(B_{p,\infty})$, their tensor product is maximal away from p and ∞, so the factoring issue becomes trivial (for further discussion, see [19, Proposition 4.1] where this is explained for quaternion orders, but the argument is the same). Furthermore, a zero divisor might not immediately give an explicit isomorphism, this is only true if it has rank 1 (when identified with an element of $\mathrm{M}_{16}(\mathbb{C})$). Luckily [35] shows that in low dimensional cases, one immediately finds a rank 1 element.

Such an explicit isomorphism provides an embedding $\iota : \mathrm{End}(A) \hookrightarrow \mathrm{M}_2(B_{p,\infty})$ together with an involution σ on $\iota(\mathrm{End}(A))$ that comes from the Rosati involution. Our next goal is to find an explicit conjugation between $\iota(\mathrm{End}(A))$ and $\mathrm{M}_2(\mathcal{O}_0)$. This is the main step that needs to be made more explicit, but a potential idea could be the following. Let $B(u, v) := \mathsf{tr}(\sigma(u)v)$ be a bilinear map that by the properties of the Rosati involution equips $\iota(\mathrm{End}(A))$ with a Euclidean norm. One can also consider $\mathrm{M}_2(\mathcal{O}_0)$ together with the conjugate transpose and

[†]Recall from Section 5.2 that we can solve this problem in sub-exponential time as soon as a smooth degree isogeny to A_0 is known, and in polynomial time in case this concerns a chain of $(2, 2)$-isogenies.

the associated bilinear form $\mathsf{tr}(u^*v)$. Now one can run HKZ and list all the shortest elements and try to match them up. For $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{M}_2(\mathcal{O}_0)$ this norm is just $\mathsf{n}(a) + \mathsf{n}(b) + \mathsf{n}(c) + \mathsf{n}(d)$, so it is easy to see that there are not too many short elements. If one can match them up, then the explicit conjugation turns into solving a system of linear equations.

We now have found an explicit isomorphism between $\iota(\mathrm{End}(A))$ and $\mathrm{M}_2(\mathcal{O}_0)$ on which we have two involutions: conjugate-transpose and the involution induced by σ. These involutions will be conjugated by a Hermitian matrix (i.e., symmetric with respect to conjugate-transpose). This conjugating element is going to be g, and it can be found via linear algebra, as explained in Remark 2.8.

Remark 6.1. The involution on $\mathrm{M}_2(\mathcal{O}_0)$ is not going to be uniquely determined; consequently the same is true for the matrix g. But this is expected, as we are looking for an equivalence class in $\mathrm{Mat}^0(A_0)$.

6.4 Generalizing other known applications of KLPT

The most celebrated application of the KLPT algorithm is building signature schemes. In this paper we have not studied this aspect, even though we were motivated by several gaps in Chu's attempt from [16, Appendix A] to build a GPS-style signature scheme in dimension two. Explicitly, he asked for efficient routines called PowersmoothMatrix, now resolved by our KLPT^2 algorithm, for MatrixToIsogenyPath, addressed in Sections 4.1 and 14, and for IsogenyPathToMatrix. In the latter case, we described a polynomial-time method for chains of $(2,2)$-isogenies in Section 14. But a general polynomial-time solution is lacking: currently we still need to resort to Chu's sub-exponential time method for the PIP in $\mathrm{M}_2(\mathcal{O}_0)$. So this is clearly a compelling research question.

An attractive goal is develop a dimension-two variant of SQIsign, for which more restricted solutions of the isogeny-to-matrix conversion problem, only targeting isogenies of very small degree, could be good enough. One important open question is whether the paths returned by KLPT^2 reveal information about the endomorphism rings of the domain and the codomain. One interesting aspect, as opposed to KLPT, is that we find a direct path between two surfaces without going through one fixed special principally polarized abelian surface. In dimension one this problem was essentially resolved by the Generalized KLPT algorithm [21, §5]. It would also be interesting to study higher dimensional analogs of the various HD versions of SQIsign [3, 20, 25, 43]. A potential advantage of higher-dimensional generalizations of SQIsign is that the complexity of endomorphism computation grows with the dimension: this could allow for schemes working with smaller field sizes (i.e., smaller p).

Bibliography

[1] Vladimir I. Arnold, Alexander B. Givental, and Sergei P. Novikov. *Symplectic Geometry*, pages 1–138. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001. https://doi.org/10.1007/978-3-662-06791-8_1.

[2] Andreas Bächle, Geoffrey Janssens, Eric Jespers, Ann Kiefer, and Doryan Temmerman. A dichotomy for integral group rings via higher modular groups as amalgamated products. *Journal of Algebra*, 604:185–223, 2022. https://doi.org/10.1016/j.jalgebra.2022.03.044.

[3] Andrea Basso, Luca De Feo, Pierrick Dartois, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. SQIsign2D-west: The fast, the small, and the safer. In *ASIACRYPT 2024 Part III*, volume 15486 of *LNCS*, pages 339–370. Springer, 2024. https://doi.org/10.1007/978-981-96-0891-1_11.

[4] Jeremy Booher, Ross Bowden, Javad Doliskani, Tako Boris Fouotsa, Steven D Galbraith, Sabrina Kunzweiler, Simon-Philipp Merz, Christophe Petit, Benjamin Smith, Katherine E Stange, Yan Bo Ti, Christelle Vincent, José Felipe Voloch, Charlotte Weitkämper, and Lukas Zobernig. Failing to hash into supersingular isogeny graphs. *The Computer Journal*, 67(8):2702–2719, 2024. https://doi.org/10.1093/comjnl/bxae038.

[5] Irene Bouw, Jenny Cooley, Kristin Lauter, Elisa Lorenzo García, Michelle Manes, Rachel Newton, and Ekin Ozman. Bad reduction of genus three curves with complex multiplication. In *Women in Numbers Europe*, volume 2, pages 109–151. Springer International Publishing, 2015. https://doi.org/10.1007/978-3-319-17987-2_5.

[6] Bradley Brock. *Superspecial curves of genera two and three*. PhD thesis, Princeton University, 1994.

[7] Wouter Castryck, Mingjie Chen, Riccardo Invernizzi, Gioella Lorenzon, and Frederik Vercauteren. Breaking and repairing SQIsign2D-East. Cryptology ePrint Archive, Paper 2024/1453, 2024. https://eprint.iacr.org/2024/1453.

[8] Wouter Castryck and Thomas Decru. Multiradical isogenies. In 18^{th} *International Conference on Arithmetic, Geometry, Cryptography and Coding Theory*, volume 779 of *Cont. Math.*, pages 57–89, 2022. https://doi.org/10.1090/conm/779.

[9] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In *EUROCRYPT 2023 Part V*, volume 14008 of *LNCS*, pages 423–447. Springer, 2023. https://doi.org/10.1007/978-3-031-30589-4_15.

[10] Wouter Castryck, Thomas Decru, and Benjamin Smith. Hash functions from superspecial genus-2 curves using Richelot isogenies. *Journal of Mathematical Cryptology*, 14(1):268–292, 2020. https://doi.org/10.1515/jmc-2019-0021.

[11] Wouter Castryck, Lorenz Panny, and Frederik Vercauteren. Rational isogenies from irrational endomorphisms. In *EUROCRYPT 2020 Part II*, volume 12106 of *LNCS*, pages 523–548. Springer, 2020. https://doi.org/10.1007/978-3-030-45724-2_18.

[12] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, 2009. https://doi.org/10.1007/s00145-007-9002-x.

[13] Jorge Chavez-Saab, Maria Corte-Real Santos, Luca De Feo, Jonathan Komada Eriksen, Basil Hess, David Kohel, Antonin Leroux, Patrick Longa,

Michael Meyer, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Francisco Rodríguez Henríquez, Sina Schaeffler, and Benjamin Wesolowski. SQIsign: Algorithm specifications and supporting documentation, v1.0. available at https://sqisign.org/.

[14] Chao Chen and Fangguo Zhang. Verifiable delay functions and delay encryptions from hyperelliptic curves. *Cybersecurity*, 6(1):54, 2023. https://doi.org/10.1186/s42400-023-00189-2.

[15] Mingjie Chen, Muhammad Imran, Gábor Ivanyos, Péter Kutas, Antonin Leroux, and Christophe Petit. Hidden stabilizers, the isogeny to endomorphism ring problem and the cryptanalysis of pSIDH. In *ASIACRYPT 2023 Part III*, volume 14440 of *LNCS*, pages 99–130. Springer, 2023. https://doi.org/10.1007/978-981-99-8727-6_4.

[16] Hao-Wei Chu. *Algorithms for abelian surfaces over finite fields and their applications to cryptography.* Phd thesis, Pennsylvania State University, 2021. Available at https://etda.libraries.psu.edu/files/final_submissions/24383.

[17] Brian Conrad. Polarizations, 2004. Notes of the VIGRE Number Theory Working Group, available at https://math.stanford.edu/~conrad/vigregroup/vigre04/polarization.pdf.

[18] Romain Cosset and Damien Robert. Computing (ℓ, ℓ)-isogenies in polynomial time on jacobians of genus 2 curves. *Mathematics of Computation*, 84(294):1953–1975, 2015. http://www.jstor.org/stable/24489183.

[19] Tímea Csahók, Péter Kutas, Mickaël Montessinos, and Gergely Zábrádi. Explicit isomorphisms of quaternion algebras over quadratic global fields. *Research in Number Theory*, 8(4):77, 2022. https://doi.org/10.1007/s40993-022-00380-3.

[20] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: New dimensions in cryptography. In *EUROCRYPT 2024 Part I*, volume 14651 of *LNCS*, pages 3–32. Springer, 2024. https://doi.org/10.1007/978-3-031-58716-0_1.

[21] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: Compact post-quantum signatures from quaternions and isogenies. In *ASIACRYPT 2020 Part I*, volume 12491 of *LNCS*, pages 64–93. Springer, 2020. https://doi.org/10.1007/978-3-030-64837-4_3.

[22] Luca De Feo, Simon Masson, Christophe Petit, and Antonio Sanso. Verifiable delay functions from supersingular isogenies and pairings. In *ASIACRYPT 2019 Part I*, volume 11921 of *LNCS*, pages 248–277. Springer, 2019. https://doi.org/10.1007/978-3-030-34578-5_10.

[23] Max Deuring. Die Typen der Multiplikatorenringe Elliptischer Funktionenkörper. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 14, pages 197–272. Springer Berlin/Heidelberg, 1941.

[24] Peter K. Draxl. *Skew Fields*, volume 81 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1983. https://doi.org/10.1017/CBO9780511661907.

[25] Max Duparc and Tako Boris Fouotsa. SQIPrime: A dimension 2 variant of SQISignHD with non-smooth challenge isogenies. In *ASIACRYPT 2024 Part III*, volume 15486 of *LNCS*, pages 396–429. Springer, 2024. https://doi.org/10.1007/978-981-96-0891-1_13.

[26] Kirsten Eisenträger, Sean Hallgren, Kristin E. Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In *EUROCRYPT 2018 Part III*, volume 10822 of *LNCS*, pages 329–368. Springer, 2018. https://doi.org/10.1007/978-3-319-78372-7_11.

[27] Jonathan Komada Eriksen, Lorenz Panny, Jana Sotáková, and Mattia Veroni. Deuring for the people: Supersingular elliptic curves with prescribed endomorphism ring in general characteristic. In *LuCaNT: LMFDB, Computation, and Number Theory*, Cont. Math., pages 339–365, 2023. https://doi.org/10.1090/conm/796/16008.

[28] E. Victor Flynn and Yan Bo Ti. Genus two isogeny cryptography. In *Post-Quantum Cryptography*, volume 11505 of *LNCS*, pages 286–306. Springer International Publishing, 2019. https://doi.org/10.1007/978-3-030-25510-7_16.

[29] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In *ASIACRYPT 2017 Part I*, volume 10624 of *LNCS*, pages 3–33. Springer, 2017. https://doi.org/10.1007/978-3-319-70694-8_1.

[30] Pierrick Gaudry, Julien Soumier, and Pierre-Jean Spaenlehauer. Isogeny-based cryptography using isomorphisms of superspecial abelian surfaces. Cryptology ePrint Archive, Paper 2025/136, 2025. https://eprint.iacr.org/2025/136.

[31] Eyal Z. Goren. On certain reduction problems concerning abelian surfaces. *Manuscripta Mathematica*, 94(1):33–43, 1997.

[32] Alexandre Gélin, Everett Howe, and Christophe Ritzenthaler. Principally polarized squares of elliptic curves with field of moduli equal to $\mathbb{Q}$. *The Open Book Series*, 2(1):257–274, 2019. http://dx.doi.org/10.2140/obs.2019.2.257.

[33] Everett W. Howe, Franck Leprévost, and Bjorn Poonen. Large torsion subgroups of split jacobians of curves of genus two or three. *Forum Mathematicum*, 12(3):315–364, 2000. https://doi.org/10.1515/form.2000.008.

[34] Tomoyoshi Ibukiyama, Toshiyuki Katsura, and Frans Oort. Supersingular curves of genus two and class numbers. *Compositio Mathematica*, 57(2):127–152, 1986. http://eudml.org/doc/89752.

[35] Gábor Ivanyos, Lajos Rónyai, and Josef Schicho. Splitting full matrix algebras over algebraic number fields. *Journal of Algebra*, 354(1):211–223, 2012. https://doi.org/10.1016/j.jalgebra.2012.01.008.

[36] Bruce W. Jordan and Yevgeny Zaytman. Isogeny graphs of superspecial abelian varieties and Brandt matrices. *Mathematische Zeitschrift*, 2024. To appear, preprint available at https://arxiv.org/pdf/2005.09031.pdf.

[37] Ernst Kani. The moduli spaces of jacobians isomorphic to a product of two elliptic curves. *Collectanea Mathematica*, 67:21–54, 2015. https://doi.org/10.1007/s13348-015-0148-9.

[38] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion isogeny path problem. *LMS Journal of Computation and Mathematics*, 17(A):418–432, 2014. https://doi.org/10.1112/S1461157014000151.

[39] 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. In *PKC 2025*, LNCS. Springer, 2025. To appear, preprint available at https://eprint.iacr.org/2024/1732.

[40] Jonathan Love and Dan Boneh. Supersingular curves with small non-integer endomorphisms. In *ANTS XIV*, volume 4 of *The Open Book Series*, pages 7–22. MSP, 2020. https://doi.org/10.2140/obs.2020.4.7.

[41] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In *EUROCRYPT 2023 Part V*, volume 14008 of *LNCS*, pages 448–471. Springer, 2023. https://doi.org/10.1007/978-3-031-30589-4_16.

[42] James Milne. Abelian varieties, version 2.0, 2008. Course notes available at https://www.jmilne.org/math/CourseNotes/av.html.

[43] Kohei Nakagawa, Hiroshi Onuki, Wouter Castryck, Mingjie Chen, Riccardo Invernizzi, Gioella Lorenzon, and Frederik Vercauteren. SQIsign2D-East: A new signature scheme using 2-dimensional isogenies. In *ASIACRYPT 2024 Part III*, volume 15486 of *LNCS*, pages 272–303. Springer, 2024. https://doi.org/10.1007/978-981-96-0891-1_9.

[44] Ryo Ohashi and Hiroshi Onuki. An efficient collision attack on Castryck–Decru–Smith's hash function. In *PQCrypto 2025*, LNCS. Springer, 2025. To appear, preprint available at https://eprint.iacr.org/2024/1776.

[45] Christophe Petit and Spike Smith. An improvement to the quaternion analogue of the ℓ-isogeny problem. *Presentation at MathCrypt*, 2018.

[46] Damien Robert. Breaking SIDH in polynomial time. In *EUROCRYPT 2023 Part V*, volume 14008 of *LNCS*, pages 472–503. Springer, 2023. https://doi.org/10.1007/978-3-031-30589-4_17.

[47] Damien Robert. On the efficient representation of isogenies: a survey for NuTMiC 2024. In *NuTMiC 2024*, volume 14966 of *LNCS*, pages 3–84, 2025. https://doi.org/10.1007/978-3-031-82380-0_1.

[48] Tetsuji Shioda. Supersingular K3 surfaces. In Knud Lønsted, editor, *Algebraic Geometry*, pages 564–591. Springer Berlin Heidelberg, 1979.

[49] Benjamin Smith. *Explicit Endomorphisms and Correspondences*. PhD thesis, University of Sydney, 2005. https://www.academia.edu/77805612/Explicit_endomorphisms_and_Correspondences.

[50] Katsuyuki Takashima. Efficient algorithms for isogeny sequences and their cryptographic applications. In *Mathematical Modelling for Next-Generation Cryptography. Mathematics for Industry*, volume 29, pages 97–114, Singapore, 2018. Springer. https://doi.org/10.1007/978-981-10-5065-7_6.

[51] John Voight. *Quaternion algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer Nature, 2021. https://doi.org/10.1007/978-3-030-56694-4.

[52] Benjamin Wesolowski. Orientations and the supersingular endomorphism ring problem. In *EUROCRYPT 2022 Part III*, volume 13277 of *LNCS*, pages 345–371. Springer, 2022. https://doi.org/10.1007/978-3-031-07082-2_13.

[53] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1100–1111. IEEE, 2022. https://doi.org/10.1109/FOCS52979.2021.00109.

A Matrices encoding (2, 2)-isogenies

Consider the elliptic curve $E_0 : y^2 = x^3 + x$ over $\mathbb{F}_{p^2}$, where $p \equiv 3 \bmod 4$. Let $i \in \mathbb{F}_{p^2}$ denote a fixed square root of -1. Recall that we identify $\mathrm{End}(E_0)$ with the maximal order

$$\mathcal{O}_0 = \left\langle 1, i, \frac{i+j}{2}, \frac{1+k}{2} \right\rangle \subset B_{p,\infty}$$

where, abusing notation, i is identified with the automorphism $(x, y) \mapsto (-x, iy)$ and j is identified with the Frobenius map $(x, y) \mapsto (x^p, y^p)$. For ease of notation, we will write

$$\omega_3 = \frac{i+j}{2}, \qquad \omega_4 = \frac{1+k}{2}.$$

We also introduce the following notation for the points of order 2:

$$P_0 = (0, 0), \qquad P_i = (i, 0), \qquad \bar{P}_i = (-i, 0).$$

A calculation reveals the following evaluation tables:

	P_0	P_i	$\bar{P}_i$
i	P_0	$\bar{P}_i$	P_i
ω_3	P_i	P_0	$\bar{P}_i$
ω_4	P_i	$\bar{P}_i$	P_0

if $p \equiv 3 \bmod 8$,

	P_0	P_i	$\bar{P}_i$
i	P_0	$\bar{P}_i$	P_i
ω_3	$\bar{P}_i$	$\bar{P}_i$	∞
ω_4	$\bar{P}_i$	∞	$\bar{P}_i$

if $p \equiv 7 \bmod 8$.

Using this, one can verify for each of the $\begin{bmatrix} 4 \\ 2 \end{bmatrix}_2 = 35$ subgroups $K \subset A_0 = E_0^2$ with $K \cong (\mathbb{Z}/2\mathbb{Z})^2$ that K can be realized as the kernel of the reduced-norm-4 matrix indicated in the table below:

subgroup K	$p \equiv 3 \bmod 8$	$p \equiv 7 \bmod 8$
$E_0[2] \times \{\infty\}$	$\begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$
$\{\infty\} \times E_0[2]$	$\begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}$

$\{(P_0,P_0),(P_i,P_i),(\bar{P}_i,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$
$\{(P_0,P_0),(P_i,\bar{P}_i),(\bar{P}_i,P_i),(\infty,\infty)\}$	$\begin{pmatrix} 1 & i \\ -1 & i \end{pmatrix}$	$\begin{pmatrix} 1 & i \\ -1 & i \end{pmatrix}$
$\{(P_0,P_0),(\infty,P_0),(P_0,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 1+i & 0 \\ 0 & 1+i \end{pmatrix}$	$\begin{pmatrix} 1+i & 0 \\ 0 & 1+i \end{pmatrix}$
$\{(P_i,P_0),(\bar{P}_i,P_0),(P_0,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 0 & 1+i \\ 1+i & i \end{pmatrix}$	$\begin{pmatrix} 0 & 1+i \\ 1+i & i \end{pmatrix}$
$\{(P_i,P_i),(\bar{P}_i,P_0),(P_0,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} \omega_4 & i \\ i+\omega_3 & 1 \end{pmatrix}$	$\begin{pmatrix} i & i-\omega_4 \\ -i & i+\omega_4 \end{pmatrix}$
$\{(P_0,P_i),(P_i,\bar{P}_i),(\bar{P}_i,P_0),(\infty,\infty)\}$	$\begin{pmatrix} \omega_4 & -1 \\ i+\omega_3 & i \end{pmatrix}$	$\begin{pmatrix} i & 1-\omega_3 \\ -i & 1+\omega_3 \end{pmatrix}$
$\{(P_0,P_i),(P_0,\bar{P}_i),(\infty,P_0),(\infty,\infty)\}$	$\begin{pmatrix} i & 1+i \\ 1+i & 0 \end{pmatrix}$	$\begin{pmatrix} i & 1+i \\ 1+i & 0 \end{pmatrix}$
$\{(P_0,P_i),(P_i,P_0),(\bar{P}_i,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} \omega_3 & 1 \\ 1+\omega_4 & i \end{pmatrix}$	$\begin{pmatrix} 1 & 1-\omega_3 \\ -1 & 1+\omega_3 \end{pmatrix}$
$\{(P_0,\bar{P}_i),(\bar{P}_i,P_i),(P_i,P_0),(\infty,\infty)\}$	$\begin{pmatrix} 1 & \omega_4 \\ -i & i+\omega_3 \end{pmatrix}$	$\begin{pmatrix} 1 & i-\omega_4 \\ -1 & i+\omega_4 \end{pmatrix}$
$\{(\bar{P}_i,P_0),(\bar{P}_i,\bar{P}_i),(\infty,P_i),(\infty,\infty)\}$	$\begin{pmatrix} 1 & i-\omega_4 \\ -1 & i+\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1 & \omega_4 \\ -i & i+\omega_3 \end{pmatrix}$
$\{(P_i,P_0),(P_i,\bar{P}_i),(\infty,P_i),(\infty,\infty)\}$	$\begin{pmatrix} i & i-\omega_4 \\ -i & i+\omega_4 \end{pmatrix}$	$\begin{pmatrix} i & \omega_4 \\ 1 & i+\omega_3 \end{pmatrix}$
$\{(P_i,P_0),(P_i,P_i),(\infty,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} i & 1-\omega_3 \\ -i & 1+\omega_3 \end{pmatrix}$	$\begin{pmatrix} 1 & 1+\omega_4 \\ -i & \omega_3 \end{pmatrix}$
$\{(\bar{P}_i,P_0),(\bar{P}_i,P_i),(\infty,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} 1 & 1-\omega_3 \\ -1 & 1+\omega_3 \end{pmatrix}$	$\begin{pmatrix} 1 & \omega_3 \\ i & 1+\omega_4 \end{pmatrix}$
$\{(P_i,P_i),(P_0,P_i),(\bar{P}_i,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 1+\omega_3 & i \\ -1+\omega_3 & i \end{pmatrix}$	$\begin{pmatrix} 1+\omega_4 & -1 \\ \omega_3 & i \end{pmatrix}$
$\{(P_0,P_i),(\bar{P}_i,P_i),(P_i,\infty),(\infty,\infty)\}$	$\begin{pmatrix} i+\omega_4 & i \\ -i+\omega_4 & i \end{pmatrix}$	$\begin{pmatrix} \omega_4 & i \\ i+\omega_3 & 1 \end{pmatrix}$
$\{(P_0,\bar{P}_i),(\bar{P}_i,\bar{P}_i),(P_i,\infty),(\infty,\infty)\}$	$\begin{pmatrix} i+\omega_4 & 1 \\ -i+\omega_4 & 1 \end{pmatrix}$	$\begin{pmatrix} \omega_4 & -1 \\ i+\omega_3 & i \end{pmatrix}$
$\{(P_0,\bar{P}_i),(P_i,\bar{P}_i),(\bar{P}_i,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 1+\omega_3 & 1 \\ -1+\omega_3 & 1 \end{pmatrix}$	$\begin{pmatrix} \omega_3 & 1 \\ 1+\omega_4 & i \end{pmatrix}$
$\{(P_i,P_i),(P_i,\bar{P}_i),(\infty,P_0),(\infty,\infty)\}$	$\begin{pmatrix} 1 & \omega_3+\omega_4 \\ -i & 1+i+\omega_3-\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1 & 1+i+\omega_3-\omega_4 \\ i & \omega_3+\omega_4 \end{pmatrix}$
$\{(\bar{P}_i,\bar{P}_i),(\bar{P}_i,P_i),(\infty,P_0),(\infty,\infty)\}$	$\begin{pmatrix} 1 & 1+i+\omega_3-\omega_4 \\ i & \omega_3+\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1 & \omega_3+\omega_4 \\ -i & 1+i+\omega_3-\omega_4 \end{pmatrix}$
$\{(P_0,P_0),(P_0,P_i),(\infty,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} 1 & i-\omega_3+\omega_4 \\ -i & i+\omega_3+\omega_4 \end{pmatrix}$	$\begin{pmatrix} i & 1+\omega_3-\omega_4 \\ -1 & 1+\omega_3+\omega_4 \end{pmatrix}$
$\{(P_0,P_0),(P_0,\bar{P}_i),(\infty,P_i),(\infty,\infty)\}$	$\begin{pmatrix} i & 1+\omega_3-\omega_4 \\ -1 & 1+\omega_3+\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1 & i-\omega_3+\omega_4 \\ -i & i+\omega_3+\omega_4 \end{pmatrix}$
$\{(\bar{P}_i,\bar{P}_i),(P_i,\bar{P}_i),(P_0,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 1+i+\omega_3-\omega_4 & 1 \\ \omega_3+\omega_4 & i \end{pmatrix}$	$\begin{pmatrix} 1+i+\omega_3-\omega_4 & -i \\ \omega_3+\omega_4 & 1 \end{pmatrix}$
$\{(P_i,P_i),(\bar{P}_i,P_i),(P_0,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 1+i+\omega_3-\omega_4 & -i \\ \omega_3+\omega_4 & 1 \end{pmatrix}$	$\begin{pmatrix} 1+i+\omega_3-\omega_4 & 1 \\ \omega_3+\omega_4 & i \end{pmatrix}$
$\{(P_0,P_0),(\bar{P}_i,P_0),(P_i,\infty),(\infty,\infty)\}$	$\begin{pmatrix} 1+\omega_3-\omega_4 & 1 \\ 1+\omega_3+\omega_4 & i \end{pmatrix}$	$\begin{pmatrix} i+\omega_3-\omega_4 & -i \\ -i+\omega_3+\omega_4 & 1 \end{pmatrix}$
$\{(P_0,P_0),(P_i,P_0),(\bar{P}_i,\infty),(\infty,\infty)\}$	$\begin{pmatrix} i+\omega_3-\omega_4 & -i \\ -i+\omega_3+\omega_4 & 1 \end{pmatrix}$	$\begin{pmatrix} 1+\omega_3-\omega_4 & 1 \\ 1+\omega_3+\omega_4 & i \end{pmatrix}$
$\{(P_i,P_i),(P_i,\infty),(\infty,P_i),(\infty,\infty)\}$	$\begin{pmatrix} 1+i+\omega_3 & i-\omega_4 \\ i+\omega_4 & -1+i+\omega_3 \end{pmatrix}$	$\begin{pmatrix} i+\omega_3 & -\omega_4 \\ \omega_4 & i+\omega_3 \end{pmatrix}$
$\{(P_i,\bar{P}_i),(P_i,\infty),(\infty,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} i+\omega_4 & 1+i+\omega_4 \\ 1+i+\omega_3 & 1+\omega_3 \end{pmatrix}$	$\begin{pmatrix} i+\omega_3 & \omega_3 \\ \omega_4 & 1+\omega_4 \end{pmatrix}$
$\{(\bar{P}_i,\bar{P}_i),(\bar{P}_i,\infty),(\infty,\bar{P}_i),(\infty,\infty)\}$	$\begin{pmatrix} 1+\omega_3 & 1+i-\omega_4 \\ -1-i+\omega_4 & 1+\omega_3 \end{pmatrix}$	$\begin{pmatrix} \omega_3 & -1-\omega_4 \\ 1+\omega_4 & \omega_3 \end{pmatrix}$
$\{(\bar{P}_i,P_i),(\bar{P}_i,\infty),(\infty,P_i),(\infty,\infty)\}$	$\begin{pmatrix} 1+i+\omega_4 & -i-\omega_4 \\ 1-i+\omega_4 & i-\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1+\omega_4 & \omega_4 \\ \omega_3 & i+\omega_3 \end{pmatrix}$
$\{(P_i,P_0),(P_i,\infty),(\infty,P_0),(\infty,\infty)\}$	$\begin{pmatrix} 1+\omega_3+\omega_4 & 1+i+\omega_3-\omega_4 \\ 1+\omega_3-\omega_4 & -\omega_3-\omega_4 \end{pmatrix}$	$\begin{pmatrix} i+\omega_3+\omega_4 & \omega_3+\omega_4 \\ i-\omega_3+\omega_4 & 1+i-\omega_3+\omega_4 \end{pmatrix}$
$\{(\bar{P}_i,P_0),(\bar{P}_i,\infty),(\infty,P_0),(\infty,\infty)\}$	$\begin{pmatrix} i+\omega_3+\omega_4 & \omega_3+\omega_4 \\ i-\omega_3+\omega_4 & 1+i-\omega_3+\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1+\omega_3+\omega_4 & 1+i+\omega_3-\omega_4 \\ 1+\omega_3-\omega_4 & -\omega_3-\omega_4 \end{pmatrix}$

$\{(P_0, \bar{P}_i), (P_0, \infty), (\infty, \bar{P}_i), (\infty, \infty)\}$	$\begin{pmatrix} 1+i+\omega_3-\omega_4 & i+\omega_3-\omega_4 \\ \omega_3+\omega_4 & -i+\omega_3+\omega_4 \end{pmatrix}$	$\begin{pmatrix} \omega_3+\omega_4 & 1+\omega_3+\omega_4 \\ 1+i+\omega_3-\omega_4 & 1+\omega_3-\omega_4 \end{pmatrix}$
$\{(P_0, P_i), (P_0, \infty), (\infty, P_i), (\infty, \infty)\}$	$\begin{pmatrix} \omega_3+\omega_4 & 1+\omega_3+\omega_4 \\ 1+i+\omega_3-\omega_4 & 1+\omega_3-\omega_4 \end{pmatrix}$	$\begin{pmatrix} 1+i+\omega_3-\omega_4 & i+\omega_3-\omega_4 \\ \omega_3+\omega_4 & -i+\omega_3+\omega_4 \end{pmatrix}$

B Probability of non-backtracking

Lemma B.1. *Inside the vector space $\mathbb{F}_\ell^4$, the probability that two random 2-dimensional subspaces intersect trivially is $\ell^4/(\ell^2+1)(\ell^2+\ell+1)$. The probability that a random 1-dimensional subspace and a random 3-dimensional subspace intersect trivially is $\ell^3/(\ell^3+\ell^2+\ell+1)$.*

Proof. We only prove the first formula; the second formula follows similarly. It is well-known that the number of 2-dimensional subspaces is given by the Gaussian binomial coefficient

$$\begin{bmatrix} 4 \\ 2 \end{bmatrix}_\ell = (\ell^2+1)(\ell^2+\ell+1).$$

The formula follows because the number of 2-dimensional subspaces trivially intersecting a given 2-dimensional subspace $V \subset \mathbb{F}_\ell^4$ equals ℓ^4. Indeed, w.l.o.g. we can assume $V = \langle (1,0,0,0), (0,1,0,0) \rangle$, and then the trivially intersecting subspaces are the ones of the form $\langle (a,b,1,0), (c,d,0,1) \rangle$ for $a, b, c, d \in \mathbb{F}_\ell$. □