

Computing Orientations from the Endomorphism Ring of Supersingular Curves and Applications

Jonathan Komada Eriksen³, Antonin Leroux^{1,2}

¹ DGA-MI, Bruz, France antonin.leroux@polytechnique.org

² IRMAR, Université de Rennes, France

³ Norwegian University of Science and Technology, Norway

Abstract. This work introduces several algorithms related to the computation of orientations in endomorphism rings of supersingular elliptic curves. This problem boils down to representing integers by ternary quadratic forms, and it is at the heart of several results regarding the security of oriented-curves in isogeny-based cryptography.

Our main contribution is to show that there exists efficient algorithms that can solve this problem for quadratic orders of discriminant n up to $O(p^{4/3})$. Our approach improves upon previous results by increasing this bound from $O(p)$ to $O(p^{4/3})$ and removing some heuristics.

We introduce several variants of our new algorithm and provide a careful analysis of their asymptotic running time (without heuristic when it is possible). The best proven asymptotic complexity of one of our variant is $O(n^{3/4}/p)$ in average. The best heuristic variant has a complexity of $O(p^{1/3})$ for big enough n .

We then introduce several results regarding the computation of ideals between oriented orders. The first application of this is a simplification of the known reduction from vectorization to computing the endomorphism ring, removing the assumption on the factorization of the discriminant. As a second application, we relate the problem of computing fixed-degree isogenies between supersingular curves to the problem of computing orientations in endomorphism rings, and we show that for a large range of degree d , our new algorithms improve on the state-of-the-art, and in important special cases, the range of degree d for which there exist a polynomial-time algorithm is increased. In the most special case we consider, when both curves are oriented by a small degree endomorphism, we show heuristically that our techniques allow the computation of isogenies of any degree, assuming they exist.

1 Introduction

Isogeny-based cryptography uses supersingular elliptic curves and isogenies between them to construct cryptographic schemes. An essential part of isogeny-based cryptography is the Deuring correspondence, relating supersingular elliptic curves over $\overline{\mathbb{F}}_p$ to maximal orders in a quaternion algebra ramified at p and ∞ , and isogenies to ideals, by passing to the endomorphism ring of the curve.

One particular flavour of isogeny-based schemes [3, 5, 9] use the extra information of an *orientation*, which is an embedding of a quadratic imaginary order inside the endomorphism ring. This subring corresponds to an embedding of an imaginary quadratic order \mathfrak{O} into the endomorphism ring, which in turn allows one to consider the action of \mathfrak{O} -ideals on the curves (primitively) oriented by \mathfrak{O} through \mathfrak{O} -oriented isogenies. It is a well known fact that $\text{Cl}(\mathfrak{O})$ acts freely on the set of primitively \mathfrak{O} -oriented curves (up to oriented isomorphisms) in one or two orbits [15].

An important part of the study of the schemes using these oriented curves and isogenies is to understand the link of oriented problems with generic non-oriented problems. One of the main object of study in this context is the embedding problem which was first studied in [19] (although not under that name). We present it as Problem 1.

Problem 1. (Quaternion order embedding problem.) Let p be a prime number, let \mathcal{O} be a maximal order inside $B_{p,\infty}$ and let t, n be such that there exists an element of norm n and trace t in \mathcal{O} . Find $\alpha \in \mathcal{O}$ with

$$n(\alpha) = n, \quad \text{tr}(\alpha) = t \tag{1}$$

Related Works. Oriented curves first appeared in isogeny-based cryptography with the CSIDH group action [3]. However, they were not defined as such at the time. The notion of orientation was introduced formally by Kohel and Colo in [5] together with a new group action called OSIDH. Some of the results of [5] were refined by Onuki [15]. These works introduced generic hard problems such as \mathfrak{O} -vectorization.

At first, the only applications of orientations were related to these group actions, but a broader link with the other areas of isogeny-based cryptography was demonstrated by De Feo et al. in [7] with the introduction of the \mathfrak{O} -uber isogeny problem. The authors of [7] provided in particular some reductions between flavours of the \mathfrak{O} -uber isogeny problem and generic isogeny computation problems.

In 2022, Wesolowski provided a much more complete picture in [19] by studying all orientation-related problems and providing several reductions between them, and generic problems such as the endomorphism ring problem. In particular, Wesolowski proposed the first algorithm to solve the quaternion order embedding problem when the discriminant is smaller than \sqrt{p} , and proved some relations between the \mathfrak{O} -vectorization, the \mathfrak{O} -uber isogeny, and problems related to the computation of endomorphism rings (with or without the knowledge of an orientation).

An improved heuristic algorithm to solve the embedding problem was introduced in [1] that increases the bound for when the embedding problem is solvable in polynomial time from disc $\mathfrak{O} = O(\sqrt{p})$ to disc $\mathfrak{O} = O(p)$.

In [12], Leroux proved a lower bound on the number of oriented curves by using quaternion orders generated by two non-commuting quadratic orders. The same ideas are going to be crucial in our new algorithms.

An algorithm to solve the embedding problem can be used to find fixed degree isogeny between supersingular elliptic curves. This is an important problem in isogeny-based cryptography that was first studied from the quaternionic perspective in [11] with the famous KLPT algorithm. This algorithm has found numerous applications in cryptography in the study of the Deuring correspondence (see the reductions of [8] or the signature scheme from [6] for instance). Understanding and improving the known algorithms to find isogenies of fixed degree between supersingular curves is an important task. While previous literature had been focusing on identifying cases for which there was a polynomial-time algorithm (such as KLPT), the recent article [2] was the first one to provide a generic analysis of the run-time of such algorithms in ranges of input where the running-time is not known to be polynomial.

1.1 Our Contributions

In this work, we study orientations purely on the quaternion side. Our main contribution is a set of new algorithms for solving the quaternion order embedding problem (Problem 1), which can be executed in polynomial time for disc \mathfrak{D} up to $O(p^{4/3})$.

`GenericOrderEmbedding`, our first algorithm, treats the generic case of an arbitrary quaternion order. Its complexity depends on the size of the first, second, and third successive minima of the ternary quadratic form associated to \mathcal{O} . When \mathcal{O} is a random quaternion order of discriminant Δ the expected running time is polynomial when $\text{disc } \mathfrak{D} = O(\Delta^{4/3})$.

From there, we deduce two other algorithms. `MaximalOrderEmbeddingEichler` uses `GenericOrderEmbedding` as a building block by applying it on several Eichler sub-orders of the maximal provided in input. We show that the average running time is asymptotically better than a direct application of `GenericOrderEmbedding`.

In some good cases where \mathcal{O} contains a particularly small element, we can go beyond the $O(p^{4/3})$ bound at the cost of using a factorization oracle, under some heuristics. The resulting algorithm `GenericOrderEmbeddingFactorization` can be seen as a generalization of the algorithm from [1], and in the best cases where the \mathcal{O} contains an element of norm $O(1)$, it runs in polynomial time for any discriminant. Further, for any order, the runtime is always upper bounded as $O(p^{1/3})$, independent of the size of the discriminant.

In the second part, we study ideals between oriented quaternion orders. We show that when the orientation of the quaternion orders induce the same orientation of K into $B_{p,\infty}$, their connecting ideal is always generated by the image of a quadratic ideal. We apply this result to give a new, simple reduction to show that the \mathfrak{D} -vectorization problem reduces to the endomorphism ring problem, a result previously only known for when the factorization of disc \mathfrak{D} was known, and assuming \mathfrak{D} has a small number of genera [19, Theorem 2].

We also give a heuristic reduction from the problem of computing equivalent ideals of a given norm to the quaternion order embedding problem, and show that in important special cases, our algorithms improves the range for which

this is solvable in polynomial time. In particular, in the special case where the two maximal orders are optimally embedded by quadratic orders of very small discriminant, it is possible to find equivalent ideals of any norm efficiently. We also obtain a heuristic improvement in the best known asymptotic complexity to solve this problem in the generic case, showing that it is always solvable in time $O(p^{2/3})$, improving on the results from [2] for a wide variety of degrees d .

We implement our algorithms in SageMath [17]. The implementation can be found at:

<https://github.com/Jonathke/Computing-Optimal-Embeddings>

1.2 Technical Overview

Let us take an order \mathcal{O} of discriminant Δ , and elements t, n, α as in Problem 1.

Overview of the algorithms. Our new algorithms to find elements of given norm and trace are mainly built upon an oracle to find trace pairings, i.e. the value of the trace of the product of the element α with some elements β of \mathcal{O} . This oracle is built by looking at the discriminant of the quaternion order $\mathbb{Z}[1, \alpha, \beta, \alpha\beta]$ and seeing that its discriminant must be divided by Δ when α and β do not commute. This gives an equation on $\text{tr}(\alpha\beta)$ modulo Δ . And this equation is enough to recover the value over \mathbb{Z} when $n(\alpha\beta) < \Delta^2$.

We obtain our algorithm `GenericOrderEmbedding` by applying this idea on a reduced basis $1, \beta_1, \beta_2, \beta_3$ of \mathcal{O} and enumerating all possible solutions until the correct one is found. As for a random order \mathcal{O} we can expect $n(\beta_1) \approx n(\beta_2) \approx n(\beta_3) \approx \Delta^{2/3}$, this will be efficient to recover α when $n = O(\Delta^{4/3})$ and we can show that asymptotically (when n grows and p remains fixed) the complexity of this algorithm is $O(n^{3/2}/p^2)$.

Our algorithm `MaximalOrderEmbeddingEichler` is obtained by trying to apply `GenericOrderEmbedding` on all Eichler sub-order of order N (where N is chosen to ensure that each execution `GenericOrderEmbedding` should be polynomial in average and that there is one execution that will succeed). We show that the average running time of this algorithm is $O(n^{3/4}/p)$.

Finally, in cases where $n(\beta_1)$ is smaller than the expected $\Delta^{2/3}$, the trace pairing $\text{tr}(\alpha\beta_1)$ will be much smaller than $\text{tr}(\alpha\beta_j)$ for $1 < j \leq 3$. Thus, it will be possible to determine $\text{tr}(\alpha\beta_1)$ exactly for values of n bigger than $\Delta^{4/3}$. In those cases, we can exploit the knowledge of $\text{tr}(\alpha\beta_1)$ to translate the embedding problem to a problem of representing some integer by some binary quadratic form. It is well known that such equation can be solved in polynomial time with the help of a factorization oracle. This yields the `GenericOrderEmbedding-Factorization` algorithm.

Acknowledgement

The authors would like to thank Benjamin Wesolowski for a fruitful discussion that motivated this work. We would also like to thank Wouter Castryck,

Gioella Lorenzon, Riccardo Invernizzi and Frederik Vercauteren, with whom useful discussions among other things lead to a cleaner description of `GenericOrderEmbeddingFactorization`, the introduction of `Heuristic 2`, and subsequent heuristic running times. Finally, we thank the members of Group project F at the 2023 Bristol isogeny workshop for useful discussions about several aspects of this work. We also thank the organizers of this workshop for providing a nice framework for these discussions.

2 Mathematical Background

A *quaternion algebra* B is a four dimensional \mathbb{Q} -algebra with a \mathbb{Q} -basis $1, i, j$, satisfying

$$i^2 = a, j^2 = b, k = ij = -ji,$$

for some $a, b \in \mathbb{Q}^\times$. Elements $\alpha = x + iy + jz + kw \in B$ have a conjugate $\bar{\alpha} = x - iy - jz - kw$, and from this we define the reduced norm $n(\alpha) := \alpha\bar{\alpha}$ and reduced trace $\text{tr}(\alpha) := \alpha + \bar{\alpha}$.

The values a, b determine the places where B *ramify*, which again determines B up to isomorphism. In this work, we fix a prime p , and focus on the quaternion algebra $B_{p,\infty}$ ramified at p and ∞ .

A *lattice* in $B_{p,\infty}$ is a \mathbb{Z} -submodule $L \subseteq B_{p,\infty}$ of rank 4. Lattices have an invariant called the discriminant, defined as

$$\text{disc } L = \det(\text{tr}(\beta_i\beta_j)_{i,j})$$

where $\beta_1, \beta_2, \beta_3, \beta_4$ is a \mathbb{Z} -basis of L . A lattice \mathcal{O} is called an *order* if it is also a subring of $B_{p,\infty}$, i.e. it contains 1, and is closed under multiplication. The discriminant of an order is always a square, hence we can define the *reduced discriminant*

$$\text{discrd } \mathcal{O} = \sqrt{\text{disc } \mathcal{O}} \in \mathbb{Z}$$

In $B_{p,\infty}$, orders \mathcal{O} always satisfy

$$\text{discrd } \mathcal{O} = pN,$$

where $N := [\mathcal{O}_0 : \mathcal{O}]$, for some maximal order \mathcal{O}_0 containing \mathcal{O} .

In the rest of this document, we will always use the reduced discriminant of quaternion order despite very often using the word *discriminant* and using the notation $\text{disc } \mathcal{O}$.

2.1 On Successive Minimas in a Quaternion Order

We define the successive minimas of a quaternion order \mathcal{O} to be the successive minimas of \mathcal{O}/\mathbb{Z} .

Below, we prove several simple results bounding the successive minimas of quaternion orders. Most of those results are folklore and/or very easy to prove but we restate them for convenience.

In all this section, \mathcal{O} is a quaternion order of reduced discriminant Δ and $\beta_1, \beta_2, \beta_3$ realizes the successive minimas of \mathcal{O} .

Proposition 1. (Minkowski) $\frac{4}{3}\Delta^2 \leq n(\beta_1)n(\beta_2)n(\beta_3) \leq 8\Delta^2$

Lemma 1. $n(\beta_1) \leq 2\Delta^{2/3}$.

Proof. This follows from combining Proposition 1 with $n(\beta_2)n(\beta_3) \geq n(\beta_1)^2$. \square

Lemma 2. $n(\beta_2) \leq 2\sqrt{2}\Delta/\sqrt{n(\beta_1)}$.

Proof. By Proposition 1, $n(\beta_1)n(\beta_2)^2 \leq n(\beta_1)n(\beta_2)n(\beta_3) \leq 8\Delta^2$, and this proves the result. \square

Lemma 3. $n(\beta_2) \geq \Delta/(4n(\beta_1))$.

Proof. The quaternion order $\mathbb{Z}[1, \beta_1, \beta_2, \beta_1\beta_2]$ is contained in \mathcal{O} . Thus, by Proposition 2, we have $\Delta \leq 4n(\beta_1)n(\beta_2)$. \square

Lemma 4. $n(\beta_1) \leq \frac{2\sqrt{2}\Delta}{\sqrt{n(\beta_3)}}$

Proof. Combining Proposition 1 with $n(\beta_1)^2n(\beta_3) \leq n(\beta_1)n(\beta_2)n(\beta_3) \leq 8\Delta^2$ proves the result. \square

Lemma 5. $n(\beta_3) \leq 32\Delta$

Proof. The result follows from the combination of Lemma 3 with Proposition 1. \square

2.2 Oriented Orders

The main focus in this paper is on *optimal embeddings*. Our main motivation is the relation to *primitively \mathfrak{D} -oriented curves*, defined as the pair (E, ι) , where E is a supersingular elliptic curve, and $\iota : K \hookrightarrow \text{End}(E) \otimes \mathbb{Q}$ is an optimal embedding of \mathfrak{D} into $\text{End}(E)$, i.e. such that

$$\iota|_{\mathfrak{D}} : \mathfrak{D} \hookrightarrow \text{End}(E)$$

satisfies

$$\iota(K) \cap \mathcal{O} = \iota(\mathfrak{D}).$$

Hence, we introduce the analogous notation, which will be used repeatedly in Section 4:

Definition 1. Let K an imaginary quadratic field, with $\mathfrak{D} \subseteq K$ an imaginary quadratic order and let B be a definite quaternion algebra over \mathbb{Q} with $\mathcal{O} \subseteq B$ an order. Given an embedding $\iota : K \hookrightarrow B$, we can define a \mathfrak{D} -oriented order to be the pair (\mathcal{O}, ι) , whenever $\iota(\mathfrak{D}) \subseteq \mathcal{O}$. Further, (\mathcal{O}, ι) is said to be a *primitively \mathfrak{D} -oriented order* if $\iota(\mathfrak{D}) = \mathcal{O}$.

2.3 On the Order Generated by two Quaternion Elements

Give two integral elements $\alpha_1, \alpha_2 \in B$ that does not commute, $\mathbb{Z}\langle\alpha_1, \alpha_2\rangle \subseteq B$ is an order, with discriminant given by the following proposition:

Proposition 2. [10, Chapter 7] Let \mathfrak{D}_i be quadratic orders equal to $\mathbb{Z}[\alpha_i]$ for $i = 1, 2$ such that α_1, α_2 are not commuting. Let $D_i = \text{disc } \mathfrak{D}_i$, $t_i = \text{tr}(\alpha_i)$ for $i \in \{1, 2\}$ and $s = \text{tr}(\alpha_1\alpha_2)$, then

$$\text{disc } \mathbb{Z}\langle\alpha_1, \alpha_2\rangle = (D_1D_2 - (t_1t_2 - 2s)^2)/4$$

3 Algorithms to Solve the Quaternion Embedding Problem

In this section, we present several algorithms to solve the quaternion embedding problem.

3.1 A First Algorithm for a Generic Order.

Our first algorithm `GenericOrderEmbedding` makes use of the formula stated in Proposition 2 on the discriminant of the quaternion order generated by two elements to produce a trace pairing oracle modulo the discriminant of the order \mathcal{O} .

Algorithm 1 `GenericOrderEmbedding`(\mathcal{O}, t, n)

Input: A quaternion order $\mathcal{O} \subset B_{p,\infty}$ of discriminant Δ , two integers $t, n \in \mathbb{Z}$ such that there exists an element of trace t and norm n in \mathcal{O} .

Output: \perp or $\alpha \in \mathcal{O}$ with $n(\alpha) = n$ and $\text{tr}(\alpha) = t$.

- 1: Compute a Minkowski reduced basis $1, \beta_1, \beta_2, \beta_3$ of \mathcal{O} .
 - 2: Compute $D_i = \text{tr}(\beta_i)^2 - 4n(\beta_i)$ for $1 \leq i \leq 3$, and $D = t^2 - 4n$.
 - 3: Compute s_i a square root of $DD_i \pmod{\Delta}$.
 - 4: **for** $t_1, t_2, t_3 \in [-\sqrt{4nn(\beta_1)}, \sqrt{4nn(\beta_1)}] \times [-\sqrt{4nn(\beta_2)}, \sqrt{4nn(\beta_2)}] \times [-\sqrt{4nn(\beta_3)}, \sqrt{4nn(\beta_3)}]$ such that $t_i = (1/2)(\pm s_i + t\text{tr}(\beta_i)) \pmod{\Delta}$ **do**
 - 5: Compute α the element such that $\text{tr}(\alpha) = t$, and $\text{tr}(\alpha\beta_i) = t_i$ for $1 \leq i \leq 3$.
 - 6: **if** $n(\alpha) = n$ **then**
 - 7: Return α .
 - 8: **end if**
 - 9: **end for**
 - 10: **return** Return \perp .
-

Proposition 3. Let $\mathcal{O} \subset B_{p,\infty}$ be a quaternion order of discriminant Δ (whose factorization is known) and Minkowski reduced basis $1, \beta_1, \beta_2, \beta_3$. Let t, n be two

integers such that there exists an element of norm n and trace t in \mathcal{O} , the `GenericOrderEmbedding` will output an element α in \mathcal{O} with the correct trace and norm and runs in

$$O\left(\left\lceil 8\frac{\sqrt{nn(\beta_1)}}{\Delta} \right\rceil \left\lceil 8\frac{\sqrt{nn(\beta_2)}}{\Delta} \right\rceil \left\lceil 8\frac{\sqrt{nn(\beta_3)}}{\Delta} \right\rceil \text{polylog}(\Delta n)\right)$$

Proof. Since $1, \beta_1, \beta_2, \beta_3$ is a basis of \mathcal{O} , any element $\alpha \in \mathcal{O}$ is uniquely determined by the values $\text{tr}(\alpha)$ and $\text{tr}(\alpha\beta_i)$ for $1 \leq i \leq 3$.

For any element α of \mathcal{O} , the quaternion order $\mathbb{Z}\langle\alpha, \beta_i\rangle$ is contained in \mathcal{O} and so its discriminant is divisible by Δ . Thus, with the formula given in Proposition 2, we get that if α has trace t and norm n , we must have $DD_i = (\text{ttr}(\beta_i) - 2\text{tr}(\alpha\beta_i))^2 \pmod{\Delta}$ which gives $\text{tr}(\alpha\beta_i) = \pm s_i + \text{ttr}(\beta_i) \pmod{\Delta}$ where $s_i^2 = DD_i \pmod{\Delta}$.

Moreover, since every quadratic order in the quaternion order \mathcal{O} has negative discriminant we must have $\text{tr}(\alpha\beta_i)^2 \leq 4nn(\beta_i)$.

Thus, assuming that there exists an element of norm n and trace t in \mathcal{O} , then there will be one triple of value t_1, t_2, t_3 that will lead to a correct element α .

For each $1 \leq i \leq 3$, there are less than $\left\lceil 8\frac{\sqrt{nn(\beta_i)}}{\Delta} \right\rceil$ values of t_i that satisfy the constraint $\pmod{\Delta}$ that are within the desired interval. When the factorization of Δ is known, it is possible to compute the square-root s_i in $O(\text{polylog}(\Delta))$ and all the operations to execute for each triple t_1, t_2, t_3 can be performed in $O(\text{polylog}(\Delta n))$. This proves the result.

Proposition 3 has three interesting corollaries. The first corollary states an asymptotic complexity of `GenericOrderEmbedding` when n is big compared to Δ .

Corollary 1. *Let \mathcal{O}, t, n be as in Proposition 3, and assume that $n > \Delta^2/64$. Then, the complexity of `GenericOrderEmbedding` is $O((n^{3/2}/\Delta^2)\text{polylog}(\Delta n))$.*

Proof. When $n > \Delta^2/64$, we have that $8\sqrt{nn(\beta_i)} > \Delta$ for all $1 \leq i \leq 3$, and so the asymptotics $\lceil 8\sqrt{nn(\beta_i)}/\Delta \rceil = \Theta(\sqrt{nn(\beta_i)}/\Delta)$ holds for any $1 \leq i \leq 3$. Then, we deduce the complexity of `GenericOrderEmbedding` from Proposition 1.

This second corollary identifies the situation where `GenericOrderEmbedding` will always be polynomial-time.

Corollary 2. *Let \mathcal{O}, t, n be as in Proposition 3. If $n = O(\Delta)$, then the complexity of `GenericOrderEmbedding` is $O(\text{polylog}(\Delta n))$.*

Proof. By Lemma 5, when $n = O(\Delta)$, $\sqrt{nn(\beta_i)}/\Delta = O(1)$ and the result follows from Proposition 3.

A direct application on maximal orders. We can obtain a first algorithm to solve Problem 1 on input \mathcal{O} , by applying directly `GenericOrderEmbedding` on the maximal order \mathcal{O} . In that case, $\Delta = p$, and Corollary 2 proves that our algorithm will be polynomial time when $n = O(p)$. This is already an improvement over

the result stated in [1] as it does not rely on any heuristic, but we expect `GenericOrderEmbedding` to be better than that in average.

In Corollary 3, we give a statement to quantify the number of maximal orders for which the running time of `GenericOrderEmbedding` is polynomial in term of the fraction $p^{4/3}/n$. The proof of Corollary 3 uses a bound on the number of maximal orders having a non-trivial endomorphism smaller than a given value m that we introduce below as Lemma 6. This result was proven in [14].

Lemma 6. *For any $0 < M < p^{2/3}$, the number of maximal order in $B_{p,\infty}$ containing an element not in \mathbb{Z} of norm smaller than M is $O(M^{3/2})$.*

Corollary 3. *Let $\mathcal{O} \subset B_{p,\infty}$ be a maximal order. Let t, n be as in Proposition 3. Assume further that the order \mathcal{O} is uniformly random among the set of maximal order types.*

There exists a polynomial $P(X) \in \mathbb{Q}[X]$, and constants C_1, C_2 such that for every $\varepsilon > 0$, if $n < C_1 p^{4/3-\varepsilon}$, then the running time of `GenericOrderEmbedding` on input \mathcal{O}, t, n is smaller than $P(\log(p))$ with probability bigger than $1 - C_2 p^{-3\varepsilon}$.

Proof. A maximal order in $B_{p,\infty}$ has discriminant $\Delta = p$.

If $n(\beta_3)n < p^2$, then Proposition 3 implies that the running time of `GenericOrderEmbedding` is poly-logarithmic in n, p and since $n = O < C_1 p^{4/3-\varepsilon}$ there exists a polynomial $P(X)$ such that the running time is smaller than $P(\log(p))$.

Thus, to prove the result, it suffices to prove that the probability of $n(\beta_3)n$ being bigger than p^2 is smaller than $C_2 p^{-3\varepsilon}$ for some constant C_2 .

Using Lemma 4, we can show that if $n(\beta_3)n \geq p^2$, then we must have $n(\beta_1) \leq C\sqrt{n} \leq CC_1 p^{2/3-2\varepsilon}$ for some $C > 0$. By Lemma 6, we know there exists $C' > 0$ such that there are at most $C' p^{1-3\varepsilon}$ maximal orders admitting a non-trivial β_1 of norm smaller than $CC_1 p^{2/3-\varepsilon}$. Since there are $O(p)$ distinct isomorphism classes of maximal orders, we conclude that the probability of finding such a bad maximal orders at random is smaller than $C_2 p^{-3\varepsilon}$ for some constant C_2 and this concludes the proof.

In Appendix A, we outline a heuristic variant of this algorithm, which works with any basis, followed by enumerating close vectors.

3.2 A Better Asymptotic Algorithm to Solve the Embedding Problem.

To solve the embedding problem, we are not restricted to the obvious solution of applying `GenericOrderEmbedding` on the maximal order given in input.

The goal of this section is to introduce another algorithm `MaximalOrderEmbeddingEichler` that applies `GenericOrderEmbedding` on Eichler orders. We will show that the average asymptotic complexity of this algorithm is the square-root of the asymptotic complexity of `GenericOrderEmbedding`. Unfortunately, despite that improvement, `MaximalOrderEmbeddingEichler` does not improve on the range of values of n for which the running time is polynomial.

The principle of `MaximalOrderEmbeddingEichler` is the following: by taking a split prime N in \mathfrak{D} , we can deduce that the element α we are looking for must be contained inside an Eichler order of level N contained in \mathcal{O} . Thus, we can compute the list of these orders and try to apply `GenericOrderEmbedding` on all of them until one works. We formalize this idea below as `MaximalOrderEmbeddingEichler`. The value of N is chosen to ensure that the expected running time of `GenericOrderEmbedding` on Eichler orders of level N (whose discriminant is pN) is polynomial in $\log(pN)$. This is why we take $N = O(n^{3/4}/p)$. In that case, we can expect the complexity of the algorithm to be $O(N) = O(n^{3/4}/p)$.

Algorithm 2 `MaximalOrderEmbeddingEichler`(\mathcal{O}, t, n)

Input: A maximal order $\mathcal{O} \subset B_{p,\infty}$, two integers $t, n \in \mathbb{Z}$ such that there exists an element of trace t and norm n in \mathcal{O} .

Output: \perp or $\alpha \in \mathcal{O}$ with $n(\alpha) = n$ and $\text{tr}(\alpha) = t$.

- 1: Set $D = t^2 - 4n$ and \mathfrak{D} as the maximal order in $\mathbb{Q}(\sqrt{-D})$.
 - 2: Select a prime N split in \mathfrak{D} such that $N/2 < n^{3/4}/p < N$.
 - 3: Compute $\mathcal{O}_1, \dots, \mathcal{O}_{N+1}$ the $N + 1$ Eichler orders of level N contained in \mathcal{O} .
 - 4: **for** $i = 1$ to M **do**
 - 5: Compute $\alpha = \text{GenericOrderEmbedding}(\mathcal{O}_i, t, n)$.
 - 6: **if** $\alpha \neq \perp$ **then**
 - 7: Return α .
 - 8: **end if**
 - 9: **end for**
 - 10: **return** \perp .
-

Despite the informal reasoning outlined above, it is not easy to prove formally what is the complexity of `MaximalOrderEmbeddingEichler` because there are Eichler orders of level N in \mathcal{O} that will have elements of norm smaller than expected. Nonetheless, we obtain a bound on the average running time by proving that executing `MaximalOrderEmbeddingEichler` on all maximal orders can be done in $O(n^{3/4+\varepsilon})$ for any $\varepsilon > 0$. This is stated in Proposition 4.

For the proof, we will need another corollary of Proposition 3 to bound the running time of `GenericOrderEmbedding` in terms of the norm of its successive minimas.

Corollary 4. *Let \mathcal{O}, t, n be as in Proposition 3, with $\beta_1, \beta_2, \beta_3$ the three successive minimas of \mathcal{O} .*

- (i) *When $n(\beta_2) < \Delta^2/64n$, the running time of `GenericOrderEmbedding` is*

$$O\left(\left\lceil 16\sqrt{2}\sqrt{\frac{n}{n(\beta_1)n(\beta_2)}} \right\rceil\right).$$
- (ii) *When $n(\beta_1) < \Delta^2/64n$ and $n(\beta_2) \geq \Delta^2/64n$, the running time of `GenericOrderEmbedding` is*

$$O\left(\frac{n}{\Delta\sqrt{n(\beta_1)}}\right).$$

Proof. For (i), when $n(\beta_2) \leq \Delta^2/64n$, then the first two factors in the complexity given in Proposition 3 are 1, and so the complexity is given by the last term.

From Proposition 1, we get $n(\beta_3) \leq 8\Delta^2/n(\beta_1)n(\beta_2)$ from which we derive

$$\lceil 8\sqrt{nn(\beta_3)}/\Delta \rceil = O(\lceil 16\sqrt{2}\sqrt{\frac{n}{n(\beta_1)n(\beta_2)}} \rceil)$$

For (ii), from $n(\beta_1) < \Delta^2/64n$, we get that the first factor in the complexity stated in Proposition 3 is 1. From $n(\beta_3) \geq n(\beta_2) \geq \Delta^2/64n$, we get that the complexity is

$$O\left(\frac{n\sqrt{n(\beta_2)n(\beta_3)}}{\Delta^2}\right)$$

which can be simplified to

$$O\left(\frac{n}{\Delta\sqrt{n(\beta_1)}}\right)$$

by applying $n(\beta_2)n(\beta_3) = O(\Delta^2/n(\beta_1))$ that we derive from Proposition 1. \square

We will also need a lemma to upper-bound the number of Eichler orders admitting an embedding of two non-commuting quadratic orders of discriminant δ_1, δ_2 .

Lemma 7. *Let \mathcal{O} be an Eichler order of level N in $B_{p,\infty}$. Let $\mathfrak{D}_1, \mathfrak{D}_2$ be two quadratic imaginary orders of discriminant δ_1, δ_2 (and conductors $f(\delta_1), f(\delta_2)$) such that the \mathfrak{D}_i are optimally embedded inside \mathcal{O} and their embedding is non-commuting.*

Let us take α_1 and α_2 two elements of \mathcal{O} such that optimal embedding of \mathfrak{D}_i is equal to $\mathbb{Z}[\alpha_i]$. Let $\mathcal{O}_{1,2}$ be the quaternion sub-order of \mathcal{O} generated by α_1, α_2 and let $s = \text{tr}(\alpha_1\alpha_2)$.

We define $T(s, \delta_1, \delta_2)$ as the number of Eichler orders of level N containing $\mathcal{O}_{1,2}$.

Then, there exists a constant C such that :

$$T(s, \delta_1, \delta_2) \leq C\tau(\Delta_{1,2}/\Delta)\tau(f(s, \delta_1, \delta_2))f(s, \delta_1, \delta_2) \quad (2)$$

where $\tau(x)$ counts the number of divisors of the integer x and $f(s, \delta_1, \delta_2)^2 = \text{gcd}(f(\delta_1)^2, f(\delta_2)^2, (1 - 2s))$ when $\text{tr}(\alpha_1)\text{tr}(\alpha_2) = 1$ and $\text{gcd}(f(\delta_1)^2, f(\delta_2)^2, s)$ otherwise.

Proof. Since α_1, α_2 are supposed to reach successive minimas, it is easy to see that their trace must be either 0 or 1 (if not then there would be an element of smaller than norm inside $\mathbb{Z} + \alpha_i$).

By [18, 24.1.4], there exists a unique integer $f(\mathcal{O}_{1,2})$ and Gorenstein order $\text{Gor}(\mathcal{O}_{1,2})$ such that $\mathcal{O}_{1,2} = \mathbb{Z} + f(\mathcal{O}_{1,2})\text{Gor}(\mathcal{O}_{1,2})$, where $f(\mathcal{O}_{1,2})$ is an integer and $\text{Gor}(\mathcal{O}_{1,2})$.

We start by showing that $f(\mathcal{O}_{1,2})^2 = \text{gcd}(f(\delta_1)^2, f(\delta_2)^2, (1 - 2s))$ when $\text{tr}(\alpha_1)\text{tr}(\alpha_2) = 1$ and $f(\mathcal{O}_{1,2})^2 = \text{gcd}(f(\delta_1)^2, f(\delta_2)^2, s)$ otherwise.

The number $f(\mathcal{O}_{1,2})$ divides all the coefficients of the ternary quadratic form associated to the trace 0 elements of $\mathcal{O}_{1,2}$ (see [18, 24.2]). In particular, this means that $f(\mathcal{O}_{1,2})$ divides the conductor of all imaginary quadratic orders contained in $\mathcal{O}_{1,2}$. This proves that $f(\mathcal{O}_{1,2})^2$ divides $\gcd(f(\delta_1)^2, f(\delta_2)^2)$.

When $\text{tr}(\alpha_i) = 1$, it is easy to verify that the conductor must be odd. When $\text{tr}(\alpha_1)\text{tr}(\alpha_2) = 1$ we know that $f(\mathcal{O}_{1,2})$ is odd. With $\text{disc } \mathbb{Z} + f\mathcal{O} = f^3 \text{disc } \mathcal{O}$ for any f, \mathcal{O} , and Proposition 2, we get that $f(\mathcal{O}_{1,2})^3 \mid (\text{tr}(\alpha_1)\text{tr}(\alpha_2) - 2s)^2$. Thus, $f(\mathcal{O}_{1,2}) \mid (\text{tr}(\alpha_1)\text{tr}(\alpha_2) - 2s)$.

Moreover, since $f(\mathcal{O}_{1,2})$ divides the conductor of $\mathbb{Z}[\alpha_1\alpha_2]$, $f(\mathcal{O}_{1,2})^2$ divides its discriminant which is equal to

$$\begin{aligned} s^2 - 4n(\alpha_1)n(\alpha_2) &= \frac{-\delta_1\delta_2 - 1 + \delta_2 + \delta_1 + 4s^2}{4} \\ &= \frac{-\delta_1\delta_2 + \delta_2 + \delta_1 + (2s-1)(2s+1)}{4} \end{aligned}$$

Since $f(\mathcal{O}_{1,2})^2$ divides δ_1, δ_2 it must divide $(2s-1)(2s+1)$ and since $f(\mathcal{O}_{1,2})$ is odd, it cannot divide $(2s+1)$ as it already divides $2s-1$. Thus, $f(\mathcal{O}_{1,2})^2$ divides $(2s-1)$.

When $\text{tr}(\alpha_1) = 1$ and $\text{tr}(\alpha_2) = 0$, we must have that $f(\mathcal{O}_{1,2})$ is odd and that $f(\mathcal{O}_{1,2})^2$ divides the conductor of $\alpha_1 + \alpha_2$. We have $\text{tr}(\alpha_1 + \alpha_2) = 1$ and $n(\alpha_1 + \alpha_2) = n(\alpha_1) + n(\alpha_2) + \text{tr}(\alpha_1\bar{\alpha}_2)$. Since $\text{tr}(\alpha_2) = 0$, $\bar{\alpha}_2 = -\alpha_2$ and so $n(\alpha_1 + \alpha_2) = n(\alpha_1) + n(\alpha_2) - s$. With $n(\alpha_1) = (1 - \delta_1)/4$ and $n(\alpha_2) = -\delta_2/4$, we get that the discriminant of $\alpha_1 + \alpha_2$ is $\delta_1 + \delta_2 - 4s$. Thus, we must have that $f(\mathcal{O}_{1,2})^2 \mid s$ which proves the result.

A similar reasoning proves the result when $\text{tr}(\alpha_1) = 0$ and $\text{tr}(\alpha_2) = 0$.

Now we need to prove that $\gcd(f(\delta_1)^2, f(\delta_2)^2, (\text{tr}(\alpha_1)\text{tr}(\alpha_2) - 2s))$ (resp. $\gcd(f(\delta_1)^2, f(\delta_2)^2, s)$) must divide $f(\mathcal{O}_{1,2})^2$.

$f(\mathcal{O}_{1,2})^2$ is the gcd of the norms of all the trace 0 element in $\mathcal{O}_{1,2}$. By expressing the ternary quadratic form corresponding to the norm of trace 0 elements given as a linear combinations of $1, \alpha_1, \alpha_2, \alpha_1\alpha_2$, it is easy to verify that $\gcd(f(\delta_1)^2, f(\delta_2)^2, (1 - 2s))$ (resp. $\gcd(f(\delta_1)^2, f(\delta_2)^2, s)$) when $\text{tr}(\alpha_1)\text{tr}(\alpha_2) = 1$ (resp. otherwise) divides the norm of all the elements of trace 0 in $\mathcal{O}_{1,2}$. This proves the result.

Then, we show that $\text{Gor}(\mathcal{O}_{1,2})$ is a Bass order. A Gorenstein order is Bass if all its superorder are Gorenstein. Thus, if $\text{Gor}(\mathcal{O}_{1,2})$ is not a Bass order, there exists a non-Gorenstein order $\mathcal{O}' = \mathbb{Z} + f(\mathcal{O}')\text{Gor}(\mathcal{O}')$ that contains $\text{Gor}\mathcal{O}_{1,2}$. Since \mathcal{O}' is non-Gorenstein, we must have $f(\mathcal{O}') > 1$. We can show that $\text{Gor}\mathcal{O}_{1,2}$ contains the quadratic imaginary orders of $\mathcal{O}_{1,2}$, but with a conductor divided by $f(\mathcal{O}_{1,2})$. We can follow the same reasoning we just led to prove that $f(\mathcal{O}')^2$ must divide $\gcd(f(\delta_1)^2, f(\delta_2)^2, (\text{tr}(\alpha_1)\text{tr}(\alpha_2) - 2s))/(f(\mathcal{O}_{1,2})^2)$ and this value is 1 which is a contradiction.

Thus $\text{Gor}\mathcal{O}_{1,2}$ is Bass.

Eichler and Brzezinski proved that the number of Eichler orders of discriminant Δ containing a given Bass suborder of discriminant $\Delta \mid D$ (in fact their

result is about the number of maximal orders containing some Bass sub-order, but it can easily be extended to Eichler orders of level N) is upper-bound by $\tau D/\Delta$ where τ is the function counting the number of divisors of any given number.

To conclude our proof, we just need a result to quantify the number of quaternion order of discriminant disc \mathcal{O} containing a given order of the form $\mathbb{Z} + f\mathcal{O}$ for any integer f and Bass order \mathcal{O} . Leroux [13, Lemma 3] provided such a result when f is prime and \mathcal{O} is a maximal order. We will adapt his proof to show that if $\mathbb{Z} + f\mathcal{O}$ is contained in \mathcal{O}' where \mathcal{O} and \mathcal{O}' are Bass orders of the same discriminant, then \mathcal{O} and \mathcal{O}' are connected with an primitive ideal of norm $f' \mid f$.

Let us consider the ideal $I = \{x \in \mathcal{O}', x\mathcal{O} \subset \mathcal{O}'\}$. It is easily verified that this is an integral ideal whose left order is \mathcal{O}' and right order is \mathcal{O} . We have $f\mathcal{O}' \subset I \subset \mathcal{O}'$, and so I is equal to $f_1 I'$ where f_1 and $n(I')$ divide f and I' is a primitive integral ideal whose left order is \mathcal{O}' and right order is \mathcal{O} .

Thus, we can bound the number \mathcal{O}' of orders containing $\mathbb{Z} + f\mathcal{O}$ with disc $\mathcal{O}' = \text{disc } \mathcal{O}$ by $C\tau(f)f$ for some constant C as the number of integral ideals of norm f is in $O(f)$.

We get the final results by multiplying the bound on the number of super-order containing the Bass order $\text{Gor}\mathcal{O}_{1,2}$ with $Cf(\mathcal{O}_{1,2})\tau(f_{\mathcal{O}_{1,2}})$.

□

Proposition 4. *Let p, t, n, \mathcal{O} be such that there exists an element of trace t and norm n inside \mathcal{O} and $n > p^{4/3}$, `MaximalOrderEmbeddingEichler` will output an element α of the correct trace and norm inside \mathcal{O} .*

For any $\varepsilon > 0$, the average complexity of `MaximalOrderEmbeddingEichler` is

$$O\left(\frac{n^{3/4+\varepsilon}}{p}\right).$$

Proof. Correctness follow from Proposition 3 and the fact that if N is split in \mathfrak{D} , then α of trace t and norm n , is contained in one of the two Eichler orders of the form $\mathbb{Z} + \mathcal{O}\mathfrak{N}$ where \mathfrak{N} is an \mathfrak{D} -ideal of norm N .

We are going to prove that the sequential executions of `MaximalOrderEmbeddingEichler` on all types of maximal orders in $B_{p,\infty}$ takes time $O(n^{3/4+\varepsilon})$. This will prove the result as there are $O(p)$ distinct maximal order types in $B_{p,\infty}$.

Since the values t, n are always the same, the value of N can be the same across all executions of `MaximalOrderEmbeddingEichler`. In that case, the sequential executions of `MaximalOrderEmbeddingEichler` on all maximal orders types simply consist in the computation of all Eichler orders of level N , and the sequential executions of `GenericOrderEmbedding` on all these orders.

The integer N is prime, so there are $O(Np)$ Eichlers orders of level N , and each one can be computed in $O(\text{polylog}(pN))$ by enumerating ideals of norm N and intersecting their left and right orders. With the choice of N , the cost of computing them all is $O(n^{3/4}\text{polylog}(pn))$.

Let us write $\mathfrak{S}_{N,p}$ the set of all Eichler orders of level N in $B_{p,\infty}$ and let us write $pN = \Delta$ the discriminant of these orders. For each $\mathcal{O} \in \mathfrak{S}_{N,p}$, the cost

of executing `GenericOrderEmbedding` on input \mathcal{O}, t, n is written $C_{\mathcal{O}}$. We write $n_1^{\mathcal{O}}, n_2^{\mathcal{O}}, n_3^{\mathcal{O}}$ the norm of the successive minima. Corollary 4 proves that there exists a function $C : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that $C_{\mathcal{O}} = O(C(\Delta, n, n_1^{\mathcal{O}}, n_2^{\mathcal{O}}))$.

By Lemma 1, we have that $n_1^{\mathcal{O}} \leq 2\Delta^{2/3}$, and by Lemmas 2 and 3 we have that $\max(\Delta/(4n_1^{\mathcal{O}}), n_1^{\mathcal{O}}) \leq n_2^{\mathcal{O}} \leq 2\sqrt{2}\Delta/\sqrt{n_1^{\mathcal{O}}}$.

Now, let us define $\delta_i^{\mathcal{O}}$ as the discriminant of $\mathbb{Z}[\beta_i^{\mathcal{O}}]$. Its value is $((\varepsilon_i^{\mathcal{O}})^2 - 4n_i^{\mathcal{O}})$ where $\varepsilon_i^{\mathcal{O}} = \text{tr}(\beta_i^{\mathcal{O}})$ is a value in $\{0, 1\}$. In particular, we have $4n_i^{\mathcal{O}} - 1 \leq -\delta_i^{\mathcal{O}} \leq 4n_i^{\mathcal{O}}$. In that case, note that we also have $C_{\mathcal{O}} = O(C(\Delta, n, -\delta_1^{\mathcal{O}}, -\delta_2^{\mathcal{O}}))$

If we write $T(\delta_1, \delta_2) = \#\{\mathcal{O} \in \mathfrak{S}_{N,p} \mid \delta_1^{\mathcal{O}} = \delta_1, \delta_2^{\mathcal{O}} = n_2\}$, and we regroup maximal orders by the discriminants corresponding to their first and second successive minimas, we can get

$$\sum_{\mathcal{O} \in \mathfrak{S}_{N,p}} C_{\mathcal{O}} \leq C_1 \sum_{-\delta_1=3}^{\lceil 8\Delta^{2/3} \rceil} \sum_{-\delta_2=\lfloor \max(4\Delta/(-\delta_1), -\delta_1) \rfloor}^{\lceil \sqrt{2}\Delta/\sqrt{-\delta_1} \rceil} T(\delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) \quad (3)$$

for some constant C_1 .

When $\delta_2 \geq \Delta^2/(16n) - 1$, the bound (ii) from Corollary 4 yields

$$\begin{aligned} & \sum_{-\delta_1=3}^{\lceil 8\Delta^{2/3} \rceil} \sum_{-\delta_2=\lceil \Delta^2/(4n) \rceil}^{\lceil \sqrt{2}\Delta/\sqrt{-\delta_1} \rceil} T(\delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) \\ & \leq C_2 \sum_{-\delta_1=3}^{\lceil 8\Delta^{2/3} \rceil} \frac{n}{\Delta\sqrt{-\delta_1}} \sum_{-\delta_2=\lceil \Delta^2/(n) \rceil}^{\lceil \sqrt{2}\Delta/\sqrt{-\delta_1} \rceil} T(\delta_1, \delta_2) \end{aligned}$$

for some constant C_2 .

We have an optimal embedding of the quadratic order of discriminant δ_1 inside every Eichler order such that $\delta_1^{\mathcal{O}} = \delta_1$. Let us write \mathfrak{D}_1 for this quadratic order.

Each optimal embedding of \mathfrak{D}_1 inside an Eichler order \mathcal{O} of level N gives an optimal embedding of \mathfrak{D}_1 in the two maximal super-orders of \mathcal{O} . There are $O(h(\mathfrak{D}_1))$ distinct optimal embeddings of \mathfrak{D}_1 inside maximal orders (see [15, Proposition 3.3] for instance) and it can be shown that each of these embeddings gives an embedding of \mathfrak{D}_1 in at most 2 Eichler orders of level N (corresponding to the at most 2 \mathfrak{D}_1 -ideals of norm N).

Thus, there are $O(h(\mathfrak{D}_1))$ distinct types of Eichler orders of level N with $\delta_1^{\mathcal{O}} = \delta_1$ and we deduce that $\sum_{-\delta_2=\lceil \Delta^2/(4n) \rceil}^{\lceil \sqrt{2}\Delta/\sqrt{-\delta_1} \rceil} T(\delta_1, \delta_2) = O(h(\mathfrak{D}_1)) = O((-\delta_1)^{1/2+\varepsilon})$.

With $\Delta = pN = O(n^{3/4})$, we deduce

$$\sum_{-\delta_1=3}^{\lceil 8\Delta^{2/3} \rceil} \sum_{-\delta_2=\lceil \Delta^2/(16n) \rceil}^{\lceil \sqrt{2}\Delta/\sqrt{-\delta_1} \rceil} T(\delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) = O(n\Delta^{-1/3+\varepsilon}) = O(n^{3/4+\varepsilon}) \quad (4)$$

For every Eichler order \mathcal{O} containing an embedding of the quadratic orders of discriminant δ_1, δ_2 , Proposition 2 tells us that their must be a value $s = (1/2) \pm \sqrt{\delta_1 \delta_2} + \varepsilon \pmod{\Delta}$ where $|s| \leq \sqrt{\delta_1 \delta_2}$. In that case, the value of $\delta_1 \delta_2 = (2s - \varepsilon)^2 + k\Delta$ for some integer $0 \leq k$.

Thus we can upper-bound the second part of our sum as follows:

$$\begin{aligned} & \sum_{\substack{[8\Delta^{2/3}] \\ -\delta_1=3}} \sum_{\substack{[\Delta^2/(16n)] \\ -\delta_2=\lfloor \max(4\Delta/(-\delta_1), -\delta_1) \rfloor}} T(\delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) \\ & \leq \sum_{|s| \leq 2\Delta/\sqrt{n}} \sum_{k=0}^{[\Delta^{5/3}/n]} \sum_{\delta_1, \delta_2 \in \{\delta_1, \delta_2 \mid \delta_1 \delta_2, (2s-\varepsilon)^2 + k\Delta\}} T(s, \delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) \end{aligned}$$

where $T(s, \delta_1, \delta_2)$ was defined in Lemma 7.

Now, we can apply Corollary 4 (i) to get

$$\begin{aligned} & \sum_{|s| \leq 2\Delta/\sqrt{n}} \sum_{k=0}^{[\Delta^{5/3}/n]} \sum_{\delta_1, \delta_2 \in \{\delta_1, \delta_2 \mid \delta_1 \delta_2, (2s-\varepsilon)^2 + k\Delta\}} T(s, \delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) \\ & \leq \sqrt{n} \sum_{|s| \leq 2\Delta/\sqrt{n}} \sum_{k=0}^{[\Delta^{5/3}/n]} \sum_{\delta_1, \delta_2 \in \{\delta_1, \delta_2 \mid \delta_1 \delta_2, (2s-\varepsilon)^2 + k\Delta\}} \frac{T(s, \delta_1, \delta_2)}{\sqrt{\delta_1 \delta_2}} \end{aligned}$$

We have $\delta_1 \delta_2 > s^2$, thus $1/\sqrt{\delta_1 \delta_2} \leq 1/|s|$. Moreover, we can apply Lemma 7 to upper-bound $T(s, \delta_1, \delta_2)$. This yields

$$\begin{aligned} & \sum_{|s| \leq 2\Delta/\sqrt{n}} \sum_{k=0}^{[\Delta^{5/3}/n]} \sum_{\delta_1, \delta_2 \in \{\delta_1, \delta_2 \mid \delta_1 \delta_2, (2s-\varepsilon)^2 + k\Delta\}} \frac{T(s, \delta_1, \delta_2)}{\sqrt{\delta_1 \delta_2}} \\ & \leq \sqrt{n} \max_{x \leq n^m} \tau(N)^2 \sum_{|s| \leq 2\Delta/\sqrt{n}} \frac{1}{|s|} \sum_{k=0}^{[\Delta^{5/3}/n]} \sum_{\delta_1, \delta_2 \in \{\delta_1, \delta_2 \mid \delta_1 \delta_2, (2s-\varepsilon)^2 + k\Delta\}} f(s, \delta_1, \delta_2) \end{aligned}$$

for some constant C_3 and integer $m > 0$. We define $\tau(x)$ to be the number of distinct divisor of any integer x .

The size of the set $\{\delta_1, \delta_2 \mid \delta_1 \delta_2, (2s - \varepsilon)^2 + k\Delta\}$ can be upper-bounded by $\tau((2s - \varepsilon)^2 + k\Delta)^2$.

By definition of $f(s, \delta_1, \delta_2)$ in Lemma 7, we see that we must have $f(s, \delta_1, \delta_2)^4 \mid k$. Thus, by writing every value k as $k_0^4 k_1$ we can upper bound $f(s, \delta_1, \delta_2)^4$ by

k_0 , and we obtain :

$$\begin{aligned} & \sqrt{n} \max_{x \leq n^m} \tau(N)^2 \sum_{|s| \leq 2\Delta/\sqrt{n}} \frac{1}{|s|} \sum_{k=0}^{\lceil \Delta^{5/3}/n \rceil} \sum_{\delta_1, \delta_2 \in \{\delta_1, \delta_2 | \delta_1 \delta_2, (2s-\varepsilon)^2 + k\Delta\}} f(s, \delta_1, \delta_2) \\ & \leq \sqrt{n} \max_{x \leq n^m} \tau(N)^4 \sum_{|s| \leq 2\Delta/\sqrt{n}} \frac{1}{|s|} \sum_{k_0=1}^{\lceil (\Delta^{5/3}/n)^{1/4} \rceil} \sum_{k_1=0}^{\lfloor \Delta^{5/3}/(nk_0^4) \rfloor} k_0 \end{aligned}$$

There exists a constant C_4 such that

$$\sum_{k_0=1}^{\lceil (\Delta^{5/3}/n)^{1/4} \rceil} \sum_{k_1=0}^{\lfloor \Delta^{5/3}/(nk_0^4) \rfloor} k_0 \leq C_4 \Delta^{5/3}/n \sum_{k_0=1}^{\lceil (\Delta^{5/3}/n)^{1/4} \rceil} \frac{1}{k_0^3}.$$

The value of $\zeta(3)$ is a constant and $\sum_{s=1}^x 1/s = O(\log x)$. Thus, there is a constant C_5 such that

$$\begin{aligned} & \sqrt{n} \max_{x \leq n^m} \tau(N)^4 \sum_{|s| \leq 2\Delta/\sqrt{n}} \frac{1}{|s|} \sum_{k_0=1}^{\lceil (\Delta^{5/3}/n)^{1/4} \rceil} \sum_{k_1=0}^{\lfloor \Delta^{5/3}/(nk_0^4) \rfloor} k_0 \\ & \leq C_5 \frac{\Delta^{5/3}}{\sqrt{n}} \log(\Delta/\sqrt{n}) \max_{x \leq n^m} \tau(N)^4 \end{aligned}$$

With $\Delta = pN = O(n^{3/4})$ and the fact that $\tau(x) = O(x^\varepsilon)$ for any $\varepsilon > 0$ [20], we conclude that

$$\sum_{-\delta_1=3}^{\lceil 8\Delta^{2/3} \rceil} \sum_{-\delta_2=\lfloor \max(4\Delta/(-\delta_1), -\delta_1) \rfloor}^{\lceil \Delta^2/(16n) \rceil} T(\delta_1, \delta_2) C(\Delta, n, -\delta_1, -\delta_2) = O(n^{3/4+\varepsilon}) \quad (5)$$

The combination of Eqs. (3) to (5) proves that executing `MaximalOrderEmbeddingEichler` on all maximal orders takes time $O(n^{3/4+\varepsilon})$ and this proves that the average running time is $O(n^{3/4+\varepsilon})/p$. \square

3.3 Another Heuristic Algorithm with Factorization

The problem with `GenericOrderEmbedding` is that it does not work well when the input order \mathcal{O} contains smaller elements that one should expect from a random order of the same discriminant. Thus, while being efficient in the average case, it is not always optimal. Interestingly, we will see that the bad cases for `GenericOrderEmbedding` are actually good cases for another algorithm that we present below as `GenericOrderEmbeddingFactorization`.

The idea of this algorithm is that since \mathcal{O} contains a very small element β_1 , it will be easier to know the value of $\text{tr}(\alpha\beta_1)$ exactly. Then, once this value is

fixed, the ternary quadratic form becomes a binary quadratic form that we know how to solve efficiently.

More precisely, let β be an element in \mathcal{O} for which we know $\text{tr}(\alpha\beta)$, and let γ be any element in \mathcal{O} orthogonal to $\mathbb{Z}[\beta]$. Now look at the order

$$\mathbb{Z}\langle\beta, \gamma\rangle \subseteq \mathcal{O}$$

and write M for the index $[\mathbb{Z}\langle\beta, \gamma\rangle : \mathcal{O}]$. Writing $x + \beta y + \gamma z + \gamma\beta w$ for a generic element in $\mathbb{Z}\langle\beta, \gamma\rangle$, the norm form of this order is of the simple form

$$Q(x, y, z, w) := f(x, y) + n(\gamma)f(z, w)$$

where $f(x, y)$ denotes the norm of $x + \beta y$. While it is unlikely that α lies in $\mathbb{Z}\langle\beta, \gamma\rangle$, we have that $M\alpha \in \mathbb{Z}\langle\beta, \gamma\rangle$, thus, we can instead solve for $M\alpha$. First, we find the values of x and y from the knowledge of $\text{tr}(\alpha)$ and $\text{tr}(\alpha\beta)$ (because z, w contribute nothing to these traces). Then, we can solve for z, w by enumerating all solutions of

$$f(z, w) = \frac{n(M\alpha) - n(\alpha_0)}{n(\gamma)}.$$

with Cornacchia's algorithm. Finally, for each potential solution of the form $\alpha' := x + \beta y + \gamma z + \gamma\beta w$, we check if $\alpha'/M \in \mathcal{O}$.

The only caveat is that Cornacchia's algorithm require the factorization of the number one is trying to represent. Furthermore, the total amount of solutions is exponential in the number of distinct prime factors of the number one is trying to represent. Thus, the best we can do, as we need to enumerate through all the solutions of each Cornacchia instance, is get a heuristic runtime for our algorithm, under the plausible assumption that the integers that we encounter will not have too many prime factors.

Heuristic 1 *All integers M (resp. N) occuring in Step 4 (resp. Step 8) of Algorithm 3 behave like four times random numbers (resp. random numbers). In particular, the number of distinct prime factors is expected to be small, i.e. $O(\text{poly}(\log \log M))$ (resp. $O(\text{poly}(\log \log N))$).*

We also introduce a second heuristic to estimate the number of expected embedding of a given quadratic order in any maximal order. This heuristic will be useful in both the proof of this algorithm and later.

Heuristic 2 *Let $\mathcal{O} \subseteq B_{p, \infty}$ be a maximal quaternion order, and let \mathfrak{D} be a quadratic order, embedding into $B_{p, \infty}$. The expected number of optimal embeddings $\iota : \mathfrak{D} \hookrightarrow \mathcal{O}$ up to conjugation by \mathcal{O}^\times is $\Theta(h(\mathfrak{D})/p)$.*

One reasoning for this heuristic comes from [18, Theorem 30.7.5], which, specialized to our case, say that summing over a representative of all isomorphism classes of maximal orders in $B_{p, \infty}$, there should be $\Theta(h(\mathfrak{D}))$ embeddings. Heuristic 2 simply says that these embeddings are randomly distributed over these representatives. In [12], Leroux proved some bounds on the number of distinct embeddings of the same quadratic order inside the same maximal order.

While these bounds are not enough to prove Heuristic 2, they are a first step in the right direction as they prove that extreme situations where all quadratic orders are embedded inside the same maximal order are not possible.

Algorithm 3 GenericOrderEmbeddingFactorization(\mathcal{O}, t, n)

Input: An order $\mathcal{O} \subset B_{p, \infty}$ of discrd $\mathcal{O} = \Delta$ (with known factorization), and two integers $t, n \in \mathbb{Z}$.

Output: \perp , or $\alpha \in \mathcal{O}$ with $n(\alpha) = n$ and $\text{tr}(\alpha) = t$.

- 1: Compute a Minkowski reduced basis $1, \beta_1, \beta_2, \beta_3$ of \mathcal{O} .
 - 2: Compute $D_1 = \text{tr}(\beta_1)^2 - 4n(\beta_1)$, and $D = t^2 - 4n$.
 - 3: Compute γ_1, γ_2 , a Minkowski reduced basis of the part of \mathcal{O} orthogonal to $\mathbb{Z}[\beta]$.
 - 4: Compute $M := [\mathbb{Z}\langle \beta_1, \gamma \rangle : \mathcal{O}]$.
 - 5: Compute s_1 a square root of $DD_1 \pmod{\Delta}$.
 - 6: **for** $t_1 \in [1, \sqrt{4nn(\beta_1)}]$ such that $t_1 = (1/2)(\pm s_1 + t\text{tr}(\beta_1)) \pmod{\Delta}$ **do**
 - 7: Let $\alpha'_0 := x + y\beta_1$ be the element in $\mathbb{Z}[\beta_1]$ with $\text{tr}(\alpha'_0) = Mt, \text{tr}(\alpha'_0\beta_1) = Mt_1$
 - 8: Set $N := \frac{M^2n - n(\alpha'_0)}{Mn(\gamma_1)}$
 - 9: **for** z, w such that $n(z + \beta_1w) = MN$ **do**
 - 10: Set $\alpha' := \alpha'_0 + \gamma(z + \beta_1w)$
 - 11: **if** $\alpha'/M \in \mathcal{O}$ **then**
 - 12: **return** α'/M
 - 13: **end if**
 - 14: **end for**
 - 15: **end for**
 - 16: **return** Return \perp .
-

Proposition 5. *Assume the existence of a factorization oracle, and that Heuristic 1 holds. Given integers t and n , GenericOrderEmbeddingFactorization outputs an element $\alpha \in \mathcal{O}$ with $\text{tr}(\alpha) = t$ and $n(\alpha) = n$ or decide that none exists in time*

$$O\left(\frac{\sqrt{nn(\beta_1)}}{\Delta} \cdot \text{polylog}(np)\right)$$

where $\Delta = \text{disc } \mathcal{O}$. Further, assuming \mathcal{O} is maximal, and that Heuristic 2 holds, the expected runtime is also upper bounded by

$$O\left(\sqrt{n(\beta_1)} \cdot \text{polylog}(np)\right) \subseteq O\left(p^{1/3} \cdot \text{polylog}(np)\right)$$

Proof. The correctness of the algorithm follows directly from the description at the start of this section. We now proceed to prove the runtime of the algorithm.

The algorithm tries the $O\left(\frac{\sqrt{nn(\beta_1)}}{\Delta}\right)$ possible values of t_1 , and for each one, attempts to derive a solution α from representations of some integer MN by the principal binary quadratic form corresponding to elements in $\mathbb{Z}[\beta_1]$.

Under Heuristic 1, Cornacchia's algorithm can find the $O(\text{polylog}(MN))$ solutions in $O(\text{polylog}(MN))$ time. Testing each candidate has the same complexity and this proves the first part of the result.

However, the runtime above is the same as the time it takes to find all solutions. Applying Heuristic 2, we expect there to be a total of

$$O(h(\mathbb{Z}[\alpha])/p) = O(\sqrt{n}/p)$$

solutions. By Heuristic 1, each value of t_1 is only expected to give a polylogarithmic number of solutions, hence the total number of values t_1 that corresponds to a value of $\text{tr}(\alpha\beta_1)$ for a solution α , divided by the total possible number of values of t_1 , is

$$\tilde{O}\left(\frac{\sqrt{n}/p}{\sqrt{nm(\beta_1)}/p}\right) = \tilde{O}\left(\frac{1}{\sqrt{n(\beta_1)}}\right)$$

This bounds the expected number of values of t_1 we have to try before a solution will be found. The final expected runtime is obtained by the bound on $n(\beta_1)$ given by Lemma 1 \square

From Proposition 5, we see that as n increases, the algorithm's runtime eventually becomes independent of n . In the cases when n is very large, we can also discard the factorization oracle altogether, by only running Cornacchia on “easy” instances (for instance when N is a prime number). Indeed, under Heuristic 1, the numbers N in Step 8 behave like random integers of the same size and so they have a probability of $1/\log N$ to be prime. This allow us to run `GenericOrderEmbeddingFactorization` without the need of a factorization oracle, and the running time is only increased by a factor $O(\log(np))$.

Remark 1. Our algorithm `GenericOrderEmbeddingFactorization` can be seen as a generalization of the method introduced in [1]. Indeed, what is done in [1, Algorithm 5.1] is equivalent to looking at $\text{tr}(N_0\omega_0\alpha)$ where ω_0 is an integral element of very small norm, and the integer N_0 is such that \mathcal{O} is connected with \mathcal{O}_0 , a maximal order containing ω_0 by an ideal of norm N_0 . Since one can expect $N_0 \approx \sqrt{p}$ when \mathcal{O} is a random maximal order, this method allows us to recover α in polynomial time when $n = O(p)$. In the case where N_0 is especially small, $D\omega_0$ might be equal to β_1 and in that case, our method is equivalent to the one of [1]. Note that in every other case, our method is strictly better. Also, note that, by replacing β_1 by other elements of small norm, we can perform a similar randomization as was explained in [1, Section 5.3], to remove the need for the factorization oracle. However, this rerandomization may not help in some cases where β_1 is much smaller than β_2, β_3 , because in that case all the small vectors will lie in the same quadratic order generated by 1 and β_1 .

4 Ideals Between Oriented Orders

In this section, we expand on results related to primitively oriented maximal orders. We do this by first considering ideals between oriented maximal orders, and show that such ideals “comes from” quadratic ideals precisely when the left and right order of the ideal admits the same orientation.

Our first result is then a new algorithm that reduces \mathfrak{D} -vectorisation to \mathfrak{D} -EndRing in polynomial time for all orders. One such reduction first appeared in the work by Castryck, Vercauteren and Panny [4] for the order $\mathbb{Z}[\sqrt{-p}]$, and this was later generalised Wesolowski [19] to arbitrary orders. However, the algorithm is only polynomial in $\#\text{Cl}(\mathfrak{D})[2]$ (which can be exponential in the discriminant of \mathfrak{D}), and requires the factorization of disc \mathfrak{D} . Our reduction does not have this caveat, and only depends on the size of the discriminant, not its factorization.

Second, we consider the problem of finding ideals of fixed degree between isomorphism classes of quaternion orders, a problem of huge importance in isogeny-based cryptography. For large norm, this is solved efficiently by the (generalised) KLPT algorithm [11] [6], while for small degree, this is efficiently solvable by simple lattice reduction. The remaining sizes of norms in between here were recently studied by Benjamin Bencina, Péter Kutas, Simon-Philipp Merz, Christophe Petit, Miha Stopar and Charlotte Weitkämper [2]. The relation to the quaternion embedding problem was mentioned in the same work [2, Appendix A]. We expand on this connection, giving a heuristic algorithm which solves this problem in time $O(p^{2/3})$ for any degree d , and we show that in the special case where both orders are oriented orders of small class number, this algorithm is polynomial time for any d , allowing the computation of optimal paths between supersingular curves with small endomorphisms. Finally, in Appendix C, we consider the case where one of the orders contain an element of very small norm, and show that this can be solved in polynomial time up to $d < p^{2/3}$.

Given an \mathfrak{D} -ideal \mathfrak{l} , and an primitively \mathfrak{D} -oriented order (\mathcal{O}, ι) (Definition 1) we can define the corresponding quaternion ideal as $\mathcal{O}\langle\iota(\mathfrak{l})\rangle$. Further, given an \mathcal{O} -ideal I , one can define the corresponding \mathfrak{D} -ideal to be $\iota^{-1}(I)$, which can be computed by intersecting I with $\iota(K)$. The relation between these operations are given by the following proposition.

Proposition 6. *Let (\mathcal{O}, ι) be a primitively \mathfrak{D} -oriented order. Then*

- Given a left \mathcal{O} -ideal I , we have that $\mathcal{O}\langle\mathfrak{n}(I)\rangle \subseteq \mathcal{O}\langle I \cap \iota(K) \rangle \subseteq I$.
- Given an invertible \mathfrak{D} -ideal \mathfrak{l} , we have that $\mathcal{O}\langle\iota(\mathfrak{l})\rangle \cap \iota(K) = \iota(\mathfrak{l})$

Proof. To prove the first statement, note that the first inequality follows from the fact that $\mathfrak{n}(I)\mathbb{Z} \subseteq I \cap \iota(\mathfrak{D})$, and the second follows from the observation that $\mathcal{O}\langle I \cap \iota(\mathfrak{D}) \rangle \subseteq I$.

To prove the second statement, following [18, Exercise 30.2.a], we see that $\mathcal{O}\langle\iota(\mathfrak{l})\rangle \cap \iota(K) \supset \iota(\mathfrak{l})$, since $1 \in \mathcal{O}$, and conversely, since \mathfrak{l} is invertible, we have find that

$$\begin{aligned} (\mathcal{O}\langle\iota(\mathfrak{l})\rangle \cap \iota(K))\iota(\mathfrak{D}) &= (\mathcal{O}\langle\iota(\mathfrak{l})\rangle \cap \iota(K))\iota(\Gamma^{-1}\mathfrak{l}) \\ &\subseteq (\mathcal{O}\langle\iota(\mathfrak{l})\iota(\Gamma^{-1})\rangle \cap \iota(K))\iota(\mathfrak{l}) = \iota(\mathfrak{D})\mathfrak{l} \end{aligned}$$

where we are using $\mathcal{O}\langle\iota(\mathfrak{l})\iota(\Gamma^{-1})\rangle \cap \iota(K) = \mathcal{O} \cap \iota(K) = \iota(\mathfrak{D})$, which follows by definition of (\mathcal{O}, ι) being primitively \mathfrak{D} -oriented. \square

The previous proposition motivates the following definition, which emphasises when a quaternion ideal is generated by the image of a quadratic ideal:

Definition 2. Let (\mathcal{O}, ι) be a primitively \mathfrak{D} -oriented maximal order. A left \mathcal{O} -ideal is said to be generated by an \mathfrak{D} -ideal if

$$I = \mathcal{O}\langle I \cap \iota(K) \rangle$$

The following lemma shows that the orientation automatically “transfer” to the right order of an ideal generated by an \mathfrak{D} -ideal.

Lemma 8. Let (\mathcal{O}, ι) be a primitively \mathfrak{D} -oriented maximal order, and let I be a left \mathcal{O} -ideal, generated by an \mathfrak{D} -ideal. Then $(\mathcal{O}_R(I), \iota)$ is a (not necessarily primitively) \mathfrak{D} -oriented maximal order.

Proof. Let ω be the image of a generator of \mathfrak{D} under ι . To prove that $(\mathcal{O}_R(I), \iota)$ is a \mathfrak{D} -oriented maximal order, it suffices to see that $\omega \in \mathcal{O}_R(I)$. But this follows from the fact that I can be given generators in $\iota(K)$, which commute with ω . \square

When we have a primitively \mathfrak{D} -oriented maximal order (\mathcal{O}, ι) , the previous lemma showed that given an \mathcal{O} -ideal I , the right order of I admitting the same orientation is a necessary condition for I to be generated by an \mathfrak{D} -ideal. Next, we show that this condition is also sufficient.

Lemma 9. Let (\mathcal{O}_1, ι) and (\mathcal{O}_2, ι) be primitively \mathfrak{D} -oriented maximal orders, and let I be a (not necessarily primitively) \mathfrak{D} -oriented maximal order. Then their connecting ideal I is generated by an \mathfrak{D} -ideal.

Proof. Let ω be a generator of \mathfrak{D} under ι , and let I be the unique primitive connecting ideal between $\mathcal{O}_1, \mathcal{O}_2$. We have that $\omega \in \mathcal{O}_1 \cap \mathcal{O}_2 = \mathbb{Z} + I$, and hence, $a + \omega \in I$ for some $a \in \mathbb{Z}$. Let

$$J = \mathcal{O}_1\langle a + \omega, \mathfrak{n}(I) \rangle.$$

Clearly, J is generated by an \mathfrak{D} -ideal, and we will show that $I = J$. First, note that $\mathcal{O}\langle \mathfrak{n}(I) \rangle \subseteq J \subseteq I$, so assume $\mathfrak{n}(J) = \mathfrak{n}(I)d$ for some $d \mid \mathfrak{n}(I)$. Assume now that $J \subsetneq I$, i.e. that $d \neq 1$. Since $\mathcal{O}_1\langle \mathfrak{n}(I) \rangle \subseteq J$, we have $\mathcal{O}_1\langle d \rangle \subset J + \mathcal{O}_1\langle d \rangle$. Since we have $\mathfrak{n}(J + \mathcal{O}_1\langle d \rangle) = \gcd(\mathfrak{n}(J), d^2)$, we see that we must have $\mathfrak{n}(J + \mathcal{O}_1\langle d \rangle) = d^2$. By equality of the norm, we must have $J + \mathcal{O}_1\langle d \rangle = \mathcal{O}_1\langle d \rangle$. Hence, $J/\langle d \rangle \subseteq \mathcal{O}_1$, implying that $\frac{a+\omega}{d} \in \mathcal{O}_1$, contradicting the assumption that \mathcal{O}_1 was primitively \mathfrak{D} -oriented. Hence $I = J$, which shows that I is generated by an \mathfrak{D} -ideal. \square

4.1 Vectorisation to Oriented Endring Reduction

From Lemma 9, we see that the only obstruction in finding an ideal generated by an \mathfrak{D} -ideal between two primitively oriented maximal orders is that the two \mathfrak{D} -oriented orders might be oriented in different ways. Fortunately, the following lemma shows that this is easy to fix.

Lemma 10. Let (\mathcal{O}, ι_1) be a primitively \mathfrak{D} -oriented maximal order, and let $\iota_2 : K \hookrightarrow B$ be another embedding. Then there exists an order $\mathcal{O}' \cong \mathcal{O}$ such that (\mathcal{O}', ι_2) is a primitively \mathfrak{D} -oriented maximal order

Proof. By the Skolem-Noether theorem, given any two embeddings $\iota_1, \iota_2 : K \hookrightarrow B$, there exists some $\alpha \in B^\times$ such that for any $\delta \in K$, we have that $\iota_1(\delta) = \alpha^{-1}\iota_2(\delta)\alpha$. Then $\mathcal{O}' := \alpha\mathcal{O}\alpha^{-1}$ is isomorphic to \mathcal{O} , and further it is clear that

$$\iota_2(K) \cap \mathcal{O}' = \iota_1(K) \cap \alpha\mathcal{O}\alpha^{-1} = \alpha^{-1}\iota_2(K)\alpha \cap \mathcal{O} = \iota_1(K) \cap \mathcal{O} = \iota(\mathfrak{D}),$$

hence (\mathcal{O}', ι_2) is a primitively oriented maximal order. \square

Thus, the reduction simply consists of fixing the orientations, and intersecting with the image of K under the orientation. We summarize this in Algorithm 4.

Algorithm 4 Vectorization $_{\mathfrak{D}}((\mathcal{O}_1, \iota_1), (\mathcal{O}_2, \iota_2))$

Input: Two primitively \mathfrak{D} -oriented maximal orders $(\mathcal{O}_1, \iota_1), (\mathcal{O}_2, \iota_2)$.

Output: An \mathfrak{D} -ideal \mathfrak{l} such that $\mathcal{O}_R(\mathcal{O}_1\langle\mathfrak{l}\rangle) \cong \mathcal{O}_2$.

Set $\omega_i := \iota_i(\omega)$ for $i = 1, 2$, where ω is any generator of \mathfrak{D} .

Compute α such that $\alpha\omega_1 - \omega_2\alpha = 0$ using linear algebra.

Set $\mathcal{O}'_2 := \alpha\mathcal{O}_2\alpha^{-1}$.

Compute the connecting $(\mathcal{O}_1, \mathcal{O}'_2)$ -ideal $I := N\mathcal{O}_1\mathcal{O}'_2$.

Set $\mathfrak{l} := \iota^{-1}(I \cap \iota(\mathfrak{D}))$.

return \mathfrak{l} .

Proposition 7. *Algorithm 4 is correct and runs in polynomial time in the length of the input.*

Proof. Lemma 10 shows both the existence of α , and that $(\mathcal{O}'_2, \iota_2)$ is a primitively \mathfrak{D} -oriented order. Thus, it follows from Lemma 9, that the connecting ideal between \mathcal{O}_1 and \mathcal{O}'_2 is generated by an \mathfrak{D} -ideal. Finally, the runtime is clear, as all operations done consists of simple linear algebra. \square

The Corollary from Proposition 7 is that \mathfrak{D} -Vectorization reduces to \mathfrak{D} -Endring in polynomial time, regardless of the size of $\text{Cl}(\mathfrak{D})[2]$, and without knowing the factorization of disc \mathfrak{D} , improving the results of Wesolowski [19].

Corollary 5. *Effective \mathfrak{D} -Vectorization reduces to \mathfrak{D} -Endring in polynomial time.*

Proof. In this proof, we reuse the notation from [19]. We are given two oriented curves $(E_1, \gamma_1), (E_2, \gamma_2) \in \mathcal{SS}_{\mathfrak{D}}(p)$, together with an ϵ -basis of $\text{End}(E_1)$ and $\text{End}(E_2)$, and our goal is to compute an \mathfrak{D} -ideal \mathfrak{a} such that $\mathfrak{a}\star(E_1, \gamma_1) = (E_2, \gamma_2)$, and an efficient representation of $\phi_{\mathfrak{a}} : (E_1, \gamma_1) \rightarrow \mathfrak{a}\star(E_1, \gamma_1)$.

First, compute optimal embeddings ι_1 and ι_2 such that $(\text{End}(E_i), \iota_i)$ are primitively \mathfrak{D} -oriented maximal orders using [19, Lemma 2]. Next, we run Algorithm 4 on $(\text{End}(E_1), \iota_1), (\text{End}(E_2), \iota_2)$, which outputs an \mathfrak{D} -ideal \mathfrak{a} solving the vectorization problem. Finally, an efficient representation of the isogeny $\phi_{\mathfrak{a}}$ can be computed unconditionally in polynomial time using [16, Theorem 2.8], or, for a more practical alternative, with [19, Proposition 9] assuming GRH. \square

4.2 Finding Fixed Norm Ideals Between Maximal Orders

When given two maximal orders $\mathcal{O}_1, \mathcal{O}_2$, we consider the problem of finding a left \mathcal{O}_1 -ideal I of norm d such that $\mathcal{O}_R(I) \cong \mathcal{O}_2$. This problem is of huge importance in isogeny-based cryptography, as it corresponds to computing isogenies of a given norm between supersingular curves, when they exist. One special case of this, is finding such an ideal of norm ℓ^k for some fixed, small prime ℓ , and the smallest $k \in \mathbb{Z}_{\geq 0}$, such that such an ideal exists. This corresponds to an optimal path between the curves in the ℓ -isogeny graph.

In this section, we give a new algorithm for solving this problem, based on our algorithms for the quaternion embedding problem. The algorithm consists of computing the ascending ideal to the correct level, and then bruteforcing the remaining horizontal part, for well chosen embeddings. To do this, we need the following Lemma.

Lemma 11. *Let (\mathcal{O}, ι) be a primitively $(\mathbb{Z} + d\mathfrak{D})$ -oriented maximal order, with $\omega = \iota(d\omega_0)$, where ω_0 is any generator of \mathfrak{D} , and let*

$$I := \mathcal{O}\langle\omega, d\rangle.$$

Then, $\mathfrak{n}(I) = d$, and $(\mathcal{O}_R(I), \iota)$ is a primitively \mathfrak{D} -oriented maximal order.

Proof. First, note that $\mathfrak{n}(I) = d$, because if not, this would contradict the primality of the embedding, by the same argument as the last part of Lemma 9. Next, we show that $(\mathcal{O}_R(I), \iota)$ is a \mathfrak{D} -oriented order, i.e. $\omega/d \in \mathcal{O}_R(I)$. To see this, note that for any element

$$\alpha\omega + \beta d \in I, \quad \alpha, \beta \in \mathcal{O}_L(I)$$

we have that

$$\begin{aligned} (\alpha\omega + \beta d)\omega/d &= \alpha\omega^2/d + \beta\omega \\ &= \alpha(\mathrm{tr}(\omega)\omega - \mathfrak{n}(\omega))/d + \beta\omega \\ &= \alpha(\mathrm{tr}(\iota(\omega_0))\omega + d\mathfrak{n}(\iota(\omega_0))) + \beta\omega \in I \end{aligned}$$

Finally, to see that the \mathfrak{D} -embedding on $\mathcal{O}_R(I)$ induced by ι is optimal, note that if it was not, this would again contradict the optimality of the $(\mathbb{Z} + d\mathfrak{D})$ -embedding on \mathcal{O} , since $\mathfrak{n}(I)$ induces the embedding $d\mathcal{O}_R(I) \subseteq \mathcal{O}_L(I)$. \square

For simplicity, we will assume factorization, and use a special purpose algorithm we call `GenericOrderEmbeddingFactorizationAll`, whose only difference with `GenericOrderEmbeddingFactorization`, is that it keeps searching and outputting solutions, until all are found. From the proof of Proposition 5, the expected runtime of this version is still $O\left(\frac{\sqrt{nm(\beta_1)}}{p}\right)$ under Heuristics 1 and 2.

Proposition 8. *Assume the existence of a factorization oracle, and that Heuristic 1 and Heuristic 2 holds. Let β_1 be the smallest non-integer in \mathcal{O}_1 , and let*

Algorithm 5 ConnectingIdealWithNorm_d($\mathcal{O}_1, \mathcal{O}_2$)

Input: Two maximal orders $\mathcal{O}_1, \mathcal{O}_2 \subset B_{p,\infty}$, and an integer d .

Output: \perp or an ideal I with $\mathcal{O}_L(I) = \mathcal{O}_1$, $\mathcal{O}_R(I) \cong \mathcal{O}_2$, and $n(I) = d$.

Let γ_1 be the element achieving the first successive minima of \mathcal{O}_2 .

for ω in GenericOrderEmbeddingFactorizationAll($\mathcal{O}_1, n(d\gamma_1), t(d\gamma_1)$) **do**

Set d' to be the biggest integer s.t. $\omega/d' \in \mathcal{O}_1$ for $i = 1, 2, 3$

Set $I_1 := \mathcal{O}_1 \langle \omega/d', d/d' \rangle$

Set $\mathcal{O}_{\text{crater}} := \mathcal{O}_R(I)$

Let $\iota : \mathbb{Z}[\gamma_1] \hookrightarrow \mathcal{O}_{\text{crater}}$ be defined by $\iota(\gamma_1) = \omega/d$.

for \mathfrak{I} in all $\mathbb{Z}[\gamma_1]$ -ideals of norm d' **do**

Set $I_2 := \mathcal{O}_{\text{crater}} \langle \iota(\mathfrak{I}) \rangle$

if $\mathcal{O}_R(I_2) \cong \mathcal{O}_2$ **then**

return $I_1 \cdot I_2$.

end if

end for

end for

return \perp

γ_1 be the smallest non-integer in \mathcal{O}_2 . ConnectingIdealWithNorm_d always returns a solution I if it exists, or \perp if a solution does not exist, and runs in expected time

$$O\left(\sqrt{n(\beta_1)n(\gamma_1)}\right).$$

Proof. First, we prove the correctness of the algorithm. Assume a solution I exists. We will prove that the solution I can be written as product $I = I_1 \cdot I_2$, where I_2 comes from a $\mathbb{Z}[\gamma_1]$ -ideal. This will also prove the correctness of the algorithm, as it runs through all embeddings of $\mathbb{Z}[d\gamma_1]$ into \mathcal{O}_1 , computes the unique corresponding ascending ideal I_1' , and then multiplies this with all the remaining ideals that comes from a $\mathbb{Z}[\gamma_1]$ -ideals.

Let us denote by ω_0 , the element in $\mathcal{O}_R(I)$ such that $\alpha\omega_0\alpha^{-1} = \gamma_1$ for some $\alpha \in B_{p,\infty}^\times$ (this exists since $\mathcal{O}_2 \cong \mathcal{O}_R(I)$). Let $\omega := d\omega_0$. The sequence of inclusions $d\mathcal{O}_R(I) \subset I \subset \mathcal{O}_L(I)$ coming from the fact that the norm of I is d implies that $\omega \in \mathcal{O}_L(I)$. Since $\mathcal{O}_L(I) \cap \mathcal{O}_R(I) = \mathbb{Z} + I$ it is easily verified that since $n(\omega) = 0 \pmod{d}$, we must have $\omega \in I$ or $\omega \in \bar{I}$. Without loss of generality we can assume that $\omega \in I$, and so we have $\mathcal{O}_L(I) \langle \omega, d \rangle \subset I$.

Let d' be the biggest integer such that $\omega/d' \in \mathcal{O}_L(I)$. It is clear that $d' \mid d$. We then set

$$I_1 := \mathcal{O}_L(I) \langle \omega/d', d/d' \rangle$$

where I_1 is a primitive ascending ideal of norm d/d' by Lemma 11, and we have $\mathcal{O}_L(I) \langle \omega, d \rangle = d'I_1 \subset I$.

The ideals I and I_1 are both primitive, contained inside $d'I_1$ and $n(I_1)$ divides $n(I)$ so it is easy to see that I must factor through I_1 and we must have $I \subset I_1$. Hence, we can define

$$I_2 := I_1^{-1} \cdot I$$

By Lemma 11, ω/d defines an optimal embedding of $\mathbb{Z}[\gamma_1]$ into $\mathcal{O}_R(I_1) = \mathcal{O}_L(I_2)$. Since we also had that $\omega/d \in \mathcal{O}_R(I) = \mathcal{O}_R(I_2)$, we conclude that by Lemma 9, I_2 comes from a $\mathbb{Z}[\gamma_1]$ -ideal.

Next, we analyse the runtime. Since we are assuming factorization, Heuristic 1 and Heuristic 2, `GenericOrderEmbeddingFactorizationAll` uses $O(\sqrt{n(\beta_1)})$ time before returning a potential solution. Enumerating the $\mathbb{Z}[\gamma_1]$ -ideals of norm d' can be done efficiently, again by factoring d' . Since there are at most $\mathcal{O}(h(\mathbb{Z}[\gamma_1]))$ isomorphism classes of maximal orders oriented by $\mathbb{Z}[\gamma_1]$, each candidate ideal we end up with has the correct right order with probability

$$O\left(\frac{1}{h(\mathbb{Z}[\gamma_1])}\right) = O\left(\frac{1}{\sqrt{n(\gamma_1)}}\right),$$

hence we get the expected runtime

$$O\left(\sqrt{n(\beta_1)n(\gamma_1)}\right)$$

to find a solution. □

The following corollary is immediate from Proposition 8, but we point it out here for convenience. The first part is generic, and gives the heuristic upper bounded runtime for finding equivalent ideals of given norm, independent of the degree. The second part says that when the orders are special, in the sense that they are both oriented by small quadratic orders, this problem can be solved efficiently, also independent of the degree.

Corollary 6. *Let $\mathcal{O}_1, \mathcal{O}_2 \subset B_{p,\infty}$ be two maximal orders. Then, assuming Heuristic 1, 2, and factorization, `ConnectingIdealWithNormd($\mathcal{O}_1, \mathcal{O}_2$)` runs in time*

$$O\left(p^{2/3}\right)$$

for any value of d . In the special case that there exists $\gamma_1 \in \mathcal{O}_1$ and $\beta_1 \in \mathcal{O}_2$ with $n(\beta_1), n(\gamma_1) \in O(1)$ `ConnectingIdealWithSmallNormd($\mathcal{O}_1, \mathcal{O}_2$)` runs in polynomial time for any d .

Proof. Immediate from combining Proposition 9, with using Lemma 1 to bound $n(\gamma_1)$ and $n(\beta_1)$ by $O(p^{2/3})$ in the generic case, or replacing them with $O(1)$ in the special case.

Appendix B illustrates why the second part of Corollary 6 is particularly interesting, namely because it allows us to compute optimal paths between such orders.

Finally, Algorithm 5 is expected to work in polynomial time for $d = O(p^{2/3})$, when only \mathcal{O}_1 contains an element of small norm. However, this expectation completely fails whenever the solution ideal comes from a $\mathbb{Z}[\gamma_1]$ -ideal, as the algorithm degenerates into bruteforcing $\mathbb{Z}[\gamma_1]$ -ideals. In Appendix C we give another algorithm, which always works in polynomial time in this case, assuming that the third successive minima of \mathcal{O}_2 is $O(p^{2/3})$, as one expects for “random” maximal orders.

References

1. Arpin, S., Clements, J., Dartois, P., Eriksen, J.K., Kutas, P., Wesolowski, B.: Finding orientations of supersingular elliptic curves and quaternion orders. arXiv preprint arXiv:2308.11539 (2023)
2. Bencina, B., Kutas, P., Merz, S., Petit, C., Stopar, M., Weitkämper, C.: Improved algorithms for finding fixed-degree isogenies between supersingular elliptic curves. IACR Cryptol. ePrint Arch. p. 1618 (2023), <https://eprint.iacr.org/2023/1618>
3. Castryck, W., Lange, T., Martindale, C., Panny, L., Renes, J.: CSIDH: an efficient post-quantum commutative group action. In: Advances in Cryptology–ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part III 24. pp. 395–427. Springer (2018)
4. Castryck, W., Panny, L., Vercauteren, F.: Rational isogenies from irrational endomorphisms. In: Canteaut, A., Ishai, Y. (eds.) Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12106, pp. 523–548. Springer (2020). https://doi.org/10.1007/978-3-030-45724-2_18, https://doi.org/10.1007/978-3-030-45724-2_18
5. Colò, L., Kohel, D.: Orienting supersingular isogeny graphs. Number-Theoretic Methods in Cryptology 2019 (2019)
6. De Feo, L., Kohel, D., Leroux, A., Petit, C., Wesolowski, B.: SQISign: Compact post-quantum signatures from quaternions and isogenies. In: Moriai, S., Wang, H. (eds.) Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7–11, 2020, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12491, pp. 64–93. Springer (2020). https://doi.org/10.1007/978-3-030-64837-4_3, https://doi.org/10.1007/978-3-030-64837-4_3
7. De Feo, L., Delpech de Saint Guilhem, C., Fouotsa, T.B., Kutas, P., Leroux, A., Petit, C., Silva, J., Wesolowski, B.: Séta: Supersingular encryption from torsion attacks. In: Advances in Cryptology–ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part IV 27. pp. 249–278. Springer (2021)
8. Eisenträger, K., Hallgren, S., Lauter, K.E., Morrison, T., Petit, C.: Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In: Nielsen, J.B., Rijmen, V. (eds.) Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III. Lecture Notes in Computer Science, vol. 10822, pp. 329–368. Springer (2018). https://doi.org/10.1007/978-3-319-78372-7_11, https://doi.org/10.1007/978-3-319-78372-7_11
9. Feo, L.D., Fouotsa, T.B., Kutas, P., Leroux, A., Merz, S.P., Panny, L., Wesolowski, B.: SCALLOP: scaling the CSI-FiSh. In: IACR International Conference on Public-Key Cryptography. pp. 345–375. Springer (2023)
10. Kohel, D.: Endomorphism rings of elliptic curves over finite fields. Ph.D. thesis, University of California at Berkeley (1996)
11. Kohel, D., Lauter, K., Petit, C., Tignol, J.P.: On the quaternion-isogeny path problem. LMS Journal of Computation and Mathematics **17**(A), 418–432 (2014)

12. Leroux, A.: An effective lower bound on the number of orientable supersingular elliptic curves. In: SAC 2022-Selected Areas in Cryptography (2022)
13. Leroux, A.: A new isogeny representation and applications to cryptography. In: Advances in Cryptology–ASIACRYPT 2022: 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5–9, 2022, Proceedings, Part II. pp. 3–35. Springer (2022)
14. Love, J., Boneh, D.: Supersingular curves with small noninteger endomorphisms. Open Book Series 4(1), 7–22 (2020)
15. Onuki, H.: On oriented supersingular elliptic curves. Finite Fields Their Appl. **69**, 101777 (2021). <https://doi.org/10.1016/J.FFA.2020.101777>, <https://doi.org/10.1016/j.ffa.2020.101777>
16. Page, A., Robert, D.: Introducing Clapoti(s): Evaluating the isogeny class group action in polynomial time. IACR Cryptol. ePrint Arch. p. 1766 (2023), <https://eprint.iacr.org/2023/1766>
17. The Sage Developers: SageMath, the Sage Mathematics Software System (version 9.7) (2022), <https://sagemath.org>
18. Voight, J.: Quaternion Algebras. Springer Graduate Texts in Mathematics series (2018)
19. Wesolowski, B.: Orientations and the supersingular endomorphism ring problem. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part III. Lecture Notes in Computer Science, vol. 13277, pp. 345–371. Springer (2022). https://doi.org/10.1007/978-3-031-07082-2_13, https://doi.org/10.1007/978-3-031-07082-2_13
20. Wigert, C.S.: Sur l'ordre de grandeur du nombre des diviseurs d'un entier. Almqvist & Wiksell (1907)

A Searching in the Unique Two-Sided Ideal of Norm p .

The main idea behind `GenericOrderEmbedding` is to exploit a system of equations on the trace pairings $\bmod \Delta$ provided by the formula from Proposition 2. In `GenericOrderEmbedding`, the final solution α is recovered by enumerating all suitable solutions of the trace pairings system until we find the solution.

One might wonder if we could try a different approach to recover the desired solution more efficiently. To simplify the reasoning we restrict hereafter to the special case where \mathcal{O} is a maximal order in $B_{p,\infty}$.

Let α be the solution we are looking for. Using the trace pairings $\bmod p$, we get values $t_1, t_2, t_3 \bmod p$. Let us take α_0 any solution to the trace pairing system, meaning that $\text{tr}(\alpha_0) = t$, and $\text{tr}(\alpha_0\beta_i) = t_i \bmod p$.

By linearity of the trace, we have that $\alpha_1 = \alpha - \alpha_0$ is an element that lies in the intersection $\bmod p$ of the trace pairing kernels.

Since

$$\text{tr}(\alpha_0\bar{\alpha}_1) \equiv \text{tr}(\alpha\bar{\alpha}_1) \pmod{p}$$

and

$$\text{tr}(\alpha\bar{\alpha}_1) = \text{tr}((\alpha_0 + \alpha_1)\bar{\alpha}_1) = \text{tr}(\alpha_0\bar{\alpha}_1) + \text{tr}(\alpha_1\bar{\alpha}_1),$$

we have that $\text{tr}(\alpha_1 \bar{\alpha}_1) = 2n(\alpha_1) \equiv 0 \pmod{p}$, hence α_1 is contained in the unique 2-sided ideal of norm p .

More precisely, it can be shown that the kernel of the trace pairing system mod p has always dimension 2. Writing the kernel as a lattice Λ , we get that α_1 must be contained in $\Lambda + p\mathcal{O}$. Thus, we end trying to find to solve the equation $\text{tr}(\alpha_1) = 0$ and $n(\alpha_1 + \alpha_0) = n$ in $\Lambda + p\mathcal{O}$. This yields a new ternary quadratic form, but it is unclear if it is any easier to solve.

However, we can use this idea to modify algorithm 1 in the case of maximal orders to work with any basis (not necessarily reduced), and achieve the same complexity under some heuristics. The idea is that once an element α_0 is found, which has the correct trace pairings modulo p , the element $-\alpha_1$ lying in the unique two-sided ideal of norm p , will heuristically be the vector in the lattice closest to p whenever $n(\alpha) = n(\alpha_0 + \alpha_1) < p^{4/3}$. We summarize this in Algorithm 6.

Algorithm 6 OrderEmbeddingCVP(\mathcal{O}, t, n)

Input: A maximal order $\mathcal{O} \subset B_{p,\infty}$, two integers $t, n \in \mathbb{Z}$ such that there exists an element of trace t and norm n in \mathcal{O} .

Output: \perp or $\alpha \in \mathcal{O}$ with $n(\alpha) = n$ and $\text{tr}(\alpha) = t$.

- 1: Compute any basis $1, \beta_1, \beta_2, \beta_3$ of \mathcal{O} .
 - 2: Compute $D_i = \text{tr}(\beta_i)^2 - 4n(\beta_i)$ for $1 \leq i \leq 3$, and $D = t^2 - 4n$.
 - 3: Compute s_i a square root of $DD_i \pmod{\Delta}$.
 - 4: Compute an element α_0 such that $\text{tr}(\alpha) = t$, and $\text{tr}(\alpha\beta_i) = t_i$ for $1 \leq i \leq 3$
 - 5: Compute the lattice Λ , the trace free part of the the unique two-sided \mathcal{O} -ideal of norm p .
 - 6: **for** Enumerate $-\alpha_1 \in \Lambda$, closest to α_0 **do**
 - 7: **if** $n(\alpha_0 + \alpha_1) = n$ **then**
 - 8: Return α .
 - 9: **end if**
 - 10: **end for**
-

Remark 2. We remark that this idea can also be used to get the same bound for the algorithm from [1]. Recall that this algorithm works by computing an HNF basis $\beta_1, \beta_2, \beta_3, \beta_4$ of the order, i.e.

$$\mathcal{O} = \langle e_{00} + e_{01}i + e_{02}j + e_{03}k, \\ e_{11}i + e_{12}j + e_{13}k, \\ e_{22}j + e_{23}k, \\ e_{33}i \rangle_{\mathbb{Z}}$$

Then one finds an element $\alpha_0 = t\beta_1 + x_0\beta_2$ by solving for the trace and norm modulo p . Then, for a solution α , one is looking for $\alpha_1 := \alpha + \alpha_0$ of the form $\alpha_1 = kp\beta_2 + y\beta_3 + z\beta_4$. It is clear that $p\beta_2, \beta_3, \beta_4$ again generates the trace free part of the unique two-sided ideal of norm p , hence, we can again expect α_1 to be the CVP solution to α_0 in this lattice whenever $n(\alpha) = n(\alpha_1 - \alpha_0) < p^{4/3}$.

B A Worked Example

We use Algorithm 7 to compute the shortest path in the 2-isogeny graph between E_0 and E_{1728} , where $j(E_i) = i$.

Let $p = 2^{55} \cdot 3 - 1 \equiv 11 \pmod{12}$. We work in the quaternion algebra

$$B_{p,\infty} = \mathbb{Q} + i\mathbb{Q} + j\mathbb{Q} + k\mathbb{Q}$$

where $i^2 = -1$ and $j^2 = -p$. Let $\mathcal{O}_{1728} \cong \text{End}(E_{1728})$ and $\mathcal{O}_0 \cong \text{End}(E_0)$. Explicitly, fix \mathcal{O}_{1728} to be

$$\mathcal{O}_{1728} = \mathbb{Z} + i\mathbb{Z} + \frac{i+j}{2}\mathbb{Z} + \frac{1+k}{2}\mathbb{Z}.$$

We also know that \mathcal{O}_0 contains the element $\omega = \frac{1+\sqrt{-3}}{2}$, where $\text{tr}(\omega) = 1$, and $\text{n}(\omega) = 1$. Hence, we look for the smallest $k \in \mathbb{N}$ such that $\mathbb{Z}[2^k\omega]$ embeds into \mathcal{O}_{1728} , by running `ConnectingIdealWithSmallNorm2k($\mathcal{O}_{1728}, \mathcal{O}_0$)` with for increasing $k \in [1, 2, \dots]$. This corresponds to running `GenericOrderEmbeddingFactorizationAll` with $t = \text{tr}(2^k\omega) = 2^k$ and $n = \text{n}(2^k\omega) = 2^{2k}$. We find that there exists an optimal embedding

$$\iota : \mathbb{Z}[2^k\omega] \hookrightarrow \mathcal{O}_{1728}$$

For $k = 54$ defined by

$$\iota(2^k\omega) = 9007199254740992 + \frac{19924704230006999}{2}i - \frac{23041705}{2}j - 34653096k,$$

and we use this element to find an ideal connecting \mathcal{O}_{1728} and \mathcal{O}_0 of norm 2^{54} . Translating this to an isogeny from

$$E_{1728} : y^2 = x^3 + x$$

We find that the point $K \in E_{1728}$ with

$$x(K) = 86739268981076750i + 69276702275648044, \quad i^2 = -1$$

generates an isogeny to E_0 of degree 2^{54} , corresponding to the shortest path between E_0 and E_{1728} in the 2-isogeny graph.

C Another Algorithm for Finding Equivalent Ideals

As mentioned, the problem with Algorithm 5 is that when (most of) the solution ideal comes from a $\mathbb{Z}[\gamma_1]$ -ideal, Algorithm 5 may end up brute-forcing through many horizontal ideals if the degree contains many distinct prime factors. In Algorithm 7, we fix this issue. The idea is to compute embeddings for all elements in a basis of \mathcal{O}_2 . Then we can recover the solution ideal using Lemma 12, which we state below.

Lemma 12. *Let $\mathcal{O} \subseteq B_{p,\infty}$ be a maximal order, and let I be a primitive right \mathcal{O} -ideal of norm d coprime to p . Given a basis $1, \gamma_1, \gamma_2, \gamma_3$ of \mathcal{O} . Let d_i be the smallest integer such that $d_i \gamma_i \in \mathcal{O}_L(I)$, and let*

$$I_i = \mathcal{O}_L(I) \langle d_1 \gamma_1, d_i \rangle.$$

Then

$$I = I_1 \cap I_2 \cap I_3$$

Proof. I is a primitive ideal connecting its left and right order of norm coprime to p , so all the ideals connecting $\mathcal{O}_L(I)$ and $\mathcal{O}_R(I)$ are contained in I . Indeed, let J be an ideal connecting $\mathcal{O}_L(I)$ and $\mathcal{O}_R(J)$, then $I \cdot \bar{J}$ is an $\mathcal{O}_L(I)$ two-sided ideal, and theory (see [18]) tells us that this ideal must be a scalar multiplied by the unique two-sided ideal of norm p in $\mathcal{O}_L(I)$. Since I has norm coprime to p , J can be factored as I times scalars times the unique two-sided ideal and so J is contained in I .

Let $J = I_1 \cap I_2 \cap I_3$.

In the proof of Proposition 8, we proved that we must have $I \subset I_i$ for $i = 1, 2, 3$ and so we have $I \subset J$.

We will now show that we must have $J \subset I$. For that, we are going to use that $\{x \in \mathcal{O}_L(I) \mid x\mathcal{O}_R(I) \subset \mathcal{O}_L(I)\} \subset I$. It is easy to see that this set $\{x \in \mathcal{O}_L(I) \mid x\mathcal{O}_R(I) \subset \mathcal{O}_L(I)\}$ is an ideal whose left order is $\mathcal{O}_L(I)$ and right order is $\mathcal{O}_R(I)$. Thus, it must be contained in I by what we proved earlier.

Let us now take $x \in J$. For each $i = 1, 2, 3$, there must be $\alpha_i, \beta_i \in \mathcal{O}_L(I)$ such that $x = \alpha_i d_i \gamma_i + \beta_i d_i$.

We are going to show that $x\mathcal{O}_R(I) \subset \mathcal{O}_L(I)$. Let us take $y \in \mathcal{O}_R(I)$. Since $1, \gamma_1, \gamma_2, \gamma_3$ is a basis of $\mathcal{O}_R(I)$, we have that $y = y_0 + \sum_{i=1}^3 y_i \gamma_i$. Thus, $xy = y_0 x + \sum_{i=1}^3 y_i (\alpha_i d_i \gamma_i + \beta_i d_i) \gamma_i$. With $\gamma_i^2 = \text{tr}(\gamma_i) \gamma_i - n(\gamma_i)$ we get

$$xy = y_0 x + \sum_{i=1}^3 y_i (\alpha_i n(\gamma_i)) d_i + (\beta_i + \alpha_i \text{tr}(\gamma_i)) d_i \gamma_i$$

and it is easy to verify that this belongs to $\mathcal{O}_L(I)$.

This proves that $J \subset I$ and this proves the result. □

We now give the algorithm.

Proposition 9. *Assume the existence of a factorization oracle, and that Heuristic 1 holds. Let β_1 be the smallest non-integer in \mathcal{O}_1 , and let $1, \gamma_1, \gamma_2, \gamma_3$ be a Minkowski-reduced basis of \mathcal{O}_2 . `ConnectingIdealWithSmallNormd` always returns a solution I if it exists, or \perp if a solution does not exist, and runs in time*

$$O \left(\max \left\{ \left\lceil \left\lfloor \frac{d\sqrt{n(\beta_1)n(\gamma_3)}}{p} \right\rfloor \right\rceil, \left\lceil \left\lfloor \frac{d\sqrt{n(\gamma_1)}}{p} \right\rfloor \right\rceil, \left\lceil \left\lfloor \frac{d\sqrt{n(\gamma_2)}}{p} \right\rfloor \right\rceil, \left\lceil \left\lfloor \frac{d\sqrt{n(\gamma_3)}}{p} \right\rfloor \right\rceil \right\} \right).$$

Algorithm 7 ConnectingIdealWithSmallNorm $_d(\mathcal{O}_1, \mathcal{O}_2)$

Input: Two maximal orders $\mathcal{O}_1, \mathcal{O}_2 \subset B_{p, \infty}$, and an integer d .

Output: \perp or an ideal I with $\mathcal{O}_L(I) = \mathcal{O}_1$, $\mathcal{O}_R(I) \cong \mathcal{O}_2$, and $n(I) = d$.

Let $1, \gamma_1, \gamma_2, \gamma_3 \in \mathcal{O}_2$ be a Minkowski-reduced basis of \mathcal{O}_2 .

Compute $\text{All}_{\omega_1} = \text{GenericOrderEmbeddingFactorizationAll}(\mathcal{O}_1, n(d\gamma_1), t(d\gamma_1))$

Compute $\text{All}_{\omega_2} = \text{GenericOrderEmbeddingFactorizationAll}(\mathcal{O}_1, n(d\gamma_2), t(d\gamma_2))$

Compute $\text{All}_{\omega_3} = \text{GenericOrderEmbeddingFactorizationAll}(\mathcal{O}_1, n(d\gamma_3), t(d\gamma_3))$

for $\omega_1, \omega_2, \omega_3$ in $\text{All}_{\omega_1} \times \text{All}_{\omega_2} \times \text{All}_{\omega_3}$. **do**

 Set d_i to be the biggest integer s.t. $\omega_i/d_i \in \mathcal{O}_1$ for $i = 1, 2, 3$

 Set $I_i := \mathcal{O}_1 \langle \omega_i/d_i, d/d_i \rangle$ for $i = 1, 2, 3$

 Set $I := I_1 \cap I_2 \cap I_3$

if $\mathcal{O}_R(I) \cong \mathcal{O}_2$ **then**

return I .

end if

end for

return \perp

Proof. First, we show the correctness of the algorithm. Assume that a solution I exists. Then I induces an embedding $d\mathcal{O}_2 \cong d\mathcal{O}_R(I) \subset \mathcal{O}_1 = \mathcal{O}_L(I)$. The isomorphism is given by an element α , i.e. $\alpha\mathcal{O}_2\alpha^{-1} = \mathcal{O}_R(I)$. Setting $\omega_i := \alpha^{-1}d\gamma_i\alpha$ for $i \in \{1, 2, 3\}$, it follows from Lemma 12 that

$$I = \bigcap_{i=1}^3 \mathcal{O}_1 \langle \omega_i/d_i, d/d_i \rangle$$

where d_i are the biggest integers such that $\omega_i/d_i \in \mathcal{O}_1$, thus showing the correctness of the algorithm.

Next, we analyse the runtime. The first potentially dominating term follows directly from running `GenericOrderEmbeddingFactorizationAll` on γ_i sequentially, and noting that $\gamma_1 < \gamma_2 < \gamma_3$. However, when β_1 is sufficiently small, the bottleneck of the algorithm becomes iterating over the cartesian product of the solutions. For each γ_i , we bound the number of solutions with Heuristic 2, giving the second dominating term. \square

Thus, from Proposition 9, we see that when $n(\beta_1) = O(1)$, and $n(\gamma_3) = O(p^{2/3})$ (as one expects for a random maximal order), Algorithm 7 runs in polynomial time for $d < p^{2/3}$.