Algorithmic Number Theory Symposium XVII, Groningen, July 6-10, 2026

THE PRINCIPAL IDEAL PROBLEM FOR ENDOMORPHISM RINGS OF SUPERSPECIAL ABELIAN VARIETIES

WOUTER CASTRYCK, JONATHAN KOMADA ERIKSEN, RICCARDO INVERNIZZI, FREDERIK VERCAUTEREN

Abstract. We describe a Las Vegas algorithm for the principal ideal problem in matrix rings $\mathrm{M}_g(\mathcal{O})$ for $g \geq 2$, over maximal orders $\mathcal{O}$ in the rational quaternion algebra $B_{p,\infty}$ ramified at ∞ and a prime number p. Under plausible heuristic assumptions, the method has expected polynomial runtime. An implementation in SageMath shows that it runs very efficiently in practice, with compact output. Our main auxiliary result is a method for finding endomorphisms of superspecial abelian varieties (i.e., powers of supersingular elliptic curves) with a prescribed kernel.

1. Introduction

In this article we study two closely related computational problems, sitting on either side of Deuring's correspondence, which establishes a bijection between supersingular elliptic curves over $\overline{\mathbb{F}}_p$ (up to isomorphism and Galois conjugacy) and maximal orders in the unique rational quaternion algebra $B_{p,\infty}$ ramified precisely at p and ∞ (up to conjugation). First, we consider the principal ideal problem (PIP) in matrix rings over maximal orders in $B_{p,\infty}$:

Problem 1.1. *Let $\mathcal{O} \subset B_{p,\infty}$ be a maximal order. Given a left ideal $I \subset \mathrm{M}_g(\mathcal{O})$, with $g \geq 2$, compute a matrix $\mathbf{M} \in \mathrm{M}_g(\mathcal{O})$ such that $I = \mathrm{M}_g(\mathcal{O})\,\mathbf{M}$.*

Some authors refer to PIP as the "principalization problem". Note that the output is uniquely determined up to left-multiplication by a unit $\mathbf{U} \in \mathrm{GL}_g(\mathcal{O})$.

Recall that, explicitly, the Deuring correspondence maps a supersingular elliptic curve E to the unique conjugacy class of maximal orders $\mathcal{O} \subset B_{p,\infty}$ that admit a ring isomorphism $\iota : \mathcal{O} \to \mathrm{End}(E)$. We can then extend ι component-wise to an isomorphism $\iota_g : \mathrm{M}_g(\mathcal{O}) \to \mathrm{M}_g(\mathrm{End}(E)) = \mathrm{End}(E^g)$. Our second problem is the following kernel-to-matrix problem for superspecial abelian varieties, i.e., powers of supersingular elliptic curves:

Problem 1.2. *Let E be a supersingular elliptic curve over $\overline{\mathbb{F}}_p$ and let $\mathcal{O} \subset B_{p,\infty}$ be a maximal order with an explicit isomorphism $\iota : \mathcal{O} \to \mathrm{End}(E)$. Given a finite subgroup $H \subset E^g$, with $g \geq 2$, compute a matrix $\mathbf{M} \in \mathrm{M}_g(\mathcal{O})$ such that the kernel of $\iota_g(\mathbf{M})$ equals H.*

We will motivate these problems shortly, but let us first discuss why they are well-posed. Famous results by Eichler imply that as soon as a central simple $\mathbb{Q}$-algebra B is either not a quaternion algebra over $\mathbb{Q}$ or unramified at ∞, it necessarily contains a unique maximal order $\mathcal{O} \subset B$ up to conjugation and its ideal class number is 1, i.e., all left (or right) ideals $I \subset \mathcal{O}$ are principal [26, Thm. 34.9]. Both properties are

intimately related. From this perspective $B_{p,\infty}$, which violates the assumptions,[1] is a rather peculiar object within the realm of central simple $\mathbb{Q}$-algebras. Crucially for us, the assumptions do apply to $\mathrm{M}_g(B_{p,\infty})$ as soon as $g \geq 2$. Since it can be argued that $\mathrm{M}_g(\mathcal{O})$ is a maximal order in $\mathrm{M}_g(B_{p,\infty})$ whenever $\mathcal{O}$ is a maximal order in $B_{p,\infty}$ [26, Thm. 8.7], this implies that $\mathrm{M}_g(\mathcal{O})$ is a principal ideal ring, hence Problem 1.1 is indeed well-posed. To draw the same conclusion for Problem 1.2, one can consider the left ideal of matrices $\mathbf{M} \in \mathrm{M}_g(\mathcal{O})$ for which $\iota(\mathbf{M})$ annihilates H; the joint kernel is precisely H (not larger). By Eichler's results, this ideal must be principal, therefore Problem 1.2 is also well-posed. It can be seen as a higher-dimensional analog of the kernel-to-ideal problem for supersingular elliptic curves that is ubiquitous in isogeny-based cryptography.

Remark 1.3. Looking ahead, our discussion will mainly flow in the opposite direction: we will address Problem 1.2 more or less directly, and then convert this into a solution to Problem 1.1. In fact, our line of thought can be used to provide an alternative proof that $\mathrm{M}_g(\mathcal{O})$ is a principal ideal ring.

It is interesting to view Problem 1.2 in the context of a result, usually attributed to Deligne (and sometimes co-attributed to Ogus and/or Shioda), stating that

$$E_1 \times \cdots \times E_g \cong E_1' \times \cdots \times E_g'$$

for any choice of supersingular elliptic curves $E_1, \ldots, E_g, E_1', \ldots, E_g'$ over $\overline{\mathbb{F}}_p$, as soon as $g \geq 2$. This fact can be deduced from Eichler's aforementioned results, see e.g. [14, 22, 29]. Consequently, the bijection $E \mapsto \mathrm{End}(E)$ between isomorphism classes of supersingular elliptic curves up to Galois conjugacy and maximal orders in $B_{p,\infty}$ up to conjugation extends in a trivial way to a bijection $A \mapsto \mathrm{End}(A)$ between g-dimensional superspecial abelian varieties over $\overline{\mathbb{F}}_p$ up to isomorphism (we do not need Galois conjugacy) and maximal orders in $\mathrm{M}_g(B_{p,\infty})$ up to conjugation, for the simple reason that if $g \geq 2$ then both sides contain a single object only. This sheds an alternative light on Problem 1.2: the isogeny $E^g \to E^g/H$ is separable, therefore the codomain is superspecial [5, App. C], hence admits an isomorphism to E^g. So one sees that H must indeed be the kernel of an endomorphism $\iota_g(\mathbf{M}) \in \mathrm{M}_g(\mathrm{End}(E))$ for some $\mathbf{M} \in \mathrm{M}_g(\mathcal{O})$. Thanks to recent work by Gaudry, Soumier and Spaenlehauer [13] this approach can in fact be made effective; see Remark 3.2.

Motivation and prior work. Problem 1.2 (kernel-to-matrix conversion) pops up naturally in the design of cryptographic protocols from superspecial principally polarized (p.p.) abelian varieties where knowledge of the endomorphism ring (including the Rosati involution) is the trapdoor; see [4, §6] and [21]. As far as we are aware, it was first studied by Chu in his PhD thesis [9] in an attempt to design a Galbraith–Petit–Silva style signature scheme from superspecial p.p. abelian surfaces. Chu's strategy was to reduce the problem to Problem 1.1 (PIP), as explained above. He then devised a PIP solver for $g = 2$ with sub-exponential runtime [9, Ch. 2], by mimicking a method due to Page [23] for the principal ideal problem in rational quaternion algebras that do not ramify at ∞ (to which Eichler's results do apply). However, this algorithm is too slow for constructive applications.

In [4, §4.2] it was observed that if $H \cong \mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z}$ for some prime number ℓ then a corresponding matrix $\mathbf{M} \in \mathrm{M}_2(\mathcal{O})$ can be written down essentially for

[1] It can be shown to contain about $p/24$ non-conjugate maximal orders, where each maximal order has about twice as many left ideal classes [32, §42.3.8].

free, without need to pass through PIP. Using recursion, this method extends to groups H admitting a subgroup series $1 = H_0 \subset H_1 \subset \ldots \subset H_{r-1} \subset H_r = H$ all of whose factor groups have the above shape. Since this covers the kernels of arbitrary polarized isogenies, this is good enough for most cryptographic applications. Nevertheless it remains a compelling question whether such a direct method exists for general groups, i.e., not passing via PIP, and whether it generalizes to arbitrary $g \geq 2$. This was the motivating question for the article at hand. We will answer the question in the affirmative, and show that the resulting ideas can be turned into an efficient PIP solver, rather than the other way around.

Remark 1.4. Even when considering kernels H of polarized isogenies only, there may be an efficiency gain in factoring H down to the level of cyclic groups, e.g., when using a list of precomputed matrices as in [4, Rem. 4.3]. Indeed, $E^2[\ell]$ contains $(\ell^2+1)(\ell^3-1)/(\ell-1)$ subgroups of the form $(\mathbb{Z}/\ell\mathbb{Z})^2$ and only $(\ell^4-1)/(\ell-1)$ subgroups of the form $\mathbb{Z}/\ell\mathbb{Z}$.

The literature on the principal ideal problem is more extensive; see Page [23, §1], Chu [9, §2.1.2] and Biasse–Fieker–Hoffman–Youmans [2, §1] for an overview of the existing PIP solvers for maximal orders in various simple algebras (not restricted to $\mathbb{Q}$-algebras), including number fields. Let us highlight the work of Kirschmer and Voight [18] who studied the problem for maximal orders in $B_{p,\infty}$, i.e., the case $g = 1$ of Problem 1.1 (with the promise that the input ideal I is principal), by reducing it to lattice reduction in dimension 4. This technique fails for $g \geq 2$, essentially because the norm form on $\mathrm{M}_g(B_{p,\infty})$ is no longer quadratic.

Of utmost relevance is independent work by Page, Robert and Soumier [24] who, also motivated by cryptography, describe a polynomial-time algorithm for Problem 1.1.[2] Our method is quite different from theirs. Firstly, it has a more geometric flavor which, depending on the reader's taste, may be found more conceptual. Secondly and most importantly, we do not rely on the KLPT algorithm [19] as a subroutine; this contrasts with [24] who invoke a g-fold application of KLPT. Since KLPT is notorious for its long output, our algorithm returns a generator that is considerably more compact; see Theorem 4.2 and Remark 4.4 for a discussion of the case $g = 2$. It also conceivably runs faster in practice. Thirdly, the algorithm from [24] has cleaner heuristic assumptions, only relying on the Generalized Riemann Hypothesis (GRH). A common ingredient in both works is the special role played by so-called extremal maximal orders $\mathcal{O}_0 \subset B_{p,\infty}$, i.e., maximal orders containing an imaginary quadratic order with a very small discriminant (in absolute value). As explained below and also pointed out in [24, §2.1], it is easy to reduce the principal ideal problem in any given maximal order $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ to the case $\Lambda = \mathrm{M}_g(\mathcal{O}_0)$ with $\mathcal{O}_0 \subset B_{p,\infty}$ an extremal maximal order.

Contributions. For matrix rings over extremal maximal orders, we prove:

Theorem 1.5 (PIP, extremal case)**.** *There exists a Las Vegas algorithm which on input a left ideal $I \subset \mathrm{M}_g(\mathcal{O}_0)$ with norm n, where $\mathcal{O}_0 \subset B_{p,\infty}$ is an extremal maximal order, outputs a matrix $\mathbf{M} \in \mathrm{M}_g(\mathcal{O}_0)$ such that $I = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{M}$. Under plausible heuristic assumptions, the algorithm runs in an expected number of*

$$O(\log^{4+\varepsilon} p + g^{3+\varepsilon} \log^{3+\varepsilon}(gnp) + g^{\omega+1+\varepsilon} \log^{1+\varepsilon}(gnp))$$

bit operations, with ω the matrix-multiplication exponent.

[2] We thank the authors of [24] for sharing an early draft of their paper.

Here it is assumed that I is given to us as a list (of constant length) of generators, where the coordinates of the entries (with respect to a fixed $\mathbb{Z}$-basis of $\mathcal{O}_0$, e.g., as specified in (1) below) are reduced mod n; this is a free assumption in any reasonable setting. The plausible heuristic assumptions have a very standard flavour. They essentially assume that the statistical arithmetic behaviour of natural numbers, such as the probability of primality, in certain non-uniform distributions does not deviate too wildly from that in uniform distributions on large intervals in the same range. Our paper is accompanied by a proof of concept implementation of this algorithm in SageMath, available at `https://github.com/KULeuven-COSIC/PIP`.

In order to handle arbitrary maximal orders, a key observation is that the algorithm from Theorem 1.5 can be used to compute a conjugation between $\mathrm{M}_g(\mathcal{O}_0)$ and any given maximal order $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$. For this, one computes the so-called connecting ideal $J_{\Lambda,\mathcal{O}_0} = \exp(G_{\Lambda,\mathcal{O}_0})\,\mathrm{M}_g(\mathcal{O}_0)\Lambda \subset \mathrm{M}_g(\mathcal{O}_0)$ where $\exp(G_{\Lambda,\mathcal{O}_0})$ denotes the exponent of the additive group

$$G_{\Lambda,\mathcal{O}_0} = \frac{\Lambda}{\mathrm{M}_g(\mathcal{O}_0) \cap \Lambda}.$$

It is not hard to check that running the algorithm on input $J_{\Lambda,\mathcal{O}_0}$ returns a matrix $\mathbf{C} \in \mathrm{M}_g(\mathcal{O}_0)$ such that $\Lambda = \mathbf{C}^{-1}\,\mathrm{M}_g(\mathcal{O}_0)\,\mathbf{C}$. This conjugation can then be used to reduce PIP from Λ to $\mathrm{M}_g(\mathcal{O}_0)$: given an ideal $I \subset \Lambda$ of norm n, another run of the algorithm from Theorem 1.5 on the ideal $\mathbf{C}\,I\,\mathbf{C}^{-1}$ (which also has norm n) returns a generator $\mathbf{M}$, and then $\mathbf{C}^{-1}\,\mathbf{M}\,\mathbf{C}$ is a generator of I. This leads to:

Corollary 1.6 (PIP, general case)**.** *Let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be a maximal order and let $\mathcal{O}_0 \subset B_{p,\infty}$ be an extremal maximal order. Let $I \subset \Lambda$ be a left ideal of norm n. Using two runs of the algorithm from Theorem 1.5, once on input $J_{\Lambda,\mathcal{O}_0}$ and once on input an ideal of norm n, we can compute $\mathbf{M} \in \Lambda$ such that $I = \Lambda\,\mathbf{M}$.*

Observe that the norm of $J_{\Lambda,\mathcal{O}_0}$ is an important complexity parameter in this reduction. In practice this quantity is easy to compute using Lemma 2.5 further down.[3] In the special case where $\Lambda = \mathrm{M}_g(\mathcal{O})$ for a non-extremal maximal order $\mathcal{O}$, then the first run of the algorithm from Theorem 1.5 can actually be done on the connecting ideal $J_{\mathrm{M}_2(\mathcal{O}),\mathcal{O}_0}$ between maximal orders in $\mathrm{M}_2(B_{p,\infty})$, rather than $\mathrm{M}_g(B_{p,\infty})$. See Remark 5.7.

Remark 1.7. This reduction also shows that a solution to Problem 1.1 can be used to find the conjugation between any pair of maximal orders in $\mathrm{M}_g(B_{p,\infty})$. These two problems are in fact equivalent. Indeed Problem 1.1 for an ideal I essentially amounts to finding the conjugation between its left and right order.

As mentioned, Theorem 1.5 is inspired by a method for solving Problem 1.2:

Theorem 1.8. *Let E be a supersingular elliptic curve over $\mathbb{F}_{p^2}$ and let $\mathcal{O} \subset B_{p,\infty}$ be a maximal order with an explicit isomorphism $\iota : \mathcal{O} \to \mathrm{End}(E)$. Let $H \subset E^g$,*

[3] It can be bounded by the index $|G_{\Lambda,\mathcal{O}_0}|$ of the Eichler order $\Lambda \cap \mathrm{M}_g(\mathcal{O}_0)$ inside Λ (which equals its index inside $\mathrm{M}_g(\mathcal{O}_0)$), but this is usually an overestimation. A local analysis shows

$$|G_{\Lambda,\mathcal{O}_0}|^{\frac{1}{2g-1}} \leq \mathsf{n}(J_{\Lambda,\mathcal{O}_0}) \leq |G_{\lambda,\mathcal{O}_0}|$$

where one sees that these inequalities become equalities if $g = 1$, reestablishing [12, Prop. 1]. In general, the lower bound is met if the connecting isogeny is cyclic.

with $g \geq 2$, be a finite subgroup. There is an algorithm that solves Problem 1.2 with complexity

$$O(F_H + gD_H + \log|H| \cdot \log^{4+\varepsilon} p + \log^{3+\varepsilon}(|H|mp) + \log^2|H| \cdot g^{\omega+\varepsilon} \log^{1+\varepsilon}(mp))$$

where F_H and D_H denote the complexity of factoring $|H|$ and solving discrete logarithms in H and m is the norm of the connecting ideal between $\mathrm{M}_2(\mathcal{O})$ and $\mathrm{M}_2(\mathcal{O}_0)$ for some extremal maximal order $\mathcal{O}_0 \subset B_{p,\infty}$.

Our repository contains a proof-of-concept implementation of this method, as well.

Roadmap. In Section 2, we gather some background material (explicit localization, quaternionic matrix norms, higher-dimensional Deuring correspondence) that is mostly standard and well-known to specialists, but it contains several explicit formulae that we did not manage to pinpoint in the existing literature. In Section 3, we show that Problem 1.2 reduces to the construction of a certain commutative square of elliptic curve isogenies. Section 4 discusses how to construct such a square on the quaternion side of the Deuring correspondence. Section 5 is devoted to the proof of Theorem 1.5, while Section 6 finalizes the proof of Theorem 1.8.

Acknowledgments. This work is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT – No. 101020788), by the Research Council KU Leuven grant C14/ 24/099, as well as by Cybersecurity Research Flanders with reference number VR20192203. Riccardo Invernizzi is funded by Research Foundation – Flanders (FWO) under a PhD Fellowship fundamental research (project number 1138925N). We thank Aurel Page, Damien Robert and Julien Soumier for sharing with us a preliminary version of their article [24]. We also thank the Isocrypt brainstorm team and the anonymous reviewers for their helpful input.

2. Preliminaries

Even though we did not manage to pinpoint every explicit formula, the vast majority of the claims in this section are standard and well-known to specialists. We refer the reader to the books by Reiner [26], Voight [32], Waterhouse [33] and the references provided in course of the text below for more background.

2.1. Representing $B_{p,\infty}$ and choosing an extremal maximal order $\mathcal{O}_0$.

Let p be an odd prime number. The quaternion algebra $B_{p,\infty}$ ramified at p, ∞ can be constructed as

$$B_{p,\infty} = \left(\frac{-q, -p}{\mathbb{Q}}\right)$$

where q is chosen as small as possible, i.e.,

- $q = 1$ if $p \equiv 3 \bmod 4$,
- $q = 2$ if $p \equiv 5 \bmod 8$,
- q is the smallest prime congruent to 3 mod 4 modulo which $-p$ is a quadratic residue if $p \equiv 1 \bmod 8$; under GRH we can assume $q < 2\log^2 p$ [1].

Elements of $B_{p,\infty}$ are expressed in terms of the usual basis $1, i, j, k$ where $i^2 = -q$, $j^2 = -p$ and $k = ij = -ji$. Recall that $B_{p,\infty}$ comes equipped with a conjugation $\alpha = a_1 + a_2 i + a_3 j + a_4 k \mapsto \bar{\alpha} = a_1 - a_2 i - a_3 j - a_4 k$ as well as $\mathbb{Q}$-valued maps $\mathsf{n} : \alpha \mapsto \alpha\bar{\alpha}, \mathsf{tr} : \alpha \mapsto \alpha + \bar{\alpha}$ that are referred to as the reduced norm and reduced trace, respectively (see Section 2.3 for the origins of the adjective "reduced", that we

will usually omit). Under any choice of $\mathbb{C}$-algebra isomorphism $B_{p,\infty} \otimes_{\mathbb{Q}} \mathbb{C} \cong \mathrm{M}_2(\mathbb{C})$ these notions match with the usual determinant and trace.

An order $\mathcal{O}$ in a finite-dimensional $\mathbb{Q}$-algebra B is a subring that is free of rank $[B : \mathbb{Q}]$ as a $\mathbb{Z}$-module. It is represented by specifying a $\mathbb{Z}$-basis. An order $\mathcal{O} \subset B$ is said to be maximal if it is not contained in a strict superorder. For $B = B_{p,\infty}$ it is convenient to fix an "extremal" maximal order $\mathcal{O}_0$ throughout, which is a maximal order containing a quadratic subring R whose discriminant (in absolute value) is very small. Concretely, we fix:

$$\begin{array}{lccc} \mathcal{O}_0 : & \langle 1, i, \frac{i+j}{2}, \frac{1+k}{2} \rangle_{\mathbb{Z}}, & \langle 1, i, \frac{1+j+k}{2}, \frac{i+2j+k}{4} \rangle_{\mathbb{Z}}, & \langle 1, \frac{1+i}{2}, \frac{ci+k}{q}, \frac{(1+i)(ci+k)}{2q} \rangle_{\mathbb{Z}}, \\ & \cup & \cup & \cup \\ R : & \mathbb{Z}[\sqrt{-1}] & \mathbb{Z}[\sqrt{-2}] & \mathbb{Z}[\frac{1+\sqrt{-q}}{2}] \end{array} \tag{1}$$

for $p \equiv 3 \bmod 4$ and $p \equiv 5, 1 \bmod 8$, respectively, where c denotes the smallest positive solution to $x^2 + p \equiv 0 \bmod q$.[4]

Remark 2.1. We have assumed that p is odd, in order to not pollute the exposition with edge cases, but the methods below are easily adapted to cover $p = 2$ as well; here the most convenient representation is $(-1, -1/\mathbb{Q})$ and, up to conjugation, there is one maximal order only: the (rational) Hurwitz quaternions $\langle 1, i, j, (1 + i + j + k)/2 \rangle_{\mathbb{Z}}$, which are clearly "extremal".

2.2. **Explicit localization.** Being ramified at p, ∞ implies that $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \cong \mathrm{M}_2(\mathbb{Z}_\ell)$ for any maximal order $\mathcal{O} \subset B_{p,\infty}$ and any prime number $\ell \neq p$; in contrast the ring $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_p$ is a domain. We now describe how to compute such an isomorphism. For simplicity we work up to finite ℓ-adic precision, i.e., writing $n = \ell^e$ with $e \geq 1$, we show how to establish an explicit isomorphism

$$\Phi_n : \mathcal{O}/n\mathcal{O} \to \mathrm{M}_2(\mathbb{Z}/n\mathbb{Z}). \tag{2}$$

It is certainly known to specialists how to do this, see e.g. [8, Prop. 4.1]; in fact, methods are available for computing such isomorphisms for more general families of algebras [28]. But we did not manage to locate the explicit formulas below, which we include for the reader's convenience.

Remark 2.2. It is easy to use these formulas for addressing the following seemingly more general targets:

- by the Chinese remainder theorem, one can efficiently combine isomorphisms of the form (2) to handle the case where n is an arbitrary positive integer coprime with p, assuming that the factorization of n is known,
- the same formulas solve the infinite-precision tasks of finding an explicit isomorphism $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \cong \mathrm{M}_2(\mathbb{Z}_\ell)$ or $B_{p,\infty} \otimes_{\mathbb{Z}} \mathbb{C} \cong \mathrm{M}_2(\mathbb{C})$: e.g., for the case $p \equiv 3 \bmod 4$ one picks c_1, c_2 satisfying (3) where one discards the "mod n", and the method goes through mutatis mutandis,
- by working component-wise, we can extend the isomorphism (2) to an isomorphism $\mathrm{M}_g(\mathcal{O})/n\,\mathrm{M}_g(\mathcal{O}) \cong \mathrm{M}_g(\mathcal{O}/n\mathcal{O}) \to \mathrm{M}_{2g}(\mathbb{Z}/n\mathbb{Z})$; in fact, we can even extend this to $\Lambda/n\Lambda$ for any maximal order $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ because, as explained in Section 1, our PIP solver can be used to find an explicit conjugation between Λ and $\mathrm{M}_g(\mathcal{O})$ for any given $\mathcal{O}$.

[4] This follows [19, Lem. 2–4], except for the third case $p \equiv 1 \bmod 8$ where the last basis vector was replaced with j. But this must have been a typo: the resulting $\mathbb{Z}$-module is not an order.

Note that, by the Skolem–Noether theorem, the isomorphism Φ_n is uniquely determined up to base change, i.e., post-composition with conjugation by an element of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$. More generally, any isomorphism $\mathrm{M}_g(\mathcal{O})/n\,\mathrm{M}_g(\mathcal{O}) \to \mathrm{M}_{2g}(\mathbb{Z}/n\mathbb{Z})$ is obtained by applying Φ_n component-wise and composing this with conjugation by an element of $\mathrm{GL}_{2g}(\mathbb{Z}/n\mathbb{Z})$.

2.2.1. *The extremal maximal order $\mathcal{O}_0$.* We focus on computing an isomorphism (2) for the extremal case $\mathcal{O} = \mathcal{O}_0$ introduced in (1) above. We will later explain how to adapt this to any maximal order. It suffices to specify the images under Φ_n of the basis elements of $\mathcal{O}_0$ given in (1), meaning that it is enough to find four matrices $\mathbf{M}_1, \ldots, \mathbf{M}_4 \in \mathrm{M}_2(\mathbb{Z}/n\mathbb{Z})$ satisfying the relations among those elements. Of course $\Phi_n(1) = \mathbf{I}_2 =: \mathbf{M}_1$. By composition of Φ_n with conjugation by an invertible matrix if needed, for $\mathbf{M}_2$ we can make the choices

$$\Phi_n(i) = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \qquad \Phi_n(i) = \begin{pmatrix} 0 & 2 \\ -1 & 0 \end{pmatrix}, \qquad \Phi_n\left(\frac{1+i}{2}\right) = \begin{pmatrix} 0 & \frac{1+q}{4} \\ -1 & 1 \end{pmatrix}$$

in the cases $p \equiv 3 \bmod 4$ resp. $p \equiv 5, 1 \bmod 8$. This can be seen as follows: since $\mathbf{M}_2$ must be linearly independent from $\mathbf{M}_1$, there must exist $v_1 \in (\mathbb{Z}/n\mathbb{Z})^2$ whose reduction mod ℓ is not an eigenvector of $\mathbf{M}_2$ mod ℓ. The desired conjugation amounts to base-change from the standard basis of $(\mathbb{Z}/n\mathbb{Z})^2$ to the basis $v_1, -\mathbf{M}_2\, v_2$. We now proceed as follows:

- When $p \equiv 3 \bmod 4$, the identity
$$i\left(\frac{i+j}{2}\right) = \frac{-1+k}{2} = -1 - \left(\frac{i+j}{2}\right) i$$
implies $\mathbf{M}_2\,\mathbf{M}_3 + \mathbf{M}_3\,\mathbf{M}_2 = -\mathbf{M}_1$ which forces us to choose
$$\Phi_n\left(\frac{i+j}{2}\right) = \mathbf{M}_3 := \begin{pmatrix} c_1 & c_2 \\ c_2 - 1 & -c_1 \end{pmatrix}$$
for $c_1, c_2 \in \mathbb{Z}/n\mathbb{Z}$ satisfying
$$\det(\mathbf{M}_3) = -c_1^2 - c_2^2 + c_2 \equiv \mathsf{n}\left(\frac{i+j}{2}\right) \equiv \frac{p+1}{4} \bmod n. \tag{3}$$
Any such c_1, c_2 will do. As a sanity check, note that the congruence is indeed solvable: modulo ℓ it defines a smooth conic in $\mathbb{P}^2(\mathbb{F}_\ell)$ (one checks that the discriminant of the corresponding ternary quadratic form equals $-p \not\equiv 0 \bmod \ell$), so it has $\ell + 1$ rational points, all of which can be lifted by Hensel's lemma. At least one of these points is affine (this is obvious for $\ell > 2$, while for $\ell = 2$ one checks that the conic has exactly one point at infinity). In practice a solution is found by sampling c_2 at random and checking that what is left for c_1^2 is indeed a square mod n. From the identity
$$\frac{1+k}{2} = 1 + i\left(\frac{i+j}{2}\right)$$
we then conclude
$$\Phi_n\left(\frac{1+k}{2}\right) = \mathbf{M}_4 := \mathbf{M}_1 + \mathbf{M}_2\,\mathbf{M}_3 = \begin{pmatrix} c_2 & -c_1 \\ -c_1 & -c_2 + 1 \end{pmatrix}.$$
One easily checks that the multiplication table of $1, i, \frac{i+j}{2}, \frac{1+k}{2}$ is fully compatible with that of $\mathbf{M}_1, \ldots, \mathbf{M}_4$. Moreover, an explicit calculation shows

that Φ_n has determinant $-4c_1^2 - 4c_2^2 + 4c_2 - 1 \equiv p$ when viewed as a $\mathbb{Z}/n\mathbb{Z}$-linear map. So since $\gcd(p, n) = 1$ it indeed concerns an isomorphism.

- When $p \equiv 5 \bmod 8$, it is more convenient to first determine $\mathbf{\Delta} := \mathbf{M}_3 - \mathbf{M}_4$, which one checks should satisfy $\mathbf{M}_2\,\mathbf{\Delta} + \mathbf{\Delta}\,\mathbf{M}_2 = \mathbf{M}_1 + \mathbf{M}_2$. We find

$$\mathbf{\Delta} = \begin{pmatrix} 1 - c_1 & 2c_2 - 1 \\ c_2 & c_1 \end{pmatrix}$$

for $c_1, c_2 \in \mathbb{Z}/n\mathbb{Z}$ such that $-c_1^2 + c_1 - 2c_2^2 + c_2 \equiv (p+3)/8 \bmod n$. From the relation $\mathbf{M}_3 = \mathbf{\Delta}(\mathbf{M}_2 + 2\,\mathbf{M}_1) - \mathbf{M}_1$ we then get

$$\Phi_n\left(\frac{1+j+k}{2}\right) = \mathbf{M}_3 := \begin{pmatrix} -2c_1 - 2c_2 + 2 & -2c_1 + 4c_2 \\ -c_1 + 2c_2 & 2c_1 + 2c_2 - 1 \end{pmatrix},$$

$$\Phi_n\left(\frac{i+2j+k}{4}\right) = \mathbf{M}_4 := \mathbf{M}_3 - \mathbf{\Delta} = \begin{pmatrix} -c_1 - 2c_2 + 1 & -2c_1 + 2c_2 + 1 \\ -c_1 + c_2 & c_1 + 2c_2 - 1 \end{pmatrix},$$

completing the description of Φ_n, which is seen to be an isomorphism along similar lines as above.

- Finally, when $p \equiv 1 \bmod 8$ we can choose $c_1, c_2 \in \mathbb{Z}/n\mathbb{Z}$ such that

$$-c_1^2 - c_1c_2 - \frac{1+q}{4}c_2^2 - cc_2 \equiv \frac{c^2+p}{q} \bmod n$$

and let

$$\Phi_n\left(\frac{ci+k}{q}\right) = \mathbf{M}_3 := \begin{pmatrix} -c_1 & c_1 + \frac{1+q}{4}c_2 + c \\ c_2 & c_1 \end{pmatrix}.$$

This follows from the requirement $\mathbf{M}_2\,\mathbf{M}_3 + \mathbf{M}_3\,\mathbf{M}_2 = \mathbf{M}_3 - c\,\mathbf{M}_1$. It then readily follows that

$$\Phi_n\left(\frac{(1+i)(ci+k)}{2q}\right) = \mathbf{M}_4 := \mathbf{M}_2\,\mathbf{M}_3 = \begin{pmatrix} \frac{1+q}{4}c_2 & \frac{1+q}{4}c_1 \\ c_1 + c_2 & -\frac{1+q}{4}c_2 - c \end{pmatrix}$$

and a similar reasoning shows that this defines an isomorphism.

2.2.2. *General maximal orders.* To compute Φ_n for a general maximal order $\mathcal{O}$, one can compute the left $\mathcal{O}_0$-ideal $I = \exp(\mathcal{O}/(\mathcal{O}_0 \cap \mathcal{O}))\mathcal{O}_0\mathcal{O} \subset \mathcal{O}_0$, which is a connecting ideal between $\mathcal{O}_0$ and $\mathcal{O}$ in the sense that its right order $\mathcal{O}_R(I) = \{\, x \in B_{p,\infty} \,|\, Ix \subset I \,\}$ equals $\mathcal{O}$. Let $\alpha \in I$ be such that $\gcd(\mathsf{n}(\alpha)/\,\mathsf{n}(I), n) = 1$, where $\mathsf{n}(I) = \gcd(\,\mathsf{n}(\alpha) \,|\, \alpha \in I\,) = [\mathcal{O}_0 : I]^{1/2}$ denotes the norm of I. It is not hard to check that the map

$$\Psi_\alpha : \mathcal{O}/n\mathcal{O} \to \mathcal{O}_0/n\mathcal{O}_0 : x \mapsto \alpha x \alpha^{-1}$$

is an isomorphism of rings, where we point out a subtlety when $\gcd(\mathsf{n}(\alpha), n) > 1$: at first sight, the reader might think that Ψ_α is not well-defined in this case because α is not invertible modulo n. But the full expression $\alpha x \alpha^{-1}$ is genuinely reducible modulo n. Indeed, $\mathcal{O} = \mathcal{O}_R(I)$ implies $\alpha x \in I$, hence there exist $y, z \in \mathcal{O}_0$ such that $\alpha x = y\alpha + z\,\mathsf{n}(I)$ and therefore $\alpha x \alpha^{-1} = y + z(\mathsf{n}(I)/\,\mathsf{n}(\alpha))\bar{\alpha}$. Now it is clear that the right-hand side makes sense modulo n in view of our assumption on $\mathsf{n}(\alpha)$. The desired isomorphism Φ_n is then obtained by composing Ψ_α with the isomorphism described in the previous paragraph.

Remark 2.3. In Lemma 2.8 we state a result, which can be turned into a more geometric method for finding an isomorphism (2), if one is given a maximal order $\mathcal{O} \subset B_{p,\infty}$ and an effective isomorphism $\iota : \mathcal{O} \to \mathrm{End}(E)$, where $E/\overline{\mathbb{F}}_p$ is a supersingular elliptic curve. Although this method is generally less efficient, it offers a more intuitive interpretation of the map Φ_n.

2.3. **Reduced norm and trace of quaternionic matrices.** Recall that every finite-dimensional $\mathbb{Q}$-algebra B comes equipped with norm and trace maps

$$\mathcal{N} : B \to \mathbb{Q} : \alpha \mapsto \det(x \mapsto x\alpha), \qquad \mathcal{T} : B \to \mathbb{Q} : \alpha \mapsto \mathrm{trace}(x \mapsto x\alpha),$$

and for $B = B_{p,\infty}$ we have $\mathcal{N}(\alpha) = \mathsf{n}(\alpha)^2$ and $\mathcal{T}(\alpha) = 2\,\mathsf{tr}(\alpha)$, which is why n and tr are referred to as the reduced norm and trace, respectively. This picture generalizes as follows: for all $g \geq 1$ and any $\mathbf{M} \in \mathrm{M}_g(B_{p,\infty})$ we have

$$\mathrm{charpol}_{x \mapsto x\,\mathbf{M}}(t) = \mathrm{charpol}_{\theta(\mathbf{M})}(t)^{2g}$$

where θ denotes any $\mathbb{C}$-algebra isomorphism $\mathrm{M}_g(B_{p,\infty}) \otimes_{\mathbb{Q}} \mathbb{C} \to \mathrm{M}_{2g}(\mathbb{C})$. For this reason the characteristic polynomial of $\theta(\mathbf{M})$ is called the reduced characteristic polynomial of $\mathbf{M}$. Its coefficients are rational numbers that do not depend on the choice of θ. One sees that $\mathcal{N}(\mathbf{M}) = \mathsf{n}(\mathbf{M})^{2g}$ and $\mathcal{T}(\mathbf{M}) = 2g\,\mathsf{tr}(\mathbf{M})$ where $\mathsf{n}(\mathbf{M})$, resp., $\mathsf{tr}(\mathbf{M})$ denote the usual determinant, resp., trace of $\theta(\mathbf{M})$, called the reduced norm, resp., the reduced trace of $\mathbf{M}$.[5] Just like in the case $g = 1$, we will usually drop the adjective "reduced", which should cause no confusion.

Lemma 2.4. *For all $\mathbf{M} \in \mathrm{M}_g(B_{p,\infty})$ and all orders $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ we have:*

1. *$\mathbf{M} \in \mathrm{M}_g(B_{p,\infty})^\times$ if and only if $\mathsf{n}(\mathbf{M}) \neq 0$;*
2. *if $\mathbf{M} \in \Lambda$ then $\mathsf{n}(\mathbf{M}) \in \mathbb{Z}$ and if moreover $\mathsf{n}(\mathbf{M}) \neq 0$ then $\mathsf{n}(\mathbf{M})\,\mathbf{M}^{-1} \in \Lambda$; in particular $\mathbf{M} \in \Lambda^\times$ if and only if $\mathsf{n}(\mathbf{M}) = 1$.*

Proof. The proof is not hard; details can be found in our online version [6]. □

Note that, as a consequence to Lemma 2.4(2), one obtains that $\mathsf{n}(\mathbf{M})\,\mathbf{I}_g$ is an element of the (left or right) principal ideal generated by $\mathbf{M}$.

Lemma 2.5. *Let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be an order and let $\mathbf{M} \in \Lambda$ have non-zero norm. Then $\mathsf{n}(\mathbf{M})^{2g} = [\Lambda : \Lambda\,\mathbf{M}]$.*

Proof. This is a standard argument; a proof for $\mathrm{M}_g(B_{p,\infty})$ can be found in [6]. □

Corollary 2.6. *Let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be a maximal order and let $\mathbf{M} \in \Lambda$ have prime norm $\mathsf{n}(\mathbf{M}) = \ell$ different from p. Consider an isomorphism $\theta : \Lambda/\ell\Lambda \to \mathrm{M}_{2g}(\mathbb{F}_\ell)$. Then $\mathrm{rank}(\theta(\mathbf{M})) = 2g - 1$.*

Proof. Again, the details are deferred to [6]. □

2.4. **Deuring correspondence.** As mentioned in the introduction, for any $g \geq 1$ the map $A \mapsto \mathrm{End}(A)$ defines a bijection between, on one hand, g-dimensional superspecial abelian varieties over $\overline{\mathbb{F}}_p$ up to isomorphism and Galois conjugacy and, on the other hand, maximal orders in $\mathrm{M}_g(B_{p,\infty})$ up to conjugation. We now recall how to incorporate isogenies in this correspondence. Let A be a g-dimensional superspecial abelian variety and let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be a maximal order admitting

[5] Alternatively, the reduced norm $\mathsf{n}(\mathbf{M})$ can be characterized as the norm of the Dieudonné determinant of $\mathbf{M}$, which implies that it takes non-negative values only (a fact that is also apparent from Lemma 2.9).

an isomorphism $\iota_g : \Lambda \to \mathrm{End}(A)$. Typically, we will work with $A = E^g$ for a supersingular elliptic curve $E/\overline{\mathbb{F}}_p$ and $\Lambda = \mathrm{M}_g(\mathcal{O})$ where $\mathcal{O} \subset B_{p,\infty}$ is a maximal order admitting an isomorphism $\iota : \mathcal{O} \to \mathrm{End}(E)$, that we extend component-wise to an isomorphism $\mathrm{M}_g(\mathcal{O}) \to \mathrm{M}_g(\mathrm{End}(E)) = \mathrm{End}(E^g)$. To every finite subgroup $H \subset A$ we can associate a left ideal

$$I_H = \{\, \mathbf{M} \in \Lambda \,|\, \forall P \in H : \iota_g(\mathbf{M})(P) = 0 \,\}$$

whose norm $\mathsf{n}(I_H) = \gcd(\, \mathsf{n}(\mathbf{M}) \,|\, \mathbf{M} \in I_H \,) = [\Lambda : I_H]^{1/2g}$ equals $|H|$. Conversely, given a left ideal $I \subset \Lambda$ one can consider the joint kernel $H_I \subset A$ of the endomorphisms $\iota_g(\mathbf{M})$ over all $\mathbf{M} \in I$. This defines a bijection between left ideals $I \subset \Lambda$ of norm coprime with p and finite subgroups $H \subset A$. The right order

$$\Lambda' = \mathcal{O}_R(I_H) = \{\, x \in \mathrm{M}_g(B_{p,\infty}) \,|\, I_H x \subset I_H \,\}$$

of I_H is naturally isomorphic to $\mathrm{End}(A')$, where $A' = A/H$ denotes the codomain (uniquely determined up to $\overline{\mathbb{F}}_p$-isomorphism) of a separable isogeny $\varphi_H : A \to A'$ with kernel H. Explicitly, following [33, §3.1], let $\psi : A' \to A$ be an isogeny with $\psi \circ \varphi_H = |H|\,\mathrm{id}_A$ and $\varphi_H \circ \psi = |H|\,\mathrm{id}_{A'}$ (we will refer to this as the pseudo-dual[6] of φ_H). The isomorphism is

$$\mathrm{End}(A') \to \Lambda' : \alpha \mapsto |H|^{-1} \iota_g^{-1}(\psi \circ \alpha \circ \varphi_H).$$

We remark that this correspondence between ideals and subgroups can in fact be extended to include ideals of norm a multiple of p (including 0), by considering arbitrary subgroup schemes instead of subgroups, and by allowing for arbitrary homomorphisms (possibly non-finite or inseparable).

Composition of isogenies corresponds to multiplication of ideals, as follows. For a finite subgroup $H' \subset A'$ one can consider the corresponding left ideal $I_{H'} \subset \Lambda'$ and isogeny $\varphi_{H'} : A' \to A'' = A'/H'$. The ideal associated with the kernel $\varphi_H^{-1}(H')$ of the composition isogeny $\varphi_{H'} \circ \varphi_H$ is given by $I_H \cdot I_{H'}$.

Lemma 2.7 (pseudo-dual at ideal level)**.** *Let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be a maximal order and let $I \subset \Lambda$ be a left ideal with right order Λ'. Then*

$$I^{\mathrm{ps}} = \{\, x \in B_{p,\infty} \,|\, Ix \subset \mathsf{n}(I)\Lambda \,\}$$

is a left Λ'-ideal with right order Λ, satisfying $I \cdot I^{\mathrm{ps}} = \mathsf{n}(I)\Lambda$ and $I^{\mathrm{ps}} \cdot I = \mathsf{n}(I)\Lambda'$. In particular $\mathsf{n}(I^{\mathrm{ps}}) = \mathsf{n}(I)^{2g-1}$.

Proof. If $g = 1$ then $I^{\mathrm{ps}} = \bar{I}$ by [32, Prop. 16.6.15] from which the statement follows immediately. If $g \geq 2$ then the statement is easy to prove by writing $I = \Lambda\,\mathbf{M}$ for some generator $\mathbf{M}$. One then readily verifies that $I^{\mathrm{ps}} = \mathsf{n}(\mathbf{M})\,\mathbf{M}^{-1}\,\Lambda$. □

The Deuring correspondence provides an important interpretation of the localization isomorphisms Φ_n discussed in Section 2.2.

Lemma 2.8. *Let A be a g-dimensional superspecial abelian variety over $\overline{\mathbb{F}}_p$ and let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be a maximal order together with an isomorphism $\iota_g : \Lambda \to \mathrm{End}(A)$. Let n be a positive integer coprime with p. Let $P_1, \ldots, P_{2g}$ be a basis of $A[n]$. Then*

$$\Lambda/n\Lambda \to \mathrm{M}_{2g}(\mathbb{Z}/n\mathbb{Z}) : \mathbf{M} \mapsto \textit{matrix of } \iota_g(\mathbf{M}) \textit{ with respect to } P_1, \ldots, P_{2g}$$

is an isomorphism of rings.

[6]Warning: if $g \geq 1$ then in general the pseudo-dual of the pseudo-dual of φ_H is not equal to φ_H itself: rather, it equals $|H|^{2g-2}\varphi_H$. Note in particular that $\deg \psi = |H|^{2g-1}$.

Proof. It is clear that this is a homomorphism between equicardinal rings, so it suffices to prove injectivity. This follows from the fact that if an endomorphism annihilates $A[n]$ it must factor through multiplication-by-n, e.g., by a componentwise application of [30, Cor. III.4.11]. □

Lemma 2.9. *Let A be a g-dimensional superspecial abelian variety over $\overline{\mathbb{F}}_p$ and let $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$ be a maximal order together with an isomorphism $\iota_g : \Lambda \to \mathrm{End}(A)$. Let $\mathbf{M} \in \Lambda$ have non-zero norm. Then $\mathsf{n}(\mathbf{M}) = \deg(\iota_g(\mathbf{M}))$.*

Proof. This can be found in [20, pp. 46-48]. Note: if $\gcd(\mathsf{n}(\mathbf{M}), p) = 1$ then this also follows from Lemma 2.8 and Corollary 2.6 (or rather the generalization of this corollary to $\mathsf{n}(\mathbf{M}) = \ell^e$ for arbitrary $e \geq 1$, which can be proven along the same lines). The case $g = 2$ can also be found in [17, Cor. 64]; see also [13, Prop. 3.9]. □

3. The main idea

We now describe the main idea underlying our algorithms, which is geometric in nature. We focus on kernel-to-matrix conversion, where this geometric intuition is the clearest. Concretely, the goal of this section is to explain why Problem 1.2 reduces to the construction of a commutative square

$$\begin{array}{ccc} E/\langle P\rangle & \longrightarrow & E \\ \downarrow & & \downarrow \\ E & \longrightarrow & E \end{array}$$

of elliptic curve isogenies, with $P \in E$ a given point of prime order, where the degrees of the horizontal isogenies are equal, the degrees of the vertical isogenies are equal, and these two degrees are coprime. In Section 4 we will then show that this square can be built efficiently in the special case where $\mathcal{O} = \mathcal{O}_0$ is an extremal maximal order (which is good enough).

3.1. Reduction to the case $|H| = \ell$.

We first reduce to the case where $|H|$ is a prime. One factors

$$|H| = \prod_{i=1}^{r} \ell_i$$

into prime numbers $\ell_i \neq p$, not necessarily distinct; in general, this step takes superpolynomial (but subexponential) time in $\log |H|$. Letting $H_0 = H$, one can then repeat the following steps for $i = 1, \ldots, r$:

(1) Pick any order-ℓ_i subgroup $K_i \subset H_{i-1}$ and solve Problem 1.2 for K_i, yielding a matrix $\mathbf{M}_i$.
(2) Let $H_i = \iota_g(\mathbf{M}_i)H_{i-1}$.

It is easy to check that the product $\mathbf{M} = \mathbf{M}_r \cdots \mathbf{M}_2 \, \mathbf{M}_1$ is of the desired kind.

Remark 3.1. As alluded to in the introduction, if H has a subgroup series with factors isomorphic to $(\mathbb{Z}/\ell\mathbb{Z})^g$, then it is advantageous not to reduce all the way to the case where $|H|$ is prime, because for groups of this form the kernel-to-matrix problem has a much easier solution. For $g = 2$ this can be found in [4, §4.2]; that method easily generalizes to arbitrary $g \geq 2$ using a trick akin to Section 3.3 below.

3.2. **Reduction to the case $H = \langle (P, 0, \ldots, 0) \rangle$.** We may now assume that H is cyclic of prime order ℓ. Hence, $H = \langle (P_1, P_2, \ldots, P_g) \rangle$ where $P_i \in E[\ell]$ for all $i = 1, \ldots, g$ and at least one point, say P_j, has exact order ℓ. Using Lemma 2.8 it is not hard to find $b_1, \ldots, b_g \in \mathcal{O}$ such that $P_i = \iota(b_i)(P_j)$ for all $i = 1, \ldots, g$, where we can (and do) take $b_j = 1$ (more details on how to find the b_i's are given in Section 6.1). Then the matrix

$$\mathbf{B} = \begin{pmatrix} -1 & 0 & 0 & \ldots & b_1 & \ldots & 0 \\ 0 & -1 & 0 & \ldots & b_2 & \ldots & 0 \\ 0 & 0 & -1 & \ldots & b_3 & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & b_j & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & b_g & \ldots & -1 \end{pmatrix} \tag{4}$$

is an element of $\mathrm{GL}_g(\mathcal{O})$ and maps H to $\langle (0, \ldots, 0, P_j, 0, \ldots, 0) \rangle$. If $\mathbf{\Pi} \in \{0,1\}^{g \times g}$ is a permutation matrix swapping entries 1 and j, then also $\mathbf{\Pi}\,\mathbf{B} \in \mathrm{GL}_g(\mathcal{O})$ and this now maps H to $H_1 = \langle (P_j, 0, \ldots, 0) \rangle$. If $\mathbf{M}$ is a valid solution to Problem 1.2 upon input H_1, then $\mathbf{M}\,\mathbf{\Pi}\,\mathbf{B}$ is a valid solution for input H.

3.3. **Reduction to the case $g = 2$.** It is straightforward to check that, if $\mathbf{M}$ is a valid solution to Problem 1.2 upon input $H = \langle (P, 0) \rangle \subset E^2$, then

$$\begin{pmatrix} \mathbf{M} & 0 \\ 0 & \mathbf{I}_{g-2} \end{pmatrix}$$

solves Problem 1.2 upon input $H = \langle (P, 0, \ldots, 0) \rangle \subset E^g$.

3.4. **Reduction to the construction of a commutative square.** As in Section 2.4 we consider the left ideal $I_P := I_{\langle P \rangle} = \{\, a \in \mathcal{O} \,|\, \iota(a)(P) = 0 \,\}$ with norm $\mathsf{n}(I_P) = \gcd(\, \mathsf{n}(a) \,|\, a \in I_P \,) = \ell$. Note that if $a, c \in I_P$ then the matrix

$$\mathbf{M} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

annihilates $H = \langle (P, 0) \rangle$, whatever b and d are. Therefore, if we succeed in finding a, c admitting b, d such that $\mathsf{n}(\mathbf{M}) = \ell$, then the kernel cannot be larger, meaning that $\mathbf{M}$ is of the desired form.

The key ingredient is the following expression for the reduced norm of $\mathbf{M}$:

$$\mathsf{n}(c)\,\mathsf{n}(\mathbf{M}) = \mathsf{n}(\mathsf{n}(c)b - a\bar{c}d). \tag{5}$$

This formula is taken from [4, §3.2] and is easy to verify.[7] Assuming $\mathsf{n}(c) \neq 0$, this shows that suitable b, d can be found if and only if the left ideal

$$J_{a,c} = \mathcal{O}\,\mathsf{n}(c) + \mathcal{O}c\bar{a},$$

[7] This can be done via explicit calculation. The reader is also invited to compare this with the high-school formula

$$\det \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \det \begin{pmatrix} C & D \\ A & B \end{pmatrix} = \det(C) \det(B - AC^{-1}D)$$

for the determinant of a 2×2 block matrix in which the blocks have even dimension (this is used in the first equality) and assuming $\det(C) \neq 0$. We finally note that this formula lies very close to the interpretation of $\mathsf{n}(\mathbf{M})$ as the norm of the Dieudonné determinant of $\mathbf{M}$.

contains an element of norm $\mathsf{n}(c)\ell$. Indeed, the conjugate of this element has norm $\mathsf{n}(c)\ell$ as well, so upon writing the element as $x\,\mathsf{n}(c)+yc\bar{a}$, one can let $b=\bar{x}, d=-\bar{y}$. But for $J_{a,c}$ to admit an element of norm $\mathsf{n}(c)\ell$ is a strong requirement because then we have $\mathsf{n}(c)\ell \mid \mathsf{n}(J_{a,c}) \mid \mathsf{n}(c)\ell$. The first divisibility can be seen as follows: as b and d vary, the right-hand side of Equation (5) runs over the norms of all possible elements of $J_{a,c}$, while the left-hand side is always divisible by $\mathsf{n}(c)\ell$, no matter what b and d are, because $\iota_2(\mathbf{M})$ always annihilates the order-ℓ point $(P,0)$. Consequently, the two divisibilities are equalities. From this one sees that $J_{a,c}$ contains an element of norm $\mathsf{n}(c)\ell$ if and only if it is a principal ideal.

Luckily, it is indeed algorithmically feasible to find $a, c \in I$ such that $J_{a,c}$ is principal. Let τ be the isogeny $\tau : E \to F = E/\langle P\rangle$ corresponding to I_P, i.e., the isogeny with $\ker(\tau) = \langle P\rangle$. Assume we can build the following commutative square

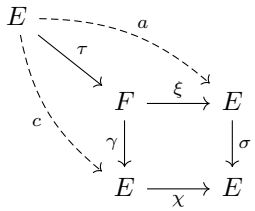

of isogenies, where $\deg\xi = \deg\chi$ and $\deg\gamma = \deg\sigma$ are coprime. Note that apart from the curve F in the upper-left corner, the square only involves the base curve E; in particular σ, χ are endomorphisms. How to construct such a square is explained in Section 4 for the special case where $\mathcal{O} = \mathcal{O}_0$ is extremal. As explained in Section 1, this special case is sufficient to cover the general case.

Let us now assume that we have constructed the square above; we can proceed as follows. We let a, c be such that

$$\iota(a) = \xi \circ \tau, \qquad \iota(c) = \gamma \circ \tau,$$

which are indeed endomorphisms annihilating H, so they are elements of I_P. One sees that the left ideal generated by

$$\mathsf{n}(c) = \ell \cdot \deg\gamma = \ell \cdot \deg\sigma = \ell \cdot \mathsf{n}(\sigma) = \bar{\sigma}\cdot\ell\sigma,$$
$$c\bar{a} = \ell \cdot \iota^{-1}(\gamma \circ \hat{\xi}) = \ell \cdot \iota^{-1}(\hat{\chi}\circ\sigma) = \bar{\chi}\cdot\ell\sigma$$

just equals the principal ideal generated by $\ell\sigma$, as wanted; here we used that $\deg\sigma$ and $\deg\chi$ are coprime, and we have abused notation and identified σ, χ with their preimages under ι. In conclusion, the quaternions x, y can be found by expressing

$$1 = x\cdot\bar{\sigma} + y\cdot\bar{\chi}$$

and we recall that one can then take $b=\bar{x}, d=-\bar{y}$.

Remark 3.2. The fact that we are reduced to computing such a commutative square, can also be conceptually explained by the fact that such a square induces an isomorphism Ψ between the products of the opposite corners of the square [3, Thm. A.1], in this case $\Psi : F \times E \to E \times E$. Thus, we could also construct our matrix $\mathbf{M}$ from the composition

$$\Psi \circ \begin{pmatrix} \tau & 0 \\ 0 & \mathrm{id}_E \end{pmatrix}.$$

Doing so yields a matrix $\mathbf{M}$ of a slightly different form; see [6, §4.5]. We note that a general algorithm with expected polynomial runtime for computing isomorphisms

between products of supersingular elliptic curves ("effective Deligne") was recently devised by Gaudry, Soumier and Spaenlehauer [13].

4. Building the commutative square

We now describe an elementary method for building this commutative square, switching to the quaternion side of the Deuring correspondence. Throughout this section, we will assume that we work with an extremal maximal order $\mathcal{O} = \mathcal{O}_0$, containing a quadratic subring R with small discriminant in absolute value. In practice, one can work with $\mathcal{O}_0 \supset R$ as in (1). Let us formalize the following problem, which was stated for $I = I_P$ at the beginning of Section 3.4:

Problem 4.1. *Let $\mathcal{O}_0$ be an extremal maximal order in $B_{p,\infty}$, and I a left $\mathcal{O}_0$-ideal of prime norm $\mathsf{n}(I) = \ell \neq p$. Find a matrix $\mathbf{M} \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\mathsf{n}(\mathbf{M}) = \ell$ and each entry in the first column $(a\ c)^t$ of $\mathbf{M}$ is an element of I.*

Then, concretely, this section is devoted to a proof of the following theorem. Here it is assumed that the input ideal I is provided to us as a list (of constant length) of generators, e.g., $I = \mathcal{O}_0\ell + \mathcal{O}_0\alpha$ for some $\alpha \in \mathcal{O}_0$.

Theorem 4.2. *There is a Las Vegas algorithm for solving Problem 4.1 which runs in expected time $O(\log^{4+\varepsilon} p + \log^3 B)$, with B the maximal norm of the given generators of I, and which under plausible heuristic assumptions outputs a matrix $\mathbf{M}$ whose entries are bounded, norm-wise, as*

$$\begin{pmatrix} O(\ell n) & O(np^{1/2}) \\ O(\ell p^{1+\varepsilon}) & O(p^{3/2+\varepsilon}) \end{pmatrix} \tag{6}$$

with $n = O(\max\{p^{1/2+\varepsilon}, p^{1+\varepsilon}/\ell\})$. If $\ell \leq p^{1/2}$ and ℓ is contained in the image of the norm map on R then this can be improved by taking $n = \ell$ in (6).

Our solution will use several standard operations on quaternion ideals as subroutines; a good account of such operations can be found in [7, §2.4]. We invite the reader to keep the pseudo-code from Algorithm 1 at hand, where for notational simplicity it is assumed that $p \equiv 3 \bmod 4$ and $\mathcal{O}_0 \supset R = \mathbb{Z}[\sqrt{-1}]$ is as in (1).

Remark 4.3. The heuristic assumptions are all related to the statistical arithmetic behaviour of norms in lattices. For example, it is assumed that the probability of primality of $\mathsf{n}(a)/\ell$ for a random element $a \in \{\, x \in I \mid \mathsf{n}(x) \leq X \,\}$ lies within a constant factor of the probability of primality of a random integer sampled from $[1, X/\ell]$. We note that small biases as described in [27] are not harmful.

If we rewrite the commutative square from Section 3.4 in terms of ideals (i.e., replacing curves with maximal orders and isogenies with ideals) we obtain

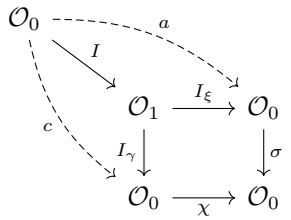

Here χ, σ, a and c are principal ideals, which we denote with their generators by abuse of notation. The degree requirements become $\mathsf{n}(I_\xi) = \mathsf{n}(\chi)$, $\mathsf{n}(I_\gamma) = \mathsf{n}(\sigma)$,

$\gcd(\mathsf{n}(\sigma), \mathsf{n}(\chi)) = 1$. Some care is needed when viewing this as a "commutative" diagram: e.g., it is obvious that the right order of χ will in general not equal $\mathcal{O}_0$, but rather its conjugate $\chi^{-1}\mathcal{O}_0\chi$. The construction below will ensure that I_ξ and I_γ have the same left order $\mathcal{O}_1$ (more precisely I_γ will arise as $[I_\xi]^*\sigma$, where $*$ denotes the pull-back as discussed in [7, §2.4.5]). But this will not necessarily equal the right order of I.

4.1. **Horizontal arrows.** The first step is finding the ideals I_ξ and χ. We let

$$I_\xi = \frac{a}{\ell}\bar{I}$$

for $a \in I$ of norm ℓn such that we can efficiently find an element in $\chi \in \mathcal{O}_0$ of norm $n = \mathsf{n}(I_\xi)$. This ideal is right-equivalent to $\bar{I}$ (or equivalently, $\bar{I}_\xi$ is left-equivalent to I), so if under the Deuring correspondence I corresponds to a degree-ℓ isogeny $\tau : E \to F$ then I_ξ corresponds a degree-n isogeny $\xi : F \to E$ with composition $\xi \circ \tau = a$. We let $\mathcal{O}_1$ be the left order of I_ξ. If possible, it is allowed to take $I_\xi = \bar{I}$ so that $n = \ell$ and $\mathcal{O}_1 = \mathcal{O}_R(I)$. In general, we sample $a \in I$ at random until n is prime and we can find a corresponding χ in the subring $R \subset \mathcal{O}_0$, where norm equations can be solved efficiently [11]. In general, a prime appears as a norm in R with probability about $1/2h$ where h is the class number of R. Notice that $h(\mathbb{Z}[\sqrt{-1}]) = h(\mathbb{Z}[\sqrt{-2}]) = 1$ and $h(\mathbb{Z}[(1+\sqrt{-q})/2]) = O(\log^{1+\varepsilon} p)$ assuming $q = O(\log^2 p)$ under GRH. If I_ξ is sampled from the right-equivalent ideals of $\bar{I}$, we expect $n = \mathsf{n}(I_\xi) = O(\sqrt{p})$ so the probability of n being a prime is roughly $2/\log p$. Thus we expect to succeed after $O(h \log p)$ attempts. These steps are detailed in lines 1-3 of Algorithm 1.

4.2. **Vertical arrows.** Once χ is fixed, σ is determined by $\bar{\chi}\sigma \in \bar{I}_\xi$. We employ the isomorphism $\Phi_n : \mathcal{O}_0/n\mathcal{O}_0 \mapsto \mathrm{M}_2(\mathbb{Z}/n\mathbb{Z})$ described in Section 2.2. Note that since n is prime, computing this isomorphism can be done efficiently. Write $\bar{I}_\xi = \mathcal{O}_0 x + \mathcal{O}_0 n$ for some $x \in \mathcal{O}_0$. Let $\mathbf{M}_{\hat{\xi}} = \Phi_n(x), \mathbf{M}_{\bar{\chi}} = \Phi_n(\bar{\chi})$. Both $\mathbf{M}_{\hat{\xi}}$ and $\mathbf{M}_{\bar{\chi}}$ will have rank one. The element we are after is an invertible matrix $\mathbf{M}_\sigma$ mapping $\ker(\mathbf{M}_{\hat{\xi}})$ to $\ker(\mathbf{M}_{\bar{\chi}})$; any representant σ of $\Phi_n^{-1}(\mathbf{M}_\sigma)$ is then an endomorphism mapping $\ker(\hat{\xi})$ to $\ker(\bar{\chi})$, and the invertibility of $\mathbf{M}_\sigma$ ensures that $\mathsf{n}(\sigma)$ is coprime with n. Matrices of the form $\mathbf{M}_\sigma$ can be sampled using linear algebra and will be invertible with high probability; to achieve the bounds (6) we will need an improved, lattice-based method discussed in Section 4.4. The left $\mathcal{O}_1$-ideal I_γ is then computed as $[I_\xi]^*\sigma$, which equals $I_\xi\sigma + \mathcal{O}_1\,\mathsf{n}(\sigma)$ by [7, §2.4.5.1]. See lines 4-8 of Algorithm 1.

4.3. **Recovering the matrix M.** We now perform the steps outlined in Section 3.4 in the "ideal world" to recover the entries a, b, c, d of $\mathbf{M}$. Recall that we already fixed a during the generation of I_ξ. To find $c = \gamma \circ \tau$, note that we cannot just compute $I \cdot I_\gamma$ because the respective right and left orders do not match. As shown in Section 3.4 we have $c\bar{a} = \ell\gamma \circ \hat{\xi}$ so it is more convenient to compute $c\bar{a}$ as a generator of the principal ideal $\ell\bar{I}_\xi I_\gamma$, where now the product glues nicely at the maximal order $\mathcal{O}_1$. We then extract c from a and $c\bar{a}$. We can now conclude as in Section 3.4 by computing x, y such that $1 = x \cdot \bar{\sigma} + y \cdot \bar{\chi}$ and letting $b = \bar{x}$, $d = -\bar{y}$. Note that these coefficients are not unique and in Section 4.4 we will explain how to make an optimized choice using lattices. We also note the following subtlety: our push-forward construction ensures that

$$\gamma \circ \hat{\xi},\ \chi \circ \sigma\ :\ E \to E$$

have the same kernel, but this only means that they coincide up to post-composition with an automorphism. Concretely for us, this implies that we might still have to left-multiply χ with a unit in $\mathcal{O}_0$. We can find the corrected version of χ by imposing the identity $c\bar{a} = \ell\bar{\chi}\sigma$. These steps are discussed in lines 9-12 of Algorithm 1.

4.4. **Norm analysis of the entries of M.** It is standard that the ideal I can be viewed as a lattice with Euclidean covolume $\mathsf{n}(I)^2 p/4 = \ell^2 p/4$; under this identification, the quaternion norm matches with the square of the Euclidean norm. Assuming the Gaussian heuristic, the expected number of elements $a \in I$ such that $\mathsf{n}(a)$ is smaller than a given bound X equals $2\pi^2 X^2/\ell^2 p$. We need in the order of $h \log p = \log^{2+\varepsilon} p$ attempts, so in general we expect that $\mathsf{n}(a) = O(\ell p^{1/2+\varepsilon})$ and consequently $n = \mathsf{n}(\chi) = O(p^{1/2+\varepsilon})$ will do. However, if ℓ is small then the lattice is skew because $\mathcal{O}_0$ is extremal: if we treat q as a constant, then in the range $\ell = O(p^{1/2})$ we expect the successive minima (with respect to the quaternion norm) to be in the orders of ℓ^2, ℓ^2, p, p, rather than $\ell p^{1/2}, \ell p^{1/2}, \ell p^{1/2}, \ell p^{1/2}$, forcing us to resort to an ellipsoidal Gaussian heuristic; the standard spherical Gaussian heuristic fails. Furthermore, with probability about $1/2h$, the prime ℓ itself is the norm of an element of R, in which case the potentially much smaller choice $a = \ell$ will do. Taking into account that q may in fact grow with p, but at most as $O(\log^2 p)$ under GRH, these considerations lead to estimating n as in the statement of Theorem 4.2.

We now discuss our improved method for Line 6 of Algorithm 1. The matrices $\mathbf{M}_\sigma$ mapping $\ker(\mathbf{M}_{\bar{\xi}})$ to $\ker(\mathbf{M}_{\bar{\chi}})$ form an additive subgroup of $\mathrm{M}_2(\mathbb{Z}/n\mathbb{Z})$ of index n, whose pre-image under Φ_n is a sublattice of $\mathcal{O}_0$ also of index n, hence of covolume $np/4$. But it concerns a skew sublattice, since it contains the element χ whose quaternion norm is n. Given that $\mathcal{O}_0$ is extremal, we therefore expect successive minima in the orders of n, n, p, p, roughly. So we expect being able to choose $\deg(\gamma) = \mathsf{n}(\sigma) = O(p^{1+\varepsilon})$ and consequently $\mathsf{n}(c) = O(\ell p^{1+\varepsilon})$.

A similar improvement applies to Line 11. To merely find x, y satisfying $1 = x\bar{\sigma} + y\bar{\chi}$, one can simply take $x = n_2\sigma$, $y = n_1\chi$ where $n_1, n_2 \in \mathbb{Z}$ are Bézout coefficients expressing 1 as a linear combination of the coprime integers $n = \mathsf{n}(\chi), \mathsf{n}(\sigma)$. But one can shorten x by subtracting from it a nearby element of the lattice

$$\{\, z \in \mathcal{O}_0 \mid z\bar{\sigma} \in \mathcal{O}_0\bar{\chi} \,\}$$

which has covolume $\mathsf{n}(\bar{\chi})^2 p/4 = n^2 p/4$. The expected Euclidean covering radius of such a lattice is $O(n^{1/2} p^{1/4})$, so we expect being able to realize $\mathsf{n}(x) = O(np^{1/2})$ and then from $y\bar{\chi} = 1 - x\bar{\sigma}$ we expect $\mathsf{n}(y) = O(p^{3/2+\varepsilon})$. We note that reasoning from y rather than from x results in the same estimates. Altogether this leads to the estimates (6).

Remark 4.4. These estimates contrast with the method from [24], which due to its reliance on the KLPT algorithm appears to produce a matrix with entries in the order of $\ell p^{7/2+\varepsilon}$ (or $\ell p^{3+\varepsilon}$ with the variant of [25]).

4.5. **Complexity.** We conclude the proof of Theorem 4.2 with a complexity analysis. For simplicity, we will not make explicit reference to the parameter q introduced in Section 2.1: since $q < 2\log^2 p$ under GRH, this is backed up by the ε's in the exponents. The most expensive operations are:

- finding χ; as argued in Section 4.1, we need $O(h \log p)$ primality tests in Line 1, where $h = O(\log^{1+\varepsilon} p)$ is the class number of R. Note that if $p \not\equiv$

Algorithm 1: Solving the PIP instance from Problem 4.1

Input: Left $\mathcal{O}_0$ ideal I of norm $\ell \neq p$, with $p \equiv 3 \bmod 4$ and $\mathcal{O}_0$ as in (1)
Output: Matrix $\mathbf{M} \in \mathrm{M}_2(\mathcal{O}_0)$ such that $\mathsf{n}(\mathbf{M}) = \ell$ and each entry in the first column of $\mathbf{M}$ is an element of I

1 Find $a \in I$ of small, prime norm $n \equiv 1 \bmod 4$
2 $I_\xi \leftarrow a\bar{I}/\ell$
3 Set $\chi \leftarrow \ell_1 + i\ell_2$ where $\ell_1^2 + \ell_2^2 = n$ // $\mathsf{n}(\chi) = n$
4 Find x such that $\bar{I}_\xi = \mathcal{O}_0 x + \mathcal{O}_0 n$
5 Let $\mathbf{M}_{\hat{\xi}} \leftarrow \Phi_n(x)$ and $\mathbf{M}_{\bar{\chi}} \leftarrow \Phi_n(\bar{\chi})$ // see Section 2.2
6 Compute $\mathbf{M}_\sigma$ mapping $\ker(\mathbf{M}_{\hat{\xi}})$ to $\ker(\mathbf{M}_{\bar{\chi}})$ // see Section 4.4
7 Let σ be any representant of $\Phi_n^{-1}(\mathbf{M}_\sigma)$
8 $I_\gamma \leftarrow I_\xi \sigma + \mathcal{O}_1\, \mathsf{n}(\sigma)$
9 $c\bar{a} \leftarrow$ a principal generator of $\ell \bar{I}_\xi I_\gamma$ // see Section 4.3
10 $c \leftarrow c\bar{a}a/\,\mathsf{n}(a)$
11 Find x, y such that $1 = x\bar{\sigma} + y\bar{\chi}$ // see Section 4.4
12 Set $b \leftarrow \bar{x}$ and $d \leftarrow -\bar{y}$
13 **Return** $\mathbf{M} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$

1 mod 8 we have $h = 1$. Since the cost of primality testing is $O(\log^{2+\varepsilon} p)$, we can always bound this step in $O(\log^{4+\varepsilon} p)$. This dominates the cost of a single run of a Cornacchia-type algorithm in Line 3.

- ideal reductions: we need to reduce the ideal I (Line 1) as well as the two other lattices discussed in Section 4.4 (Lines 6 and 11), using LLL on their corresponding bases. In fact there is an implicit fourth reduction, because the last lattice is an intersection $\mathcal{O}_0 \cap \mathcal{O}_0 \bar{\chi}\bar{\sigma}^{-1}$ of two lattices, which in practice is computed as the dual of the sum of the dual lattices; this sum is computed using LLL. If S is such that $|\beta_i|^2 \leq S$ for all input basis elements β_i of any of these four lattices, then the complexity of LLL-reduction is $O(\log^3 S)$ [10, Ch. 2.6.1]. For our complexity claim it suffices to argue that S can be chosen polynomially bounded by p and B, which is easily verified.

All other steps are dominated by these two. The total complexity is $O(\log^{4+\varepsilon} p + \log^3 B)$ as claimed.

5. Solving PIP for extremal orders

We now convert the ideas from Section 3 into a proof of Theorem 1.5. Recall that we are given a left ideal $I \subset \mathrm{M}_g(\mathcal{O}_0)$ of norm n, where $\mathcal{O}_0 \subset B_{p,\infty}$ is an extremal maximal order as in (1), and the goal is to output a matrix $\mathbf{M} \in \mathrm{M}_g(\mathcal{O}_0)$ such that $I = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{M}$. This is established through two reductions:

(i) We first reduce to the case where I is a two-generated left ideal of prime norm ℓ. By "two-generated" we mean that I is given explicitly as $I = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{A} + \mathrm{M}_g(\mathcal{O}_0)\ell$, where the entries of $\mathbf{A}$ can then of course be assumed reduced mod ℓ, i.e., their coordinates with respect to the $\mathbb{Z}$-basis from (1) all lie in $\{0, 1, \ldots, \ell - 1\}$.

(ii) We then reduce to the case $g = 2$ by mimicking the procedure from Sections 3.2–3.3. It is then easy to convert this into an instance of Problem 4.1.

For (i), instead of seeking inspiration in the reduction from Section 3.1, we will use a rerandomization technique that seems interesting in its own right and can be viewed as a higher-dimensional variant of [19, §3.1], see the proof of Lemma 5.3. Nevertheless, as we will discuss in Lemma 5.4, it is possible to mimic the reduction from Section 3.1, which is actually both simpler and more efficient, but this assumes that the factorization of n is known (or easy to compute, e.g., true if n is smooth).

5.1. **Quaternionic matrix arithmetic.** We begin with some comments on the cost of arithmetic in $\mathrm{M}_g(B_{p,\infty})$. Expand the $\mathbb{Z}$-basis $1, b_2, b_3, b_4$ from (1) to a $\mathbb{Z}$-basis

$$(1,0,\ldots,0)^t,\ (b_2,0,\ldots,0)^t,\ \ldots,\ (0,0,\ldots,b_4)^t \tag{7}$$

of $\mathcal{O}_0^g$. This is then, of course, also a $\mathbb{Q}$-basis of $B_{p,\infty}^g$. To be able to invoke standard results from computational linear algebra, we make use of the injective ring homomorphism

$$e\ :\ \mathrm{M}_g(B_{p,\infty}) \to \mathrm{M}_{4g}(\mathbb{Q})\ :\ \mathbf{M}\ \mapsto\ \text{matrix of } \mathbf{v} \mapsto \mathbf{M}\cdot\mathbf{v} \text{ with respect to (7)}.$$

Both evaluating e and recovering $\mathbf{M}$ from $e(\mathbf{M})$ are cheap operations. The entries of $\mathbf{M}$ and $e(\mathbf{M})$ have comparable sizes: if the (absolute values of the) coordinates of the entries of $\mathbf{M}$ are in $O(n)$ then the first $2g$ columns of $e(\mathbf{M})$ are in $O(nq)$ while the entries of the remaining columns may grow up to $O(npq)$. We also have:

Lemma 5.1. $\det(e(\mathbf{M})) = \mathsf{n}(\mathbf{M})^2$.

Proof. By working with respect to the natural generalization of (7) to a basis of $\mathrm{M}_g(B_{p,\infty})$ it is easy to check that $\mathcal{N}(\mathbf{M}) = \det(e(\mathbf{M}))^g$ and a quick sign analysis shows that indeed $\det(e(\mathbf{M})) = \mathsf{n}(\mathbf{M})^2$. □

Here are some example invocations of computational linear algebra results:

Lemma 5.2. *Let $\mathbf{M}_1, \mathbf{M}_2 \in \mathrm{M}_g(\mathcal{O}_0)$ have entries that are reduced mod n. Then $\mathbf{M}_1 \cdot \mathbf{M}_2$, $\mathsf{n}(\mathbf{M}_1)$, $\mathbf{M}_1^{-1}$ can be computed using*

$$O(g^{\omega+\varepsilon}\log^{1+\varepsilon}(npq)),\quad O(g^{\omega+1+\varepsilon}\log^{1+\varepsilon}(npq)),\quad O(g^{\omega+1+\varepsilon}\log^{1+\varepsilon}(npq))$$

bit operations, respectively.

Proof. This follows from the corresponding complexities for matrices over $\mathbb{Z}$, see for instance [15, 16, 31]. □

5.2. **Reduction to prime norm.** Rerandomization works as follows.

Lemma 5.3. *The PIP problem on input a left $\mathrm{M}_g(\mathcal{O}_0)$-ideal of norm n can be solved with one call to a PIP solver on input a two-generated left $\mathrm{M}_g(\mathcal{O}_0)$-ideal of prime norm $\ell \neq p$, $\ell \in O(g^g n^{2g-1} p^{g+\varepsilon})$. Under plausible heuristic assumptions this runs in expected time $T + O(g^{3+\varepsilon}\log^{3+\varepsilon}(gnp) + g^{\omega+1+\varepsilon}\log^{1+\varepsilon}(\beta gnp))$, with T the runtime of the PIP solver on norm-ℓ-input and with β a bound on the (absolute values of the) coordinates of the entries of the matrix it returns.*

Proof. Let I be a left $\mathrm{M}_g(\mathcal{O}_0)$-ideal of norm $\mathsf{n}(I) = n$. By sampling random elements of I, we expect to find an element $\mathbf{A} \in I$ such that $\mathsf{n}(\mathbf{A}) = n\ell$ with ℓ a prime number different from p. Once a suitable $\mathbf{A}$ is found, we can construct the ideal

$$J = \frac{I\,\mathsf{n}(\mathbf{A})\,\mathbf{A}^{-1}}{n}$$

which is a left $\mathrm{M}_g(\mathcal{O}_0)$-ideal by Lemma 2.4. It is easy to verify that

$$J = \mathrm{M}_g(\mathcal{O}_0)\,\mathsf{n}(\mathbf{A})\,\mathbf{A}^{-1} + \mathrm{M}_g(\mathcal{O}_0)\ell$$

and that $\mathsf{n}(J) = \ell^{2g-1}$. Taking conjugate-transposes, we obtain the right $\mathrm{M}_g(\mathcal{O}_0)$-ideal

$$J^* = \mathsf{n}(\mathbf{A})\,(\mathbf{A}^*)^{-1}\,\mathrm{M}_g(\mathcal{O}_0) + \ell\,\mathrm{M}_g(\mathcal{O}_0)$$

which using Lemma 2.7 is verified to equal the pseudo-dual K^{ps} of the left $\mathrm{M}_g(\mathcal{O}_0)$-ideal $K = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{A}^* + \mathrm{M}_g(\mathcal{O}_0)\ell$. This is an ideal of prime norm $\mathsf{n}(K) = \ell$ and it suffices to compute a generator $\mathbf{M}$ for this ideal. Indeed, one then checks that $I = \mathrm{M}_g(\mathcal{O}_0)(\mathbf{M}^*)^{-1}\,\mathbf{A}$, which means that we have solved the PIP problem for I.

The complexity of sampling $\mathbf{A}$ is dominated by the computation of the norm $\mathsf{n}(\mathbf{A})$, which by Lemma 5.2 can be done in time $O(g^{\omega+1+\varepsilon}\log^{1+\varepsilon}(np))$, where we have used $q = O(p)$; recall that we even have $q = O(\log^2 p)$ under GRH. From the Hadamard bound it follows that $\mathsf{n}(\mathbf{A}) = \det(e(\mathbf{A}))^{1/2} = O(g^g n^{2g} p^{g+\varepsilon})$. In particular, we heuristically expect to succeed after $O(g\log(gnp))$ attempts, resulting in $\ell = O(g^g n^{2g-1} p^{g+\varepsilon})$. Each norm computation is succeeded by a primality test with complexity $O(g^{2+\varepsilon}\log^{2+\varepsilon}(gnp))$. We arrive at a total expected cost $O(g^{3+\varepsilon}\log^{3+\varepsilon}(gnp) + g^{\omega+1+\varepsilon}\log^{1+\varepsilon}(gnp))$ for sampling a good $\mathbf{A}$.

Finally, another use of Lemma 5.2 shows that computing $(\mathbf{M}^*)^{-1}$ can be done in time $O(g^{\omega+1+\varepsilon}\log^{1+\varepsilon}(\beta p))$. If $n \le \beta$ then this dominates the cost of the final multiplication with $\mathbf{A}$, while if $n > \beta$ then the cost of this multiplcation is dominated by that of sampling $\mathbf{A}$. □

As announced, if the factorization of n is known or easy to compute, then there is a simpler method, simulating the reduction from Section 3.1.

Lemma 5.4. *Solving the PIP problem on input a left $\mathrm{M}_g(\mathcal{O}_0)$-ideal of norm n, coprime to p and with known factorization, can be done by running a PIP solver on input a left, two-generated $\mathrm{M}_g(\mathcal{O}_0)$-ideal of norm ℓ, for each prime factor $\ell \mid n$. This runs in expected time $\sum_\ell T_\ell + O(g^{\omega+1+\varepsilon}\log^{3+\varepsilon}(\beta np))$ with T_ℓ the runtime of the PIP solver for the prime ℓ, and with β an overall bound on the coordinates of the (absolute values of the) entries of the matrices it returns.*

Proof. Let I be our input ideal, write $n = \ell_1\ell_2\cdots\ell_r$, let $I_0 = I$, and repeat the following steps for $i = 1, \ldots, r$:

(1) Compute the isomorphism $\Phi_{\ell_i} : \mathrm{M}_g(\mathcal{O}_0)/\ell_i\,\mathrm{M}_g(\mathcal{O}_0) \to \mathrm{M}_{2g}(\mathbb{F}_{\ell_i})$ from Section 2.2. Find a non-zero vector $\mathbf{v}$ in the joint kernel of $\Phi_{\ell_i}(I_{i-1})$. Let $\mathbf{N}_i$ be a preimage under Φ_{ℓ_i} of a corank-1 matrix with kernel $\mathbf{v}$. Run the PIP solver on $\mathrm{M}_g(\mathcal{O}_0)\,\mathbf{N}_i + \mathrm{M}_g(\mathcal{O}_0)\ell_i$ to find a generator $\mathbf{M}_i$.
(2) Let I_i be the ideal obtained from I_{i-1} by replacing each generator $\mathbf{A}$ with $\mathbf{A}\,\mathbf{M}_i^{-1}$. Observe that these generators are indeed elements of $\mathrm{M}_g(\mathcal{O}_0)$, and $\mathsf{n}(I_i) = n/(\ell_1\cdots\ell_i)$.

Then $\mathbf{M} = \mathbf{M}_r\cdots\mathbf{M}_2\,\mathbf{M}_1$ is our desired output. Apart from the runs of our PIP solver, the cost is dominated by the computations of $\mathbf{A}\,\mathbf{M}_i^{-1}$ and $\mathbf{M} = \mathbf{M}_r\cdots\mathbf{M}_2\,\mathbf{M}_1$. From Lemma 5.2 we see that the r inverses $\mathbf{M}_i^{-1}$ take $O(rg^{\omega+1+\varepsilon}\log^{1+\varepsilon}(\beta p))$ time to compute. Using an accumulation tree the product $\mathbf{M}$ can be computed in time $O(r^2 g^{\omega+\varepsilon}\log^{1+\varepsilon}(\beta p))$. The lemma follows from $r = O(\log n)$ and the fact that each $\mathbf{A}$ is reduced mod n. □

5.3. **Reduction to $g = 2$.** We mimic the procedure outlined in Sections 3.2–3.3.

Lemma 5.5. *Solving the PIP problem on input a left, two-generated $\mathrm{M}_g(\mathcal{O}_0)$-ideal of prime norm ℓ can be done by running a PIP solver on input a left, two-generated $\mathrm{M}_2(\mathcal{O}_0)$-ideal of norm ℓ. This takes expected time $T + O(g^{\omega+\varepsilon} \log^{1+\varepsilon}(\beta\ell p))$ with T the time taken by the PIP solver and β a bound on the (absolute values of the) coordinates of the entries of the matrix it returns.*

Proof. Let $I = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{A} + \mathrm{M}_g(\mathcal{O}_0)\ell$ be our left ideal, assumed to be of prime norm $\mathsf{n}(I) = \ell$. Let $\Phi_\ell : \mathcal{O}_0/\ell\mathcal{O}_0 \to \mathrm{M}_2(\mathbb{F}_\ell)$ be an isomorphism as in Section 2.2 and extend it block-wise to an isomorphism $\mathrm{M}_g(\mathcal{O}_0)/\ell\,\mathrm{M}_g(\mathcal{O}_0) \to \mathrm{M}_{2g}(\mathbb{F}_\ell)$; abusingly, we denote the resulting map by Φ_ℓ as well. Write $\mathbf{R} = \Phi_\ell(\mathbf{A})$. By Corollary 2.6 we have $\mathrm{rank}(\mathbf{R}) = 2g - 1$, i.e., the kernel of $\mathbf{R}$ is one-dimensional. We can find an invertible matrix $\mathbf{S}$ such that $\ker(\mathbf{R}\,\mathbf{S}) = \langle (w_1, w_2, 0, \ldots, 0)^t \rangle$ for some $w_1, w_2 \in \mathbb{F}_\ell$. Moreover, by mimicking Section 3.2 we see that $\mathbf{S}$ can be chosen of the form $\Phi_\ell(\mathbf{\Pi}\,\mathbf{B})$ with $\mathbf{B} \in \mathrm{GL}_g(\mathcal{O}_0)$ as in (4) and $\mathbf{\Pi}$ a row-swapping permutation matrix. It is clear that if we manage to find a generator $\mathbf{M}$ of the ideal $I' = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{A}\,\mathbf{\Pi}\,\mathbf{B} + \mathrm{M}_g(\mathcal{O}_0)\ell$, then $\mathbf{M}\,\mathbf{B}\,\mathbf{\Pi}$ is a generator of I (here we have used that $\mathbf{B}$ and $\mathbf{\Pi}$ are self-inverse).

Thus we can assume w.l.o.g. that $\mathbf{R} = \Phi_\ell(\mathbf{A})$ has a kernel of the foregoing form. If we then pick a rank-3 matrix $\mathbf{R}_2 \in \mathrm{M}_4(\mathbb{F}_\ell)$ with kernel $\langle (w_1, w_2, 0, 0) \rangle$ and we let $\mathbf{A}_2 \in \mathrm{M}_2(\mathcal{O}_0)$ be any representant of its preimage under Φ_ℓ (now applied to $g = 2$), then we can consider the ideal $I_2 = \mathrm{M}_2(\mathcal{O}_0)\,\mathbf{A}_2 + \mathrm{M}_2(\mathcal{O}_0)\ell$. If $\mathbf{M}_2$ is a generator of this ideal, then the matrix

$$\mathbf{M} = \begin{pmatrix} \mathbf{M}_2 & 0 \\ 0 & \mathbf{I}_{g-2} \end{pmatrix}$$

is a generator of I, as wanted.

All linear algebra over $\mathbb{F}_\ell$ can be done in time $O(g^\omega \log^{1+\varepsilon} \ell)$. The matrix $\mathbf{B}$ can be chosen to have entry coordinates in $O(\ell)$. By Lemma 5.2 the computation of $\mathbf{A}\,\mathbf{\Pi}\,\mathbf{B}$, $\mathbf{M}\,\mathbf{B}\,\mathbf{\Pi}$ can be done in time $O(g^{\omega+\varepsilon} \log^{1+\varepsilon}(\beta\ell p))$. □

5.4. **An instance of Problem 4.1.** Given $I_2 = \mathrm{M}_2(\mathcal{O}_0)\,\mathbf{A}_2 + \mathrm{M}_2(\mathcal{O}_0)\ell$ as in the proof of Lemma 5.5 above, it now suffices to find a matrix $\mathbf{M}_2 \in \Lambda$ such that $\ker(\Phi_\ell(\mathbf{M}_2)) = \ker(\mathbf{R}_2)$ and $\mathsf{n}(\mathbf{M}_2) = \ell$. Indeed, since I_2 and $\mathrm{M}_2(\mathcal{O}_0)\,\mathbf{M}_2$ have the same norm, it is sufficient to show one inclusion. Because $\Phi_\ell(\mathbf{M}_2)$ and $\mathbf{R}_2$ have the same kernel, we can find a matrix $\mathbf{T} \in \mathrm{M}_4(\mathbb{F}_\ell)$ such that $\mathbf{T}\,\mathbf{R}_2 = \Phi_\ell(\mathbf{M}_2)$. By taking inverse images, $\Phi_\ell^{-1}(\mathbf{T})\,\mathbf{A}_2 \equiv \mathbf{M}_2 \bmod \ell$, i.e., $\mathbf{M}_2 \in \mathrm{M}_2(\mathcal{O}_0)\,\mathbf{A}_2 + \mathrm{M}_2(\mathcal{O}_0)\ell$. Finding $\mathbf{M}_2$ is indeed an instance of Problem 4.1 applied to the $\mathcal{O}_0$-ideal generated by ℓ and the elements in the first column of $\mathbf{A}_2$.

Combining the complexity of Lemma 5.5 with the complexity of Theorem 4.2 proves the following special case of Theorem 1.5.

Theorem 5.6. *There exists a Las Vegas algorithm which on input a left ideal $I = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{A} + \mathrm{M}_g(\mathcal{O}_0)\ell$ of prime norm ℓ, where $\mathcal{O}_0 \subset B_{p,\infty}$ is an extremal maximal order and the entries of $\mathbf{A}$ are reduced mod ℓ, outputs a matrix $\mathbf{M} \in \mathrm{M}_g(\mathcal{O}_0)$ such that $I = \mathrm{M}_g(\mathcal{O}_0)\,\mathbf{M}$. Under plausible heuristic assumptions, the expected runtime of the algorithm is*

$$O(\log^{4+\varepsilon} p + \log^3(\ell p) + g^{\omega+\varepsilon} \log^{1+\varepsilon}(\ell p)).$$

Here we have used that the norms of the elements in the first column of $\mathbf{A}_2$, which serve as input to Problem 4.1, can be assumed to be in $O(\ell^2 p)$, and that the matrix returned by Theorem 4.2 has entries with norms in $O(\ell p^{3/2+\varepsilon})$.

5.5. **Proof of Theorem 1.5.** From the the proof of Lemma 5.5 and the fact that the matrix returned by Theorem 4.2 has entries with norm $O(\ell p^{3/2+\varepsilon})$, it is easy to check that the algorithm from Theorem 5.6 returns a matrix whose entries have norms in $O(g^2\ell^2 p^{3/2+\varepsilon})$. For Theorem 1.5 we need to run this algorithm with $\ell \in O(g^g n^{2g-1} p^{g+\varepsilon})$, so it takes time

$$O(\log^{4+\varepsilon} p + g^3 \log^3(gnp) + g^{\omega+1+\varepsilon} \log^{1+\varepsilon}(gnp))$$

and outputs a matrix with entry coordinates bounded by $\beta = O(g^{g+1} n^{2g-1} p^{g+3/4+\varepsilon})$. The statement then follows by adding the overhead stated in Theorem 5.6.

Remark 5.7. As stated in Corollary 1.6, this immediately gives an algorithm for solving Problem 1.1 for arbitrary maximal orders $\Lambda \subset \mathrm{M}_g(B_{p,\infty})$. Of course, away from $\mathrm{M}_g(\mathcal{O}_0)$, the main case of interest is $\Lambda = \mathrm{M}_g(\mathcal{O})$ for an arbitrary maximal order $\mathcal{O} \subset B_{p,\infty}$. This case allows for an important optimization: up to and including our reduction to the case $g = 2$, none of the reasonings rely on the order being extremal. Therefore the reduction to the extremal case can be postponed until after this step, which means that the first run of the algorithm from Theorem 1.5 is on $J_{\mathrm{M}_2(\mathcal{O}),\mathcal{O}_0}$ rather than on $J_{\mathrm{M}_g(\mathcal{O}),\mathcal{O}_0}$.

6. Finalizing kernel-to-matrix conversion

In Section 3 we have explained how to reduce Problem 1.2 to the case $g = 2$, $H = \langle (P, 0) \rangle$, and $\mathrm{ord}(P) = \ell$ a prime. We will now fill in the last missing details towards establishing Theorem 1.8.

6.1. **Missing algorithmic details.** In Section 3.2, the quaternions b_i can be found by extending P_j to a basis P_j, Q_j of $E[\ell]$ and computing the isomorphism $\Phi_\ell : \mathcal{O}/\ell\mathcal{O} \to \mathrm{M}_2(\mathbb{Z}/\ell\mathbb{Z})$ from Lemma 2.8 with respect to this basis. This comes at the cost of computing discrete logarithms in a group of size ℓ, which is also needed to express each point P_i as $x_i P_j + y_i Q_j$ (hence, with the current state-of-the-art, these steps are not polynomial-time in $\log \ell$). Then b_i can be taken as any preimage under Φ_ℓ of a matrix sending $(1, 0)$ into (x_i, y_i).

To reduce the case $g = 2$ to Problem 4.1 we need to compute the ideal $I_P \subset \mathcal{O}$ annihilating P, which is done by extending P to a basis $P, Q \in E[\ell]$, fixing a $\mathbb{Z}$-basis $1, b_2, b_3, b_4 \in \mathcal{O}$, expressing $\iota(b_i)(P)$ in terms of P and Q (discrete logarithm computations again), and some linear algebra.

If $\mathcal{O} = \mathcal{O}_0$, then we conclude by running the method described in Section 4. But to handle the case of arbitrary $\mathcal{O}$, it is in fact convenient to view this step as a PIP problem, allowing to drop the condition on $\mathcal{O}$. As already observed, the matrix

$$\mathbf{A} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

with $a, c \in I_P$ will annihilate H regardless of the choice of b and d. From Lemma 2.9 we know that $\mathsf{n}(\mathbf{A}) = \ell n$ for some positive integer n. By letting a, c be genuine generators of I_P and resampling b, d if necessary, we can assume $\gcd(n, \ell) = 1$. Now a generator of the principal left $\mathrm{M}_2(\mathcal{O})$-ideal generated by $\mathbf{A}$ and ℓ will be a matrix annihilating $(P, 0)$ with norm ℓ, hence a solution to Problem 1.2.

Remark 6.1 (immediate conversion to PIP). Of course, it is also possible to convert Problem 1.2 into a PIP instance in $\mathrm{M}_g(\mathcal{O})$ right away, as discussed in the introduction; this is the approach taken by Chu [9, Alg. A.2.2, Step 17]. It involves $O(g)$ discrete log computations in groups of size $O(|H|)$; to some extent, one can view the reduction from Section 3.1 as a simultaneous Pohlig–Hellman reduction over all components of $E^g \supset H$, which is somewhat more efficient (also the linear algebra required to find generators of the ideal drops from dimension $4g^2$ to 4).

6.2. **Complexity.** To run the reduction described in Section 3 we first need to factor $|H|$, with heuristic worst-case complexity $L_{|H|}(1/3)$ when using the general number field sieve. As explained above, we also need to compute $O(g)$ discrete logarithms in $E[\ell]$, for all prime factors $\ell \mid |H|$. Let k be minimal such that the points of H are defined over $\mathbb{F}_{p^{2k}}$. By twisting if needed, we can assume that the p^2-Frobenius endomorphism on E equals the multiplication-by-$\pm p$ map, implying that $E[\ell] \subset E(\mathbb{F}_{p^{2k}})$ for all $\ell \mid |H|$. The worst-case complexity of the discrete-logarithm computations is then $gL_{p^{2k}}(1/3)$, again using the general number field sieve. But in many cases of interest, e.g., if $|H|$ is smooth, these steps can be done much more efficiently. Nevertheless, they always dominate the remaining costs in Sections 3.1-3.3 , which we therefore suppress from the analysis.

This reduces Problem 1.2 to principalizing one two-generated ideal I_ℓ for every prime factor $\ell \mid |H|$. To do this, as explained, we first apply Theorem 1.5 to the connecting ideal $J_{\mathrm{M}_2(\mathcal{O}),\mathcal{O}_0}$, say with output $\mathbf{C}$, then solve PIP for the ideals

$$\mathbf{C}\, I_\ell\, \mathbf{C}^{-1} \subset \mathrm{M}_2(\mathcal{O}_0)$$

using Theorem 5.6, and then assemble the outputs into the desired matrix $\mathbf{M} \in \mathrm{M}_g(\mathcal{O})$. Writing $m = \mathtt{n}(J_{\mathrm{M}_2(\mathcal{O}),\mathcal{O}_0})$, $r = O(\log |H|)$ for the number of prime factors of $|H|$, and using that the entries of $\mathbf{C}$ are polynomially bounded by mp, this leads to a total complexity of $O(F_H + gD_H + r\log^{4+\varepsilon} p + \log^{3+\varepsilon}(|H|mp))$ before the assembling, which is dominated by the cost of computing $\mathbf{M}_r \cdots \mathbf{M}_2\, \mathbf{M}_1$. This can be done at cost $O(r^2 g^{\omega+\varepsilon} \log^{1+\varepsilon}(mp))$ as in Lemma 5.4 and establishes Theorem 1.8.

7. Implementation

We have implemented our main algorithms in SageMath; the code is available at `https://github.com/KULeuven-COSIC/PIP`. Our proof of concept implementation solves Problem 1.1 and Problem 1.2 in the case $g = 2$ and the input is an $\mathcal{O}_0$-ideal of prime norm ℓ, resp., a prime-order-ℓ point $(P, Q) \in E^2$ with E a supersingular elliptic curve such that $\mathrm{End}(E) \cong \mathcal{O}_0$. Solving PIP where p and ℓ are 256-bit primes takes about 0.05 seconds on an Intel i5 Alder Lake; solving kernel-to-matrix where p is a 256-bit prime and ℓ is a 20-bit prime takes about 0.1 seconds, of which 0.05 seconds are spent on the commutative square step. This shows the effectiveness of our algorithms even for cryptographic sizes.

References

[1] Eric Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.

[2] Jean-François Biasse, Claus Fieker, Tommy Hofmann, and William Youmans. An algorithm for solving the principal ideal problem with subfields. *Advances in Mathematics of Communications*, 18(6):1785–1809, 2024. `https://doi.org/10.3934/amc.2023021`.

[3] Dan Boneh, Darren B. Glass, Daniel Krashen, Kristin E. Lauter, Shahed Sharif, Alice Silverberg, Mehdi Tibouchi, and Mark Zhandry. Multiparty non-interactive key exchange and more from isogenies on elliptic curves. *J. Math. Cryptol.*, 14(1):5–14, 2020.

[4] Wouter Castryck, Thomas Decru, Péter Kutas, Abel Laval, Christophe Petit, and Yan Bo Ti. KLPT^2: Algebraic pathfinding in dimension two and applications. In *CRYPTO 2025 Part I*, volume 16000 of *Lecture Notes in Computer Science*, pages 167–200, 2025. https://doi.org/10.1007/978-3-032-01855-7_6.

[5] Wouter Castryck, Thomas Decru, and Benjamin Smith. Hash functions from superspecial genus-2 curves using Richelot isogenies. In *Proceedings of NutMiC 2019*, volume 14 of *Journal of Mathematical Cryptology*, pages 268–292, 2020.

[6] Wouter Castryck, Riccardo Invernizzi, Jonathan Komada Eriksen, and Frederik Vercauteren. The principal ideal problem for endomorphism rings of superspecial abelian varieties (online version). Cryptology ePrint Archive, Paper 2026/454, 2026.

[7] 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/.

[8] 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 *Lecture Notes in Computer Science*, pages 99–130, 2023. https://doi.org/10.1007/978-981-99-8727-6_4.

[9] 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.

[10] Henri Cohen. *A course in computational algebraic number theory*, volume 138. Springer Science & Business Media, 2013.

[11] David A. Cox. *Primes of the Form $x^2 + ny^2$: Fermat, Class Field Theory, and Complex Multiplication.* John Wiley & Sons, 2nd edition, 2013.

[12] 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.

[13] Pierrick Gaudry, Julien Soumier, and Pierre-Jean Spaenlehauer. Computing isomorphisms between products of supersingular elliptic curves. Cryptology ePrint Archive, Paper 2025/136, 2025.

[14] Daniel Goldstein and Murray Schacher. Norms in central simple algebras. *Pacific Journal of Mathematics*, 292(2):373–388, 2018. https://doi.org/10.2140/pjm.2018.292.373.

[15] David Harvey and Joris van der Hoeven. On the complexity of integer matrix multiplication. *Journal of Symbolic Computation*, 89:1–8, 2018.

[16] Erich Kaltofen and Gilles Villard. On the complexity of inverting integer and polynomial matrices. *Computational complexity*, 13:91–130, 2004.

[17] 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.

[18] Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM Journal on Computing*, 39(5):1714–1747, 2010. https://doi.org/10.1137/080734467.

[19] 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.

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

[21] Mickaël Montessinos. Equivalent computational problems for superspecial abelian surfaces. Cryptology ePrint Archive, Paper 2026/120, 2026.

[22] Arthur Ogus. Supersingular K3 crystals. In *Journées de Géométrie Algébrique de Rennes 1978 Part II*, volume 64 of *Astérisque*, pages 3–86, 1979.

[23] Aurel Page. An algorithm for the principal ideal problem in indefinite quaternion algebras. In *Proceedings of ANTS-XI*, volume 17 of *LMS Journal of Computation and Mathematics*, pages 366–384, 2014.

[24] Aurel Page, Damien Robert, and Julien Soumier. On the conversion of module representations for higher dimensional supersingular isogenies. Cryptology ePrint Archive, Paper 2026/276, 2026.

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

[26] Irving Reiner. *Maximal Orders*, volume 28 of *London Mathematical Society Monographs*. Oxford University Press, 2003. https://doi.org/10.1007/978-3-030-56694-4.

[27] Robbe Rietveld. *Modular Arithmetic of Quaternion Norms*. Bachelor thesis, Leiden University, 2025. Available at https://studenttheses.universiteitleiden.nl/handle/1887/4262328.

[28] Lajos Rónyai. Computing the structure of finite algebras. *Journal of Symbolic Computation*, 9(3):355–373, 1990.

[29] Tetsuji Shioda. Supersingular K3 surfaces. In *Algebraic Geometry Copenhagen 1978*, volume 732 of *Lecture Notes in Mathematics*, pages 564–591, 1979.

[30] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer New York, 2009.

[31] Arne Storjohann. On the complexity of inverting integer and polynomial matrices. *Computational complexity*, 24:777–821, 2015.

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

[33] William C. Waterhouse. Abelian varieties over finite fields. *Annales Scientifiques de l'École Normale Supérieure*, 2(4):521–560, 1969.