

THE SUPERSINGULAR ISOGENY PATH AND ENDOMORPHISM RING PROBLEMS ARE EQUIVALENT

BENJAMIN WESOŁOWSKI

ABSTRACT. We prove that the path-finding problem in ℓ -isogeny graphs and the endomorphism ring problem for supersingular elliptic curves are equivalent under reductions of polynomial expected time, assuming the generalised Riemann hypothesis. The presumed hardness of these problems is foundational for isogeny-based cryptography. As an essential tool, we develop a rigorous algorithm for the quaternion analog of the path-finding problem, building upon the heuristic method of Kohel, Lauter, Petit and Tignol. This problem, and its (previously heuristic) resolution, are both a powerful cryptanalytic tool and a building-block for cryptosystems.

1. INTRODUCTION

We consider two problems of foundational importance to isogeny-based cryptography, a branch of post-quantum cryptography: the endomorphism ring problem and the path-finding problem in isogeny graphs, for supersingular elliptic curves. The hardness of the first is necessary for isogeny-based cryptography to be secure [GPST16, CPV20]. Reciprocally, some cryptosystems (the earliest of which being [CLG09]) are proven secure if the second is hard. Both problems are believed to be equivalent, thereby constituting the bedrock of isogeny-based cryptography. However, known reductions rely on a variety of heuristic assumptions [PL17, EHM17, EHL⁺18]. To arithmeticians, the endomorphism ring problem is simply the computational incarnation of the Deuring correspondence [Deu41]. This arithmetic theory met graph theory in the work of Mestre [Mes86] and Pizer [Piz90], and the related computational questions have been studied since [Koh96], yet the literature still heavily relies on heuristics.

This paper aims for a rigorous study of these problems from the generalised Riemann hypothesis (henceforth, GRH). As tools, we develop a rigorous algorithm to solve norm equations in quaternion algebras, and a rigorous variant of the heuristic algorithm from [KLPT14] for the quaternion analog of the path-finding problem, overcoming obstacles previously deemed “beyond the reach of existing analytic number theory techniques” [GPS20]. As an application we prove that the path-finding problem in ℓ -isogeny graphs and the endomorphism ring problem for supersingular elliptic curves are equivalent under reductions of polynomial expected time.

1.1. Hard problems for isogeny-based cryptography. The first isogeny-based cryptosystems were proposed by Couveignes in 1997 [Cou06]. This work was only made public in 2006, when the idea reemerged in [CLG09]. The latter introduced the path-finding problem in supersingular ℓ -isogeny graphs as a possible hard problem upon which cryptosystems can be constructed.

To any primes p and ℓ are associated a so-called *supersingular ℓ -isogeny graph*. It is a regular graph of degree $\ell + 1$ and counting approximately $p/12$ vertices. Each vertex of the graph is a *supersingular elliptic curve*, and edges correspond to ℓ -isogenies between them (a particular kind of morphisms between elliptic curves). Most importantly, these graphs are Ramanujan, i.e., optimal expander graphs. This implies that random walks quickly reach the uniform distribution. Starting from an elliptic curve E , one can compute a chain of random ℓ -isogenies until the endpoint E' is uniformly distributed. Then, given only E and E' , it seems hard to recover a path connecting them. This is the key of the preimage-resistant CGL hash function [CLG09], and the first of our problems of interest.

Problem 1.1 (ℓ -ISOGENYPATH). Given a prime p , and two supersingular elliptic curves E and E' over \mathbf{F}_{p^2} , find a path from E to E' in the ℓ -isogeny graph.

Isogeny-based cryptography has since grown considerably, when Jao and De Feo [JD11] noticed that it allows to build “post-quantum” cryptosystems, supposed to resist an adversary equipped with a quantum computer. There is today a wealth of other public-key protocols [CLM⁺18, DKPS19, Cos20] (including a Round 3 candidate [JAC⁺17] for NIST’s standardisation effort), signature schemes [BKV19, DG19, GPS20, DKL⁺20] or other cryptosystems [DMPS19, BKW20] built on the presumed hardness of finding isogenies connecting supersingular elliptic curves.

The precise relation between the security of these schemes and the supposedly hard problem ℓ -ISOGENYPATH is a critical question. Some of these schemes, like [CLG09] or [GPS20], are known to be secure if finding isogeny paths is hard. The reciprocal has been unclear: if one can solve ℓ -ISOGENYPATH efficiently, is all of isogeny-based cryptography broken? The first element of response was discovered in [GPST16] by taking a detour through another problem. They prove that an efficient algorithm to solve the closely related *endomorphism ring problem* allows to break the Jao–De Feo key exchange, and essentially all schemes of this type (see [FKM21]). Similarly, it was proven in [CPV20] that the security of CSIDH [CLM⁺18] and its variants (an *a priori* very different family of cryptosystems) also reduces to the endomorphism ring problem, via a sub-exponential reduction.

Given an elliptic curve E , an endomorphism is an isogeny $\varphi : E \rightarrow E$ from E to itself. The set of all endomorphisms of E , written $\text{End}(E)$, is a ring, where the addition is pointwise and multiplication is given by composition. Loops in ℓ -isogeny graphs provide endomorphisms, hence the connection between path-finding problems and computing endomorphism rings. Since the curves considered are supersingular, the endomorphism rings are always generated by four elements (as a lattice), and they are isomorphic to certain subrings of a quaternion algebra $B_{p,\infty}$, called *maximal orders*. The problem of computing the endomorphism ring comes in two flavours. The first actually looks for endomorphisms.

Problem 1.2 (ENDRING). Given a prime p , and a supersingular elliptic curves E over \mathbf{F}_{p^2} , find four endomorphisms of E (in an efficient representation) that generate $\text{End}(E)$ as a lattice.

By an efficient representation for endomorphisms α , we mean that there is an algorithm to evaluate $\alpha(P)$ for any $P \in E(\mathbf{F}_{p^k})$ in time polynomial in the length of the representation of α and in $k \log(p)$. We also assume that an efficient representation of α has length $\Omega(\log(\deg(\alpha)))$. The second version asks for an abstract description of $\text{End}(E)$.

Problem 1.3 (MAXORDER). Given a prime p , and a supersingular elliptic curves E over \mathbf{F}_{p^2} , find four quaternions in $B_{p,\infty}$ that generate a maximal order \mathcal{O} such that $\mathcal{O} \cong \text{End}(E)$.

Neither of them clearly reduces to the other, and in [GPST16], it is only proven that solving both simultaneously allows to break cryptosystems. Many works have been studying the three problems ℓ -ISOGENYPATH, ENDRING and MAXORDER, as early as [Koh96], originally motivated by the importance of these structures in arithmetic geometry. With the increasing practical impact of these problems, it has become critical to understand their relations. It was shown in [EHL⁺18] that, under several heuristic assumptions, all three appear to be equivalent.

1.2. Contributions. We prove that the problems ℓ -ISOGENYPATH, ENDRING and MAXORDER are equivalent under reductions of polynomial expected time, assuming the generalised Riemann hypothesis. In doing so, we develop new tools for a rigorous study of these problems.

Most importantly, we develop a new, rigorous variant of the heuristic algorithm of [KLPT14] for QUATERNIONPATH, a quaternion analog of ℓ -ISOGENYPATH. This algorithm (and its variants) is a crucial component of the reductions, but is also a powerful cryptanalytic tool [GPST16] and a building-block for cryptosystems [DKPS19, GPS20, DKL⁺20]. More precisely, we solve in polynomial time the following problem for very flexible choices of \mathcal{N} , including the most important variants ℓ -QUATERNIONPATH and B -PSQUATERNIONPATH.

Problem 1.4 (QUATERNIONPATH). Given two maximal orders \mathcal{O}_1 and \mathcal{O}_2 in $B_{p,\infty}$ and a set \mathcal{N} of positive integers, find a left \mathcal{O}_1 -ideal I such that $\text{Nrd}(I) \in \mathcal{N}$ and $\mathcal{O}_R(I) \cong \mathcal{O}_2$ (definitions provided in Section 2.2). If \mathcal{N} is the set of powers of a prime ℓ , we call the corresponding

problem ℓ -QUATERNIONPATH. If \mathcal{N} is the set of B -powersmooth integers for some $B > 0$, we call the corresponding problem B -PSQUATERNIONPATH.

The design and analysis of this new algorithm spans several sections of the present article.

- In Section 3, we combine some algorithmic considerations in euclidean lattices and the Chebotarev density theorem to prove that given an ideal in a maximal order, one can efficiently find an equivalent prime ideal (Theorem 3.7). This serves as a preconditioning step in our algorithm, and has a heuristic analog in [KLPT14].
- In Section 4, we prove bounds in the number of ways to represent an integer n as a linear combination of a prime and a quadratic form. This is a generalisation of a classic problem of Hardy and Littlewood [HL23] on representing integers as $p + x^2 + y^2$. The proof resorts to analytic number theory, and the result, Theorem 4.2, unlocks the analysis of algorithms to solve certain diophantine equations in the following section.
- In Section 5, we design and analyse an algorithm (Theorem 5.1) to find integral solutions (s, t, x, y) of equations of the form

$$\det(\gamma)^2 f(s, t) + b f^\gamma(x, y) = n,$$

where n and b are positive integers, f is a positive definite, integral, binary quadratic form, and γ is a 2×2 integral matrix. The key allowing a rigorous analysis is to randomise the class of f^γ within its genus using random walks, and apply the results of the previous section. As a first application, we use this algorithm to solve norm equations in special maximal orders in Corollary 5.8.

- Finally, we piece everything together in Section 6, solving QUATERNIONPATH in Theorem 6.3. The power-of- ℓ case is an immediate consequence, and we specialise to the powersmooth case in Theorem 6.4.

Note that our efforts are focused on obtaining rigorous, polynomial-time algorithms, with little consideration for practical efficiency, hence we spend little energy on calculating or optimising the hidden constants. A fast implementation should certainly follow the heuristic algorithm [KLPT14], only resorting to our rigorous variant when unexpected obstructions are encountered.

This new algorithm at hands, we then tackle the various reductions between ℓ -ISOGENYPATH, ENDRING and MAXORDER. They are similar to heuristic methods from the literature, and notably [EHL⁺18], with a number of substantial differences that allow a rigorous analysis. Note that our chain of reductions has a different structure from [EHL⁺18].

- We start in Section 7 by proving that ℓ -ISOGENYPATH and MAXORDER are equivalent. To do so, we adapt previous heuristic methods, essentially replacing their reliance on [KLPT14] with the new rigorous variants. In particular, we prove that there is a polynomial time algorithm to convert certain ideals of prime power norm into isogenies.
- Finally, we prove in Section 8 that MAXORDER and ENDRING are equivalent. The reduction from ENDRING to MAXORDER is essentially the same as the heuristic reduction from [EHL⁺18], adapted to our new rigorous tools. The converse requires more work: the reduction from MAXORDER to ENDRING in [EHL⁺18] encounters several large random numbers which are hoped to be easy to factor with good probability. We propose a strategy that provably avoids hard factorisations, exploiting the tools developed in Section 3.

Note that we do not *a priori* restrict the size of solutions to the three problems; however, our reductions polynomially preserve bounds on the output size. In particular, all reductions preserve the property of having a polynomially bounded output size, a requirement in [EHL⁺18]. This allows the reductions to be more versatile, and apply for instance if one discovers an algorithm that solves ℓ -ISOGENYPATH with paths of superpolynomial length.

1.3. Notation. The statements $f = O(g)$, $f \ll g$ and $g = \Omega(f)$ are synonymous, where O is the classic big O notation. We write O_ε to signify that the hidden constants depend on ε .

We denote by \mathbf{Z} , $\mathbf{Z}_{>0}$, and \mathbf{Q} the ring of integers, the set of positive integers, and the field of rational numbers. For any prime power q , we denote by \mathbf{F}_q the finite field with q elements. The function \log denotes the natural logarithm. The size of a set S is denoted by $\#S$. If a and b are two integers, the greatest common divisor of a and b is written $\gcd(a, b)$. We write $a \mid b$ if a divides b , or $a \mid b^\infty$ if all prime factors of a divide b , or $a \parallel b$ if $a \mid b$ and $\gcd(a, b/a) = 1$. The number of divisors of a is denoted by $\tau(a)$, and the number of prime divisors by $\omega(a)$, and Euler's totient is $\phi(a)$. If R is a ring and n a positive integer, $M_{n \times n}(R)$ is the ring of $n \times n$ matrices with coefficients in R . All statements containing the mention (GRH) assume the generalised Riemann hypothesis.

2. PRELIMINARIES

2.1. Quadratic forms. We will extensively use the theory of quadratic forms; the reader can find more details on the theory, with a computational perspective, in [Coh13]. A quadratic form of dimension r is a polynomial in r variables whose terms all have degree 2. A quadratic form $f(x)$ in the variable $x = (x_1, \dots, x_r)$ is determined by its *Gram matrix* $G = (g_{ij})$, a symmetric $r \times r$ matrix such that

$$f(x) = x^t G x = \sum_i g_{ii} x_i^2 + 2 \sum_i \sum_{j>i} g_{ij} x_i x_j.$$

For computational purposes, we assume that quadratic forms are represented as their Gram matrix, and we let $\text{length}(f)$ be the total binary length of its coefficients. The form is *integral* if $f(x) \in \mathbf{Z}$ for any $x \in \mathbf{Z}^r$, or equivalently, if $g_{ij} \in \frac{1}{2} \mathbf{Z}$ and $g_{ii} \in \mathbf{Z}$. If f is integral and $n \in \mathbf{Z}$, we say that f represents n if there exists $x \in \mathbf{Z}^r$ such that $f(x) = n$. The form f is *definite* if $f(x) = 0$ implies $x = 0$, and it is *positive* if $f(x) \geq 0$ for all x . It is *primitive* if the greatest common divisor of all integers represented by f is 1. It is *binary* if $r = 2$. The *discriminant* of f is

$$\text{disc}(f) = \begin{cases} (-1)^{\frac{r}{2}} \det(2G) & \text{if } r \text{ is even,} \\ \frac{1}{2} (-1)^{\frac{r+1}{2}} \det(2G) & \text{if } r \text{ is odd.} \end{cases}$$

To any quadratic form f is associated a symmetric bilinear form

$$\langle x, y \rangle_f = \frac{1}{2} (f(x+y) - f(x) - f(y)).$$

Given the bilinear form, one can recover the Gram matrix as $g_{ij} = \langle e_i, e_j \rangle_f$, where $(e_i)_{i=1}^r$ is the canonical basis. If $\gamma \in M_{r \times r}(\mathbf{Q})$, let f^γ be the quadratic form defined by $f^\gamma(x) = f(\gamma x)$, with Gram matrix $\gamma^t G \gamma$. A *quadratic space* V is a \mathbf{Q} -vector space of finite dimension together with a *quadratic map* $q : V \rightarrow \mathbf{Q}$ such that for any (hence all) basis $(b_i)_{i=1}^r$ of V , we have that $q(\sum_i x_i b_i)$ is a quadratic form in x . A *lattice* is a full-rank \mathbf{Z} -submodule in a positive definite quadratic space. The discriminant of a lattice is the discriminant of the quadratic form induced by any of its bases. Any positive definite f induces a lattice structure on \mathbf{Z}^r , via the canonical basis. The geometric invariants of this lattice induce invariants of f . The *volume* of f is $\text{Vol}(f) = |\det(G)|^{1/2}$. The *covering radius* $\mu(f)$ is the smallest μ such that for any $y \in \mathbf{R}^r$, we have $\min_{x \in \mathbf{Z}^r} f(x - y) \leq \mu^2$. We will use the following bound.

Lemma 2.1. *If f is integral, then $\mu(f) \leq \frac{1}{2} r^{1/2} \gamma_r^{r/2} \text{Vol}(f)$, where γ_r is Hermite's constant.*

Proof. Let λ_i be the successive minima of f . We have $\frac{1}{2} \lambda_r \leq \mu(f) \leq \frac{r^{1/2}}{2} \lambda_r$. By Minkowski's second theorem, $\prod_{i=1}^r \lambda_i \leq \gamma_r^{r/2} \text{Vol}(f)$, and since f is integral, $\lambda_i \geq 1$, hence $\lambda_r \leq \gamma_r^{r/2} \text{Vol}(f)$. \square

2.2. Quaternion algebras. An algebra B is a *quaternion algebra over \mathbf{Q}* if there exist $a, b \in \mathbf{Q}^\times$ and $i, j \in B$ such that $(1, i, j, ij)$ is a \mathbf{Q} -basis for B and

$$i^2 = a, \quad j^2 = b, \quad \text{and } ji = -ij.$$

Given a and b , the corresponding algebra is denoted by $\left(\frac{a,b}{\mathbf{Q}}\right)$. Write an arbitrary element of B as $\alpha = x_1 + x_2i + x_3j + x_4ij$ with $x_i \in \mathbf{Q}$. The quaternion algebra B has a canonical involution $\alpha \mapsto \bar{\alpha} = x_1 - x_2i - x_3j - x_4ij$. It induces the *reduced trace* and the *reduced norm*

$$\mathrm{Trd}(\alpha) = \alpha + \bar{\alpha} = 2x_1, \quad \mathrm{Nrd}(\alpha) = \alpha\bar{\alpha} = x_1^2 - ax_2^2 - bx_3^2 + abx_4^2.$$

The latter is a quadratic map, which makes B a quadratic space, and endows its \mathbf{Z} -submodules with a lattice structure. The corresponding bilinear form is

$$\langle \alpha, \beta \rangle = \frac{1}{2} (\alpha\bar{\beta} + \beta\bar{\alpha}).$$

If Λ is a full-rank lattice in B , the *reduced norm* of Λ is $\mathrm{Nrd}(\Lambda) = \gcd(\mathrm{Nrd}(\alpha) \mid \alpha \in \Lambda)$. We associate to Λ the normalised quadratic map

$$q_\Lambda : \Lambda \longrightarrow \mathbf{Z} : \lambda \longmapsto \frac{\mathrm{Nrd}(\lambda)}{\mathrm{Nrd}(\Lambda)}.$$

An *order* \mathcal{O} in B is a full-rank lattice that is also a subring. It is *maximal* if it is not contained in any other order. For any lattice $\Lambda \subset B$, we define the *left order* of Λ and the *right order* of Λ as

$$\mathcal{O}_L(\Lambda) = \{\alpha \in B \mid \alpha\Lambda \subseteq \Lambda\}, \quad \text{and} \quad \mathcal{O}_R(\Lambda) = \{\alpha \in B \mid \Lambda\alpha \subseteq \Lambda\}.$$

If \mathcal{O} is a maximal order, and I is a left ideal in \mathcal{O} , then $\mathcal{O}_L(I) = \mathcal{O}$ and $\mathcal{O}_R(I)$ is another maximal order. Given two maximal orders \mathcal{O}_1 and \mathcal{O}_2 , their *connecting ideal* is the ideal

$$I(\mathcal{O}_1, \mathcal{O}_2) = \{\alpha \in B \mid \alpha\mathcal{O}_2\bar{\alpha} \subseteq [\mathcal{O}_2 : \mathcal{O}_1 \cap \mathcal{O}_2]\mathcal{O}_1\},$$

which satisfies $\mathcal{O}_L(I) = \mathcal{O}_1$ and $\mathcal{O}_R(I) = \mathcal{O}_2$.

Let \mathcal{O} be a maximal order. Two left \mathcal{O} -ideals I and J are *equivalent* if there exists $\alpha \in B$ such that $I = \alpha J$. The set of classes for this equivalence relation is the *(left) ideal class set* of \mathcal{O} , written $\mathrm{Cls}(\mathcal{O})$. The class of I is written $[I]$.

To any prime number p , one associates a quaternion algebra $B_{p,\infty}$. In algebraic terms, $B_{p,\infty}$ is defined as the unique quaternion algebra over \mathbf{Q} ramified exactly at p and ∞ . Explicitly, it is given by the following lemma, from [Piz80].

Lemma 2.2. *Let $p > 2$ be a prime. Then, $B_{p,\infty} = \left(\frac{-q,-p}{\mathbf{Q}}\right)$, where*

$$q = \begin{cases} 1 & \text{if } p \equiv 3 \pmod{4}, \\ 2 & \text{if } p \equiv 5 \pmod{8}, \\ q_p & \text{if } p \equiv 1 \pmod{8}, \end{cases}$$

where q_p is the smallest prime such that $q_p \equiv 3 \pmod{4}$ and $\left(\frac{p}{q_p}\right) = -1$. Assuming GRH, we have $q_p = O((\log p)^2)$, which can thus be computed in polynomial time in $\log p$.

For a given quaternion algebra, the defining pair (a, b) is not unique. However, in the rest of this article, the algebra $B_{p,\infty}$ will always be associated to the pair $(-q, -p)$ given in Lemma 2.2, and the induced basis $(1, i, j, ij)$. For each p , we distinguish a maximal order \mathcal{O}_0 in $B_{p,\infty}$, and a useful suborder $R + Rj$ in the following lemma. This order \mathcal{O}_0 will be referred to as the *special maximal order* of $B_{p,\infty}$.

Lemma 2.3. *For any $p > 2$, the quaternion algebra $B_{p,\infty}$ contains the maximal order*

$$\mathcal{O}_0 = \begin{cases} \left\langle 1, i, \frac{i+ij}{2}, \frac{1+j}{2} \right\rangle & \text{if } p \equiv 3 \pmod{4}, \\ \left\langle 1, i, \frac{2-i+ij}{4}, \frac{-1+i+j}{2} \right\rangle & \text{if } p \equiv 5 \pmod{8}, \\ \left\langle \frac{1+i}{2}, \frac{j+ij}{2}, \frac{i+cij}{q}, ij \right\rangle & \text{if } p \equiv 1 \pmod{8}, \end{cases}$$

where in the last case c is an integer such that $q \mid c^2p + 1$. Assuming GRH, the maximal order \mathcal{O}_0 contains the suborder $R + Rj$ with index $O((\log p)^2)$, where R is the ring of integers of $\mathbf{Q}(i)$. If ω is a reduced generator of R , then

$$\mathrm{Nrd}(s + t\omega + xj + y\omega j) = f(s, t) + pf(x, y),$$

where f is a principal, primitive, positive definite, integral binary quadratic form of discriminant $\mathrm{disc}(\mathbf{Q}(i)) = O((\log p)^2)$.

Proof. This lemma summarises [KLPT14, Section 2.2], itself based on [Piz80] and [LO77]. \square

If \mathcal{O} is any maximal order in $B_{p,\infty}$, then $\mathrm{disc}(\mathcal{O}) = p^2$. In fact, for any left \mathcal{O} -ideal I , we have $\#(\mathcal{O}/I) = \mathrm{Nrd}(I)^2$ and the normalised quadratic map q_I has discriminant p^2 . The following lemma tells us that the integers represented by q_I are the norms of ideals equivalent to I .

Lemma 2.4 ([KLPT14, Lemma 5]). *Let I be a left \mathcal{O} -ideal, and $\alpha \in I$. Then, $I\bar{\alpha}/\mathrm{Nrd}(I)$ is an equivalent left \mathcal{O} -ideal of norm $q_I(\alpha)$.*

2.3. Supersingular elliptic curves. A detailed account of the theory of elliptic curves can be found in [Sil86]. An *elliptic curve* is an abelian variety of dimension 1. More explicitly, given a field k of characteristic $p > 3$, an elliptic curve E can be described as an equation $y^2 = x^3 + Ax + B$ for $A, B \in k$ with $4A^3 + 27B^2 \neq 0$. The k -rational points of E is the set $E(k)$ of pairs $(x, y) \in k^2$ satisfying the curve equation, together with a point ∞_E ‘at infinity’. They form an abelian group, written additively, where ∞_E is the neutral element. The *geometric points* of E are the \bar{k} -rational points, where \bar{k} is the algebraic closure of k .

Let E_1 and E_2 be two elliptic curves defined over k . An *isogeny* $\varphi : E_1 \rightarrow E_2$ is a non-constant rational map that sends ∞_{E_1} to ∞_{E_2} . It is then a group homomorphism from $E_1(\bar{k})$ to $E_2(\bar{k})$, and its kernel over the algebraic closure, written $\ker(\varphi)$, is finite. The *degree* $\deg(\varphi)$ is the degree of φ as a rational map. When $\deg(\varphi)$ is coprime to p , then $\deg(\varphi) = \#\ker(\varphi)$. The degree is multiplicative, in the sense that $\deg(\psi \circ \varphi) = \deg(\psi)\deg(\varphi)$. For any integer $n \neq 0$, the multiplication-by- m map $[m] : E \rightarrow E$ is an isogeny. For any isogeny $\varphi : E_1 \rightarrow E_2$, its *dual* is the unique isogeny $\hat{\varphi} : E_2 \rightarrow E_1$ such that $\hat{\varphi} \circ \varphi = [\deg(\varphi)]$. If $\deg(\varphi) = \ell$ is prime, we say that φ is an ℓ -isogeny. Any isogeny factors as a product of isogenies of prime degrees, hence ℓ -isogenies are basic building blocks. An isogeny of degree coprime to p is uniquely determined by its kernel. Given this kernel, one can compute equations for the isogeny in time polynomial in $\deg(\varphi)$ and $\log p$ via Vélú’s formula [Vél71]. An isogeny can be represented in size polynomial in $\log p$ and $\deg(\varphi)$, for instance as a rational map, or by a generator of its kernel. The output of ℓ -ISOGENYPATH is a chain of ℓ -isogenies of length k ; it corresponds to an isogeny of degree ℓ^k , but should be represented as a sequence of ℓ -isogenies (so that the length of the representation is polynomial in ℓ and k instead of ℓ^k).

An *isomorphism* is an isogeny $\varphi : E_1 \rightarrow E_2$ of degree 1. We say that E_1 and E_2 are isomorphic over K (an extension of k) if there is an isomorphism between them that is defined over K . The j -invariant of E is $j(E) = \frac{256 \cdot 27 \cdot A^3}{4A^3 + 27B^2}$. We have $j(E_1) = j(E_2)$ if and only if E_1 and E_2 are isomorphic over the algebraic closure of k . It is then simple to test \bar{k} -isomorphism. It is also simple to compute explicit isomorphisms.

An endomorphism of E is an isogeny $\varphi : E \rightarrow E$ from E to itself. The endomorphism ring $\mathrm{End}(E)$ is the collection of these endomorphisms, together with the trivial map $\varphi(x, y) = \infty_E$. It is a ring for pointwise addition, and for composition of maps. The map $\mathbf{Z} \rightarrow \mathrm{End}(E) : m \mapsto [m]$ is an embedding. In that sense, $\mathrm{End}(E)$ contains \mathbf{Z} as a subring, but it is always larger (in positive characteristic). The curve E is *supersingular* if $\mathrm{End}(E)$ has rank 4 as a \mathbf{Z} -module. Then, $\mathrm{End}(E)$ is isomorphic to a maximal order in the quaternion algebra $B_{p,\infty}$, defined in Section 2.2. Up to \bar{k} -isomorphism, all supersingular elliptic curves are defined over \mathbf{F}_{p^2} , and there are $\lfloor p/12 \rfloor + \varepsilon$ of them, with $\varepsilon \in \{0, 1, 2\}$. Fix a prime $\ell \neq p$. The supersingular ℓ -isogeny graph (for p) is the graph whose vertices are these supersingular elliptic curves (up to isomorphism), and there is an edge from E_1 to E_2 for each ℓ -isogeny from E_1 to E_2 . It is a regular graph of degree $\ell + 1$ (because any E has $\ell + 1$ subgroups H of order ℓ , each inducing an isogeny of kernel H). The ℓ -isogeny graph is Ramanujan. In particular, random walks rapidly

converge to the uniform distribution, and any two curves of the graph are connected by an isogeny of degree ℓ^m with $m = O(\log p)$.

2.4. The Deuring correspondence. As already mentioned, given a supersingular elliptic curve E over \mathbf{F}_{p^2} , its endomorphism ring $\text{End}(E)$ is isomorphic to a maximal order in $B_{p,\infty}$. This *Deuring correspondence* is in fact a bijection

$$\left\{ \begin{array}{c} \text{Isomorphism classes of} \\ \text{maximal orders } \mathcal{O} \text{ in } B_{p,\infty} \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \text{Isomorphism classes of} \\ \text{supersingular elliptic curves } E \text{ over } \mathbf{F}_{p^2} \end{array} \right\} / \text{Gal}(\mathbf{F}_{p^2} / \mathbf{F}_p).$$

A more detailed account of the theory can be found in [Voi21, Chapter 42].

We have identified a special order \mathcal{O}_0 in Lemma 2.3, and it is natural to wonder what the corresponding elliptic curve may be. If $p \equiv 3 \pmod{4}$, then the curve E_0 defined by $y^2 = x^3 - x$ is supersingular. It is defined over \mathbf{F}_p , so has the Frobenius endomorphism $\pi : (x, y) \mapsto (x^p, y^p)$. Furthermore, if $\alpha \in \mathbf{F}_{p^2}$ satisfies $\alpha^2 = -1$, it is easy to check that $\iota : (x, y) \mapsto (-x, \alpha y)$ is also an endomorphism. These endomorphisms generate almost all $\text{End}(E_0)$: we actually have

$$\text{End}(E_0) = \mathbf{Z} \oplus \mathbf{Z} \iota \oplus \mathbf{Z} \frac{\iota + \iota\pi}{2} \oplus \mathbf{Z} \frac{1 + \pi}{2}.$$

Since $\iota^2 = [-1]$ and $\pi^2 = [-p]$, we have $\text{End}(E_0) \cong \mathcal{O}_0$. More generally, we have the following result.

Lemma 2.5 ([EHL⁺18, Proposition 3]). *Let \mathcal{O}_0 as in Lemma 2.3. There is an algorithm that for any prime $p > 2$ computes an elliptic curve E_0 over \mathbf{F}_p and $\iota \in \text{End}(E_0)$ such that*

$$\mathcal{O}_0 \longrightarrow \text{End}(E_0) : 1, i, j, ij \longmapsto [1], \iota, \pi, \iota\pi$$

is an isomorphism, and runs in time polynomial in $\log p$ (if $p \equiv 1 \pmod{8}$, we assume GRH).

The Deuring correspondence runs deeper than a simple bijection: it also preserves morphisms between the two categories. Given any isogeny $\varphi : E_1 \rightarrow E_2$, let $I_\varphi = \text{Hom}(E_2, E_1)\varphi$, where $\text{Hom}(E_2, E_1)$ is the set of isogenies from E_2 to E_1 . This object I_φ is a left $\text{End}(E_1)$ -ideal, hence $\mathcal{O}_L(I_\varphi) \cong \text{End}(E_1)$. Furthermore, $\mathcal{O}_R(I_\varphi) \cong \text{End}(E_2)$. In other words, I_φ connects $\text{End}(E_1)$ to $\text{End}(E_2)$, just as φ connects E_1 to E_2 . This construction preserves the ‘quadratic structure’, in the sense that $\text{Nrd}(I_\varphi) = \deg(\varphi)$.

Conversely, suppose I is a left $\text{End}(E_1)$ -ideal. Then, we can construct an isogeny φ_I as the unique isogeny with kernel $\bigcap_{\alpha \in I} \ker(\alpha)$. These two constructions are mutual inverses, meaning that for any I and φ , we have $I_{\varphi_I} = I$ and $\varphi_{I_\varphi} = \varphi$. The translation from I to φ_I can be computed efficiently, provided that I is an ideal in the special order \mathcal{O}_0 from Lemma 2.3, and that $\text{Nrd}(I)$ is powersmooth (its prime-power factors are polynomially bounded). This is the following lemma. Only the case $p \equiv 3 \pmod{4}$ is considered in [GPS20], but as noted in [EHL⁺18], it easily extends to arbitrary p .

Lemma 2.6 ([GPS20, Lemma 5]). *Let \mathcal{O}_0 as in Lemma 2.3, and E_0 as in Lemma 2.5. There exists an algorithm which, given a left \mathcal{O}_0 -ideal I of norm $N = \prod_i \ell_i^{e_i}$, returns the corresponding isogeny $\varphi_I : E_0 \rightarrow E_1$. The complexity of this algorithm is polynomial in $\log p$ and $\max_i(\ell_i^{e_i})$ (if $p \equiv 1 \pmod{8}$, we assume GRH).*

3. QUADRATIC FORMS AND PRIME SAMPLING

In this section, we consider the following problem: given an integral, primitive, positive definite quadratic form f of rank r , find $x \in \mathbf{Z}^r$ such that $f(x)$ is prime. We then give a first application of this problem, for finding ideals of prime norm in a given ideal class of a maximal order of $B_{p,\infty}$.

3.1. Sampling primes. Let f be an integral, primitive, positive definite quadratic form. In this section, we discuss the problem of sampling vectors in $\{x \in \mathbf{Z}^r \mid f(x) \leq \rho\}$ so that $f(x)$ is prime. Let us first focus on the binary case, for which the following theorem tells that an important proportion of vectors represent primes. It is a classical consequence of the effective Chebotarev density theorem under GRH, due to Lagarias and Odlyzko [LO77].

Theorem 3.1 (GRH). *If f is an integral, primitive, positive definite, binary quadratic form of discriminant D , the number of primes at most ρ represented by f is*

$$\pi_f(\rho) = \frac{\delta \rho}{h(D) \log \rho} + O(\rho^{1/2} \log(|D|\rho)),$$

where δ is 1 if $f(x, y)$ is equivalent to $f(x, -y)$, and $1/2$ otherwise.

The quantity $\pi_f(\rho)$ should be compared to the cardinality of $\{(x, y) \in \mathbf{Z}^2 \mid f(x, y) \leq \rho\}$, which we estimate in the following lemma, in a slightly more general form for later purposes.

Lemma 3.2. *For any integral, positive definite, binary quadratic form f , any $x_0 \in \mathbf{R}^2$ and any $\rho \geq 0$, we have*

$$\left| \#\{x \in \mathbf{Z}^2 \mid f(x + x_0) \leq \rho\} - \frac{\pi}{\text{Vol}(f)} \rho \right| \leq 2\pi \sqrt{\frac{2}{3}} \rho^{1/2} + \frac{2\pi}{3} \text{Vol}(f).$$

Proof. For any $z \geq 0$, let $V_2(z) = \pi z^2$ be the volume of the standard 2-ball of radius z . It is a classical application of the covering radius $\mu(f)$ that

$$V_2(\rho^{1/2} - \mu(f)) \leq \text{Vol}(f) \cdot \#\{x \in \mathbf{Z}^2 \mid f(x + x_0) \leq \rho\} \leq V_2(\rho^{1/2} + \mu(f)).$$

This comes from the fact that Voronoi cells of f have volume $\text{Vol}(f)$ and diameter $2\mu(f)$. From Lemma 2.1 with Hermite's constant $\gamma_2 = 2/\sqrt{3}$, we have $\mu(f) \leq \sqrt{\frac{2}{3}} \text{Vol}(f)$. We obtain

$$\left| \left(\rho^{1/2} \pm \mu(f) \right)^2 - \rho \right| \leq 2\mu(f)\rho^{1/2} + \mu(f)^2 \leq 2\sqrt{\frac{2}{3}} \text{Vol}(f)\rho^{1/2} + \frac{2}{3} \text{Vol}(f)^2,$$

from which the result follows. \square

Lemma 3.3. *Let f be a primitive, positive definite, integral, binary quadratic form, and let $\rho > 0$. There is an algorithm that samples uniformly random elements from*

$$\{(x, y) \in \mathbf{Z}^2 \mid f(x, y) \leq \rho\}$$

in polynomial time in $\log \rho$ and in $\text{length}(f)$.

Proof. Let $B_f(r) = \{v \in \mathbf{R}^2 \mid f(v) \leq r^2\}$ be the ball of radius r around the origin. Let $r = \rho^{1/2}$, and we wish to sample uniformly in $B_f(r) \cap \mathbf{Z}^2$. First, compute a Minkowski-reduced basis (b_1, b_2) of f with $f(b_1) \leq f(b_2)$. If $\rho < f(b_2)$, then $B_f(r) \cap \mathbf{Z}^2 \subset \mathbf{Z} b_1$, and we can uniformly sample $k \in \mathbf{Z}$ such that $k^2 \leq \rho/f(b_1)$ and return kb_1 . We may now assume that $\rho \geq f(b_2)$, which implies $r \geq \sqrt{2}\mu$, with μ the covering radius of f . Let $\mathcal{V} = \{v \in \mathbf{R}^2 \mid f(v) = \min_{\lambda \in \mathbf{Z}^2} f(v + \lambda)\}$ be the Voronoi cell around the origin. Given any $v \in \mathbf{R}^2$, a closest lattice vector is an element $\lambda(v) \in \mathbf{Z}^2$ such that $v \in \mathcal{V} + \lambda(v)$. This closest vector can be computed efficiently in dimension 2, and is unique for almost all v : only the boundaries of Voronoi cells are ambiguous. We sample as follows:

- (1) Sample $v \in B_f(r + \mu)$ uniformly.
- (2) Solve the closest vector problem for v , resulting in $\lambda(v)$ (unique with probability 1).
- (3) If $\lambda(v) \in B_f(r)$, return it; otherwise restart.

Let us analyse the distribution of $\lambda(v)$ when $v \in B_f(r + \mu)$ is uniform. For any $u \in \mathbf{Z}^2 \cap B_f(r)$, we have $\mathcal{V} + u \subset B_f(r + \mu)$, hence

$$\Pr[\lambda(v) = u] = \frac{\text{Vol}((\mathcal{V} + u) \cap B_f(r + \mu))}{\text{Vol}(B_f(r + \mu))} = \frac{\text{Vol}(\mathcal{V})}{\text{Vol}(B_f(r + \mu))}.$$

In particular, $\Pr[\lambda(v) = u]$ for $u \in \mathbf{Z}^2 \cap B_f(r)$ does not depend on u , which proves that the output of the sampling procedure is uniform in $\mathbf{Z}^2 \cap B_f(r)$. Finally,

$$\Pr[\lambda(v) \in B_f(r)] \geq \Pr[v \in B_f(r - \mu)] = \frac{(r - \mu)^2}{(r + \mu)^2} \geq \left(\frac{1 - 2^{-1/2}}{1 + 2^{-1/2}} \right)^2 \geq 0.028,$$

which proves that the procedure succeeds after an expected constant number of trials. \square

Proposition 3.4 (GRH). *Let f be a primitive, positive definite, integral, binary quadratic form. For any $\varepsilon > 0$, there is an algorithm that finds integers $x, y \in \mathbf{Z}$ such that $f(x, y)$ is a prime number at most $O_\varepsilon(|\text{disc}(f)|^{1+\varepsilon})$, and runs in polynomial time in $\text{length}(f)$.*

Proof. From Lemma 3.3, one can sample uniformly random pairs of integers $(x, y) \in \mathbf{Z}^2$ such that $f(x, y) \leq \rho$. We conclude by combining Lemma 3.2 and Theorem 3.1 (with $h(D) = O(|D|^{1/2} \log |D|)$, see for instance [Coh08, p. 138]), which imply that a uniformly random vector represents a prime with good probability. The ε in the exponent comes from the crossover point between the main term and the error term in Theorem 3.1. \square

Proposition 3.5 (GRH). *Let f be a primitive, positive definite, integral quadratic form of dimension $r \geq 3$. For any $\varepsilon > 0$, there is an algorithm that finds a vector $x \in \mathbf{Z}^r$ such that $f(x)$ is a prime number at most $O_\varepsilon\left(\left(2^{r(r-1)}|\text{disc}(f)|\right)^{1+\varepsilon}\right)$ (or $O_\varepsilon(|\text{disc}(f)|^{4/3+\varepsilon})$ if $r = 3$), and runs in polynomial time in $\text{length}(f)$.*

Proof. We are looking for two integral vectors u and v that generate a primitive binary quadratic form $g_{uv}(x, y) = f(xu + yv)$. We then apply Proposition 3.4 to g_{uv} .

Compute an LLL-reduced [LLL82] basis (b_1, b_2, \dots, b_r) of f so that $f(b_1) \leq 2^{r-1} \text{Vol}(f)^{2/r}$ and $\prod_i f(b_i) \leq 2^{r(r-1)} \text{Vol}(f)^2$. Let $u = b_1$. Then, factor $f(u) = \prod_i a_i^{e_i}$ where each a_i seems hard to factor further, and they are pairwise coprime. For each a_i , we now describe a procedure that will either reveal new factors of a_i (in which case we can restart with this new piece of information), or find a vector v_i such that $f(v_i) = \langle v_i, v_i \rangle$ is coprime to a_i . We proceed as follows:

- (1) We compute the greatest common divisor of a_i with each of $\langle b_j, b_j \rangle$ and $2\langle b_j, b_k \rangle$ (i.e., the coefficients of f in the basis b_1, \dots, b_r).
- (2) These common divisors cannot all be equal to a_i since f is primitive. So either one of them is a non-trivial factor of a_i (and we restart), or one of them is 1.
- (3) If there is an index j such that $\gcd(\langle b_j, b_j \rangle, a_i) = 1$, we return $v_i = b_j$.
- (4) Otherwise, there are indices j and k with $\gcd(\langle b_j, b_j \rangle, a_i) = \gcd(\langle b_k, b_k \rangle, a_i) = a_i$, and $\gcd(2\langle b_j, b_k \rangle, a_i) = 1$. Then, we return $v_i = b_j + b_k$.

Now, let $v = \sum_i v_i \prod_{j \neq i} a_j$, and as desired, $\gcd(f(u), f(v)) = 1$. We have $f(v) \ll f(b_1)^2 f(b_r)$, and the form g_{uv} is primitive. It has volume at most $\sqrt{f(u)f(v)}$. If $r > 3$, we have $\sqrt{f(u)f(v)} \ll \sqrt{f(b_1)f(b_2)f(b_3)f(b_r)} \leq 2^{r(r-1)/2} \text{Vol}(f)$, and if $r = 3$, we have $\sqrt{f(u)f(v)} \ll \text{Vol}(f)^{4/3}$. The result then follows from Proposition 3.4. \square

In applications, we will often need to find vectors representing primes that are *large enough* (but not too large). This can be done in a straightforward adaptation of the above strategy.

Proposition 3.6 (GRH). *There exists a constant c and an algorithm \mathcal{A} such that the following holds. Let f be a primitive, positive definite, integral quadratic form of dimension r . For any $\rho > (2^{r^2}|\text{disc}(f)|)^c$, the algorithm $\mathcal{A}(f, \rho)$ outputs a vector $x \in \mathbf{Z}^r$ such that $f(x)$ is a prime number between ρ and ρ^2 , and runs in polynomial time in $\text{length}(f)$ and $\log \rho$.*

Proof. As in the proof of Proposition 3.5, \mathcal{A} can compute a sub-basis of f that induces a primitive binary quadratic form g of discriminant at most $O_\varepsilon\left(\left(2^{r(r-1)}|\text{disc}(f)|\right)^{1+\varepsilon}\right)$. Applying Lemma 3.3, one can sample uniformly random pairs (x, y) such that $g(x, y) \leq \rho^2$. From Theorem 3.1, $g(x, y)$ is prime and larger than ρ with good probability, provided that ρ is large enough. \square

3.2. Computing equivalent ideals of prime norm. The above results will be important in the rest of the article, and we can already prove them useful with a first important application. Consider a maximal order \mathcal{O} in $B_{p,\infty}$ and a left \mathcal{O} -ideal I . We can compute an equivalent ideal J of prime norm as an immediate consequence of Proposition 3.5.

Algorithm 1 EQUIVPRIMEIDEAL $_{\varepsilon}(I)$

Require: A left ideal I in a maximal order \mathcal{O} .

Ensure: An ideal J of prime norm, and an element $\alpha \in I$ such that $J = I\bar{\alpha}/\text{Nrd}(I)$.

1: $\alpha \leftarrow$ an element $\alpha \in I$ such that $q_I(\alpha)$ is prime; {Proposition 3.5}

2: **return** $J = I\bar{\alpha}/\text{Nrd}(I)$, and α .

Theorem 3.7 (GRH). *For any $\varepsilon > 0$, Algorithm 1 is correct and runs in expected polynomial time in $\log \text{Nrd}(I)$ and $\log p$, and the output J has reduced norm $\text{Nrd}(J) = O_{\varepsilon}(p^{2+\varepsilon})$.*

Proof. It follows from Proposition 3.5, and the fact that $q_I : I \rightarrow \mathbf{Z}$ is a primitive, positive-definite, integral quadratic map of discriminant p^2 . \square

Remark 1. Recall that our efforts are focused on provability, and the constants we obtain are certainly not tight. In [KLPT14, Section 3.1], the analogue heuristic algorithm is expected to return J of norm $\text{Nrd}(J) = \tilde{O}(p^{1/2})$ most of the time, and they argue that in the worst case, one could possibly obtain $\text{Nrd}(J) = \tilde{O}(p)$.

For our applications, we need a slightly more powerful version.

Proposition 3.8 (GRH). *There is a constant c and an algorithm which on input a left ideal I , a bound $\rho > p^c$, and a prime $\ell \neq p$, returns an ideal J equivalent to I such that $\text{Nrd}(J)$ is a prime between ρ and ρ^2 , and ℓ is a non-quadratic residue modulo $\text{Nrd}(J)$, and runs in polynomial time in $\log \text{Nrd}(I)$, $\log p$, and ℓ .*

Proof. Apply Algorithm 1 with two modifications. First, we use Proposition 3.6 instead of Proposition 3.5. Second, assuming $\ell \neq 4$ we consider a sublattice $4\ell I \subset \Lambda \subset I$ in place of I , where the quotient $\Lambda/4\ell I$ is generated by any element x such that $q_I(x) \equiv 1 \pmod{4}$ and $q_I(x)$ is a non-quadratic residue modulo ℓ . It follows from quadratic reciprocity that for any y in the lattice, when $q_I(y)$ is prime, then ℓ is a non-quadratic residue modulo $q_I(y)$. Similarly, if $\ell = 2$, we consider a sublattice $8I \subset \Lambda \subset I$ where the quotient $\Lambda/8I$ is generated by any element x such that $q_I(x) \equiv 3$ or $5 \pmod{8}$. \square

4. REPRESENTING INTEGERS WITH QUADRATIC FORMS AND PRIMES

In this section, we count the number of ways to represent an integer n in the form $al + bf(x, y)$, where the integers a and b and the quadratic form f are fixed, and ℓ is required to be prime. The bounds we obtain are key to the analysis of algorithms designed in the following sections. The proof resorts to analytic number theory. The reader only interested in computational applications can safely read up to Corollary 4.3 before skipping to the next section.

We fix the following notation for the rest of the section. Let f be a primitive, integral, positive definite, binary quadratic form of discriminant $f_{\chi}f_0^2$ where f_{χ} is fundamental. Let $v = f_{\chi}f_0$. Let a, b, u be positive integers with $\gcd(a, b) = \gcd(b, v) = \gcd(u, v) = \gcd(a, f_{\chi}) = 1$. Let χ be the Kronecker symbol $\chi(m) = \left(\frac{f_{\chi}}{m}\right)$, primitive of conductor f_{χ} . Let n be a positive integer such that $\gcd(a, n) = \gcd(b, n) = \gcd(n - au, f_0) = 1$. Finally, let

$$\mathcal{S}_n(f) = \{(\ell, x, y) \mid a\ell + bf(x, y) = n, \text{ where } x, y, \ell \in \mathbf{Z}, \ell \text{ is prime, and } \ell \equiv u \pmod{v}\}.$$

The goal of this section is to obtain lower bounds on the size of $\mathcal{S}_n(f)$.

The problem at hand is a generalisation of the classic problem of Hardy and Littlewood [HL23] of representing integers as $\ell + x^2 + y^2$, where ℓ is prime. The *number of representations* of an

integer N by f is $r(N; f) = \#\{(x, y) \in \mathbf{Z}^2 \mid f(x, y) = N\}$. Following a classical approach to the Hardy and Littlewood problem, we can write

$$\#\mathcal{S}_n(f) = \sum_{\substack{\ell \leq n/a \\ \ell \text{ prime} \\ \ell \equiv n/a \pmod{b} \\ \ell \equiv u \pmod{v}}} r\left(\frac{n - a\ell}{b}; f\right).$$

Unfortunately, controlling $r(N, f)$ is in general a difficult task. However, we know more about the number of representations of N in the genus of f . We indeed have the following classical theorem (see for instance [Pal33]).

Theorem 4.1. *Let f be a primitive, integral, binary quadratic form of discriminant $D = dm^2$, where d is fundamental. For any $N > 0$ such that $\gcd(N, m) = 1$, the number of representations of N by forms of discriminant D is*

$$w \sum_{\nu|N} \chi(\nu),$$

where $\chi(\nu) = \left(\frac{d}{\nu}\right)_K$ is the Kronecker symbol, and $w = 4$ when $d = 4$, $w = 6$ when $d = 3$, and $w = 2$ otherwise.

Let $(f_i)_{i=1}^t$ be a list of class representatives for each form of same discriminant as f . The main result of this section is the following theorem.

Theorem 4.2 (GRH). *There exists an absolute constant $\delta > 0$ such that for any integer $n \geq \max(a, b, v)^{1/\delta}$, we have*

$$\left| \sum_{i=1}^t \#\mathcal{S}_n(f_i) - \frac{n}{a \log(n/a)} \frac{w}{\phi(b)\phi(v)} \left(1 + \chi\left(\frac{n - ua}{b}\right)\right) L(1, \chi) C(\chi, an f_0, b) \right| = O(n^{1-\delta}),$$

where $L(s, \chi)$ is the Dirichlet L -function, the integer w is as in Theorem 4.1, and

$$C(\chi, m, s) = \prod_{\ell|ms} \left(1 + \frac{\chi(\ell)}{\ell(\ell-1)}\right) \prod_{\ell|m} \left(1 - \frac{\chi(\ell)}{\ell}\right).$$

The following (immediate) corollary is more convenient for the forthcoming applications.

Corollary 4.3 (GRH). *There exists a constant $c > 0$ such that for any positive integer n with $\log(n) \geq c \log \max(a, b, v)$, we have*

$$\sum_{i=1}^t \#\mathcal{S}_n(f_i) \geq \frac{n}{abv} \frac{1 + \chi\left(\frac{n-ua}{b}\right)}{(\log n)^c}.$$

Remark 2. If $\#\mathcal{S}_n(f_i) \neq 0$ and $\#\mathcal{S}_n(f_j) \neq 0$, then f_i and f_j must be in the same genus. Therefore, the sums in Theorem 4.2 and Corollary 4.3 are actually sums over class representatives of a single genus.

4.1. Preliminary results. We first present a theorem that will be a central tool in the proof of Theorem 4.2. It is essentially [ABL20, Theorem 2.1] with minor tweaks.

Theorem 4.4 (GRH). *There exists a positive constant δ with the following property. Let $x \geq 2$, $b, c, d \in \mathbf{Z}_{>0}$, $c_0, d_0 \in \mathbf{Z}$, $\gcd(d_0, d) = \gcd(c_0, c) = 1$, $a_1, a_2 \in \mathbf{Z} \setminus \{0\}$, $\gcd(b, da_1a_2) = 1$ such that*

$$Q \leq x^{1/2+\delta}, \quad a_1 \leq x^{1+\delta}, \quad a_2 \leq x^\delta, \quad b, c, d \leq x^\delta.$$

Then we have

$$\sum_{\substack{q \leq Q \\ \gcd(q, a_1 a_2 d) = 1 \\ q \equiv c_0 \pmod{c}}} \left(\sum_{\substack{n \leq x \\ n \equiv a_1/a_2 \pmod{bq} \\ n \equiv d_0 \pmod{d}}} \Lambda(n) - \frac{1}{\phi(qbd)} \sum_{\substack{n \leq x \\ \gcd(n, qbd) = 1}} \Lambda(n) \right) \ll x^{1-\delta},$$

where Λ is the von Mangoldt function.

Proof. Observe that if $d \mid c^\infty$, and $b = 1$, this is [ABL20, Theorem 2.1]. First, the condition $d \mid c^\infty$ is removed, and replaced with $\gcd(q, d) = 1$, with the following simple trick. Let δ_0 be as in [ABL20, Theorem 2.1], and $\delta = \delta_0/2$. Let D be the product of prime factors of d that do not divide c . Since $d \mid (Dc)^\infty$ and $Dc \leq x^{\delta_0}$, we can apply [ABL20, Theorem 2.1] and obtain that our sum (still with $b = 1$) is

$$\sum_{\substack{c'_0 \bmod Dc \\ c'_0 \equiv c_0 \bmod c \\ \gcd(c'_0, D)=1}} \sum_{\substack{q \leq Q \\ (q, a_1 a_2)=1 \\ q \equiv c'_0 \bmod Dc}} \left(\sum_{\substack{n \leq x \\ n \equiv a_1/a_2 \bmod q \\ n \equiv d_0 \bmod d}} \Lambda(n) - \frac{1}{\phi(qd)} \sum_{\substack{n \leq x \\ \gcd(n, qd)=1}} \Lambda(n) \right) \ll \varphi(D)x^{1-\delta_0} \leq x^{1-\delta}.$$

This proves the theorem in the case $b = 1$.

Second, let us deal with the case where $b \neq 1$. Let δ_1 be such that the theorem holds with $b = 1$. Let $\delta_2 = \delta_1/5$, and assume the conditions of the theorem are met for δ_2 . In particular, $bdx^{3\delta_2} \leq x^{\delta_1}$, and for any $\varepsilon > 0$ we have

$$\begin{aligned} & \sum_{\substack{q_b \leq Q \\ q_b \mid b^\infty \\ \gcd(q_b, c)=1}} \sum_{\substack{q \leq Q/q_b \\ \gcd(q, a_1 a_2 db)=1 \\ q \equiv c_0/q_b \bmod c}} \left(\sum_{\substack{n \leq x \\ n \equiv a_1/a_2 \bmod q \\ n \equiv a_1/a_2 \bmod bq_b \\ n \equiv d_0 \bmod d}} \Lambda(n) - \frac{1}{\phi(qq_bbd)} \sum_{\substack{n \leq x \\ \gcd(n, qbd)=1}} \Lambda(n) \right) \\ & \ll \sum_{\substack{q_b \leq x^{3\delta_2} \\ q_b \mid b^\infty \\ \gcd(q_b, c)=1}} x^{1-\delta_1} + \sum_{\substack{x^{3\delta_2} \leq q_b \leq Q \\ q_b \mid b^\infty \\ \gcd(q_b, c)=1}} \sum_{q \leq Q/q_b} \left(\sum_{\substack{n \leq x \\ n \equiv a_1/a_2 \bmod qq_b}} \Lambda(n) \right) \\ & \ll_\varepsilon x^{1-\delta_1+3\delta_2} + \sum_{\substack{x^{3\delta_2} \leq q_b \leq Q \\ q_b \mid b^\infty \\ \gcd(q_b, c)=1}} \sum_{q \leq Q/q_b} \left(\frac{x}{\phi(qbq_b)} + x^{1/2+\varepsilon} \right) \\ & \leq x^{1-2\delta_2} + x^{1-3\delta_2} \sum_{\substack{x^{3\delta_2} \leq q_b \leq Q \\ q_b \mid b^\infty \\ \gcd(q_b, c)=1}} \sum_{q \leq Q/q_b} \frac{1}{\phi(q)} + x^{1/2+\varepsilon} \sum_{\substack{x^{3\delta_2} \leq q_b \leq Q \\ q_b \mid b^\infty \\ \gcd(q_b, c)=1}} Q/q_b \\ & \leq x^{1-2\delta_2} + x^{1-3\delta_2} t(Q; b, c) \log(Q) + x^{1+\varepsilon-2\delta_2} t(Q; b, c), \end{aligned}$$

where $t(Q; b, c)$ is the number of positive integers at most Q whose prime factors divide b but not c . Using the estimates $t(Q; b, c) \leq (\log Q)^{\omega(b)}$ and $\omega(b) \ll \frac{\log(b)}{\log \log(b)}$, get that $t(Q; b, c) \ll_\varepsilon x^{\delta_2+\varepsilon}$, and we deduce that there is a constant $\delta_3 > 0$ such that the above is dominated by $x^{1-\delta_3}$. \square

We now prove an elementary lemma, which is a slight generalisation of [Hoo57, Lemma 3] or [ABL20, Lemma 5.2].

Definition 4.5. For any positive integers m and s , and any $x > 0$, let

$$\Phi_{m,s}(x) = \sum_{\substack{d \leq x \\ \gcd(d, m)=1}} \frac{\chi(d)}{\phi(sd)}.$$

Lemma 4.6. For any $\varepsilon > 0$,

$$\Phi_{m,s}(x) = \frac{1}{\phi(s)} L(1, \chi) C(\chi, m, s) + O_\varepsilon \left(\frac{f_\chi^{1/2+\varepsilon} (msx)^\varepsilon (\log x)^{\omega(s;m)}}{x} \right),$$

where $C(\chi, m, s)$ is as in Theorem 4.2, where $\omega(s; m)$ is the number of prime divisors of s that do not divide m .

Proof. Decomposing d as $d_s d$ where $(d, s) = 1$, then applying [ABL20, Lemma 5.2], we get

$$\begin{aligned} \Phi_{m,s}(x) &= \sum_{\substack{d_s \leq x \\ d_s | s^\infty, \gcd(d_s, m) = 1}} \frac{\chi(d_s)}{\phi(sd_s)} \sum_{\substack{d \leq x/d_s \\ \gcd(d, ms) = 1}} \frac{\chi(d)}{\phi(d)} \\ &= \sum_{\substack{d_s \leq x \\ d_s | s^\infty, \gcd(d_s, m) = 1}} \frac{\chi(d_s)}{\phi(sd_s)} \left(L(1, \chi) C(\chi, ms, 1) + O \left(\frac{\tau(ms) f_\chi^{1/2} \log(f_\chi) \log(x/d_s)}{x/d_s} \right) \right), \end{aligned}$$

where $\tau(n)$ is the number of divisors of n . The error term is dominated by

$$\frac{\tau(ms) f_\chi^{1/2} \log f_\chi \log(x)}{x} \sum_{\substack{d_s \leq x \\ d_s | s^\infty, \gcd(d_s, m) = 1}} \frac{d_s}{\phi(sd_s)} = \frac{\tau(ms) f_\chi^{1/2} \log(f_\chi) t(x; s, m) \log(x)}{\phi(s)x},$$

where $t(x; s, m)$ is the number of integers at most x whose prime factors divide s but not m . In particular, $t(x; s, m) \leq \log(x)^{\omega(s;m)}$. The main term in the lemma follows from the equality

$$\sum_{\substack{d_s | s^\infty \\ \gcd(d_s, m) = 1}} \frac{\chi(d_s)}{\phi(sd_s)} = \frac{1}{\phi(s)} \sum_{\substack{d_s | s^\infty \\ \gcd(d_s, m) = 1}} \frac{\chi(d_s)}{d_s} = \frac{1}{\phi(s)} \prod_{p|s, p \nmid m} \left(1 - \frac{\chi(p)}{p} \right)^{-1}.$$

The latter sum can be cut off to $d_s < x$ with an error dominated by

$$\frac{1}{\phi(s)} \sum_{\substack{d_s | s^\infty \\ \gcd(d_s, m) = 1 \\ d_s > x}} \frac{1}{d_s}.$$

An elementary recursion on the number of prime factors of s not dividing m yields

$$\sum_{\substack{d_s | s^\infty \\ \gcd(d_s, m) = 1 \\ d_s > x}} \frac{1}{d_s} \leq \frac{1}{x} (\log x)^{\omega(s)}.$$

Overall, the contributed error is dominated by

$$\frac{1}{\phi(s)} L(1, \chi) C(\chi, ms, 1) \sum_{\substack{d_s | s^\infty \\ \gcd(d_s, m) = 1 \\ d_s > x}} \frac{1}{d_s} \ll \frac{\log(f_\chi) \log \log(ms) (\log x)^{\omega(s)}}{\phi(s)x},$$

where we used the estimates $L(1, \chi) = O(\log f_\chi)$ and $C(\chi, m, s) = O(\log \log m)$, a consequence of the formula $\prod_{\ell|n} (1 - \frac{1}{\ell})^{-1} = O(\log \log n)$. \square

Corollary 4.7. For any $\varepsilon > 0$, there exists δ such that if $f_\chi, s \leq x^\delta$, then

$$\Phi_{m,s}(x) = \frac{1}{\phi(s)} L(1, \chi) C(\chi, m, s) + O_\varepsilon \left(\frac{(mx)^\varepsilon}{x} \right).$$

Proof. It follows from Lemma 4.6 and the estimate $\omega(s) = O\left(\frac{\log s}{\log \log s}\right)$. \square

4.2. Proof of Theorem 4.2. Following a classical approach to the Hardy and Littlewood problem, and using that $v = f_\chi f_0$ and $\gcd(n - au, f_0) = 1$, we have

$$\sum_{i=1}^t \#\mathcal{S}_n(f_i) = \sum_{\substack{\ell \leq n/a \\ \ell \text{ prime} \\ \ell \equiv n/a \pmod{b} \\ \ell \equiv u \pmod{v}}} \sum_{i=1}^t r\left(\frac{n - a\ell}{b}; f_i\right) = w \sum_{\substack{\ell \leq n/a \\ \ell \text{ prime} \\ \ell \equiv n/a \pmod{b} \\ \ell \equiv u \pmod{v}}} \sum_{d|\frac{n-a\ell}{b}} \chi(d).$$

The rest of the proof is dedicated to estimating the related sum

$$\sum_{\substack{\ell \leq n/a \\ \ell \equiv n/a \pmod{b} \\ \ell \equiv u \pmod{v}}} \Lambda(\ell) \sum_{d|\frac{n-a\ell}{b}} \chi(d),$$

where again, Λ is the von Mangoldt function. The theorem then follows by partial summation.

Let $\delta > 0$ be a parameter to be adjusted. For any $\varepsilon > 0$, The terms where $\ell < n^{1-\delta}$ can trivially be bounded by $O_\varepsilon(x^{1-\delta+\varepsilon})$. We can deal similarly with the terms where $\ell > n/a - n^{1-\delta}$. Therefore, it is sufficient to consider sums of the form

$$(1) \quad \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod{v} \\ \ell \equiv n/a \pmod{b}}} \Lambda(\ell) \sum_{d|\frac{n-a\ell}{b}} \chi(d),$$

where $n^{1-\delta} \leq Y < X \leq n/a - n^{1-\delta}$, and $n/(2a) < X$. Since $(a, n) = 1$, the terms where $\gcd(d, an) \neq 1$ can be discarded, and the remaining terms are distributed into three parts as

$$\begin{aligned} \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod{v} \\ \ell \equiv n/a \pmod{b}}} \Lambda(\ell) \sum_{\substack{d|\frac{n-a\ell}{b} \\ \gcd(d, an)=1}} \chi(d) &= \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod{v} \\ \ell \equiv n/a \pmod{b}}} \Lambda(\ell) \left(\sum_{\substack{d \leq D \\ db|n-a\ell \\ \gcd(d, an)=1}} \chi\left(\frac{n-a\ell}{db}\right) + \sum_{\substack{d < \frac{n-a\ell}{Db} \\ db|n-a\ell \\ \gcd(d, an)=1}} \chi(d) \right) \\ &= S_1 + S_2 + S_3, \end{aligned}$$

where $D = X^{1/2}$, and

$$\begin{aligned} S_1 &= \sum_{\substack{d \leq D \\ \gcd(d, an)=1}} \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod{v} \\ \ell \equiv n/a \pmod{db}}} \Lambda(\ell) \chi\left(\frac{n-a\ell}{db}\right), \\ S_2 &= \sum_{\substack{d \leq \frac{n-Xa}{Db} \\ \gcd(d, an)=1}} \chi(d) \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod{v} \\ \ell \equiv n/a \pmod{db}}} \Lambda(\ell), \\ S_3 &= \sum_{\substack{\frac{n-Xa}{Db} < d \leq \frac{n-Ya}{Db} \\ \gcd(d, an)=1}} \chi(d) \sum_{\substack{Y < \ell \leq \frac{n-bdD}{a} \\ \ell \equiv u \pmod{v} \\ \ell \equiv n/a \pmod{db}}} \Lambda(\ell). \end{aligned}$$

Estimation of S_1 . Decomposing d as dd_χ where $(d, f_\chi) = 1$ and $d_\chi \mid f_\chi^\infty$, we have

$$\begin{aligned} S_1 &= \sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \sum_{\substack{d \leq D/d_\chi \\ \gcd(d, f_\chi an)=1}} \bar{\chi}(d) \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod v \\ \ell \equiv n/a \pmod{dd_\chi b}}} \Lambda(\ell) \chi\left(\frac{n - a\ell}{d_\chi b}\right) \\ &= \sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \sum_{y \pmod{f_\chi}} \bar{\chi}(y) \sum_{x \pmod{f_\chi}} \chi(x) \sum_{\substack{d \leq D/d_\chi \\ \gcd(d, an)=1 \\ d \equiv y \pmod{f_\chi}}} \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod v \\ \ell \equiv n/a \pmod{db} \\ \ell \equiv (n - xbd_\chi)/a \pmod{f_\chi d_\chi}}} \Lambda(\ell). \end{aligned}$$

Since $\gcd(n - au, f_0) = 1$, the terms where $\gcd(d, f_0) \neq 1$ are zero. Applying Theorem 4.4 and the prime number theorem, we have

$$\begin{aligned} S_1 &= \sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \sum_{y \pmod{f_\chi}} \bar{\chi}(y) \sum_{\substack{x \pmod{f_\chi} \\ ua \equiv n - xbd_\chi \pmod{\gcd(f_\chi d_\chi, v)}}} \chi(x) \sum_{\substack{d \leq D/d_\chi \\ \gcd(d, an f_0)=1 \\ d \equiv y \pmod{f_\chi}}} \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \pmod v \\ \ell \equiv n/a \pmod{db} \\ \ell \equiv (n - xbd_\chi)/a \pmod{f_\chi d_\chi}}} \Lambda(\ell) \\ &\approx \sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \sum_{y \pmod{f_\chi}} \bar{\chi}(y) \sum_{\substack{x \pmod{f_\chi} \\ ua \equiv n - xbd_\chi \pmod{f_\chi \gcd(d_\chi, f_0)}}} \chi(x) \sum_{\substack{d \leq D/d_\chi \\ \gcd(d, an f_0)=1 \\ d \equiv y \pmod{f_\chi}}} \frac{\sum_{\substack{Y < \ell \leq X \\ (\ell, dbvf_\chi)=1}} \Lambda(\ell)}{\phi(\text{dblcm}(v, f_\chi d_\chi))} \\ &\approx \sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \sum_{y \pmod{f_\chi}} \bar{\chi}(y) \sum_{\substack{x \pmod{f_\chi} \\ ua \equiv n - xbd_\chi \pmod{f_\chi \gcd(d_\chi, f_0)}}} \chi(x) \sum_{\substack{d \leq D/d_\chi \\ \gcd(d, an f_0)=1 \\ d \equiv y \pmod{f_\chi}}} \frac{X - Y}{\phi(\text{db}f_\chi \text{lcm}(f_0, d_\chi))}, \end{aligned}$$

where the error introduced by the approximations is dominated by

$$\sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \sum_{\substack{y \pmod{f_\chi} \\ \gcd(y, f_\chi)=1}} \sum_{\substack{x \pmod{f_\chi} \\ \gcd(x, f_\chi)=1 \\ ua \equiv n - xbd_\chi \pmod{f_\chi \gcd(d_\chi, f_0)}}} X^{1-\delta_0} \leq X^{1-\delta_0} \phi(f_\chi)^2 t(D; f_\chi, an),$$

where δ_0 is the constant of Theorem 4.4, and $t(D; f_\chi, an)$ is the number of integers at most D whose prime factors divide f_χ but not an . Reorganising the terms, our estimation of S_1 becomes

$$S_1 \approx (X - Y) \sum_{\substack{d_\chi \leq D \\ d_\chi \mid f_\chi^\infty \\ \gcd(d_\chi, an)=1}} \left(\sum_{\substack{x \pmod{f_\chi} \\ ua \equiv n - xbd_\chi \pmod{f_\chi \gcd(d_\chi, f_0)}}} \chi(x) \right) \left(\sum_{\substack{d \leq D/d_\chi \\ \gcd(d, an f_0)=1}} \frac{\bar{\chi}(d)}{\phi(\text{db}f_\chi \text{lcm}(f_0, d_\chi))} \right).$$

Let us focus on the first inner sum. Let $m = \gcd(n - ua, f_\chi \gcd(d_\chi, f_0), d_\chi)$. We have

$$\begin{aligned} &\sum_{\substack{x \pmod{f_\chi} \\ xbd_\chi \equiv n - ua \pmod{f_\chi \gcd(d_\chi, f_0)}}} \chi(x) \\ &= \begin{cases} \sum_{\substack{x \pmod{f_\chi} \\ x \equiv (n - ua)/bd_\chi \pmod{f_\chi \gcd(d_\chi, f_0)/m}} \chi(x) & \text{if } \begin{cases} \gcd(d_\chi, f_\chi \gcd(d_\chi, f_0)) = m \\ \text{and } \gcd(n - ua, f_\chi \gcd(d_\chi, f_0)) = m, \end{cases} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

In the situation where $\gcd(d_\chi, f_\chi \gcd(d_\chi, f_0)) = \gcd(n - ua, f_\chi \gcd(d_\chi, f_0))$, let $\alpha = (n - ua)/bd_\chi \bmod f_\chi \gcd(d_\chi, f_0)/m$. From [IK04, (3.9)], we have

$$\begin{aligned} \sum_{\substack{x \bmod f_\chi \\ x \equiv \alpha \bmod f_\chi \gcd(d_\chi, f_0)/m}} \chi(x) &= \sum_{y \bmod m/\gcd(d_\chi, f_0)} \chi(\alpha + y(f_\chi \gcd(d_\chi, f_0)/m)) \\ &= \frac{m}{f_\chi \gcd(d_\chi, f_0)} \sum_{y \bmod f_\chi} \chi(\alpha + y(f_\chi \gcd(d_\chi, f_0)/m)) \\ &= \begin{cases} \chi(\alpha) & \text{if } m = \gcd(d_\chi, f_0), \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

In summary, and using that $\gcd(n - ua, f_0) = 1$, we get

$$\sum_{\substack{x \bmod f_\chi \\ xbd_\chi \equiv n-ua \bmod v}} \chi(x) = \begin{cases} \chi\left(\frac{n-ua}{b}\right) & \text{if } \gcd(n-ua, f_\chi) = 1 \text{ and } d_\chi = 1, \\ 0 & \text{otherwise.} \end{cases}$$

We deduce that our estimation of S_1 is

$$S_1 \approx (X - Y) \frac{1}{\phi(f_\chi f_0)} \chi\left(\frac{n-ua}{b}\right) \left(\sum_{\substack{d \leq D \\ \gcd(d, an f_0) = 1}} \frac{\chi(d)}{\phi(db)} \right).$$

Estimation of S_2 . Similarly, using $\frac{n-Xa}{Db} \leq \frac{n}{Db} \leq \frac{2aX^{1/2}}{b}$, we can apply Theorem 4.4 and get up to an admissible error

$$\begin{aligned} S_2 &= \sum_{y \bmod f_\chi} \chi(y) \sum_{\substack{d \leq \frac{n-Xa}{Db} \\ \gcd(d, an f_0) = 1 \\ d \equiv y \bmod f_\chi}} \sum_{\substack{Y < \ell \leq X \\ \ell \equiv u \bmod f_\chi f_0 \\ \ell \equiv n/a \bmod db}} \Lambda(\ell) \\ &\approx \sum_{y \bmod f_\chi} \chi(y) \sum_{\substack{d \leq \frac{n-Xa}{Db} \\ \gcd(d, an f_0) = 1 \\ d \equiv y \bmod f_\chi}} \frac{X - Y}{\phi(db f_\chi f_0)} \\ &\approx (X - Y) \frac{1}{\phi(f_\chi f_0)} \sum_{\substack{d \leq \frac{n-Xa}{Db} \\ \gcd(d, an f_0) = 1}} \frac{\chi(d)}{\phi(db)}. \end{aligned}$$

Main term of the estimation. Anticipating that S_3 will disappear in the error term, we get that the main term of our estimation of the sum (1) is

$$S_1 + S_2 \approx (X - Y) \frac{1}{\phi(b)} \frac{1}{\phi(f_\chi f_0)} \left(1 + \chi\left(\frac{n-ua}{b}\right) \right) L(1, \chi) C(\chi, an f_0, b),$$

From Corollary 4.7, for any $\varepsilon > 0$, the error introduced in the above estimation is dominated by

$$\frac{X - Y}{\phi(f_\chi f_0)} \frac{(an f_0 D')^\varepsilon}{D'},$$

for $D' = \min(D, \frac{n-Xa}{Db}) \geq n^{\frac{1}{2} - \frac{3\delta}{2}}$.

Estimation of S_3 . We now prove that the third and last term is absorbed in the error. We have

$$S_3 = \sum_{\substack{\frac{n-Xa}{Db} < d \leq \frac{n-Ya}{Db} \\ \gcd(d, an) = 1}} \chi(d) \sum_{\substack{Y < \ell \leq \frac{n-bdD}{a} \\ \ell \equiv u \bmod v \\ \ell \equiv n/a \bmod db}} \Lambda(\ell) = \sum_{y \bmod f_\chi} \chi(y) \sum_{\substack{\frac{n-Xa}{Db} < d \leq \frac{n-Ya}{Db} \\ \gcd(d, an f_0) = 1 \\ d \equiv y \bmod f_\chi}} \sum_{\substack{Y < \ell \leq \frac{n-bdD}{a} \\ \ell \equiv u \bmod v \\ \ell \equiv n/a \bmod db}} \Lambda(\ell).$$

To split this sum into manageable components, let

$$T_\chi(s, t, y, x) = \sum_{y \bmod f_\chi} \chi(y) \sum_{\substack{s < d \leq t \\ \gcd(d, anf_0)=1 \\ d \equiv y \bmod f_\chi}} \sum_{\substack{y < \ell \leq x \\ \ell \equiv u \bmod v \\ \ell \equiv n/a \bmod db}} \Lambda(\ell).$$

Assuming that $\frac{n-Xa}{Db} \leq s \leq t \leq n/Db$ and $x \geq Y$, and recalling that $Y \geq n^{1-\delta}$ where δ is adjustable, we can ensure that Theorem 4.4 and Corollary 4.7 apply in the following estimation:

$$\begin{aligned} T_\chi(s, t, y, x) &= \sum_{y \bmod f_\chi} \chi(y) \sum_{\substack{s < d \leq t \\ \gcd(d, anf_0)=1 \\ d \equiv y \bmod f_\chi}} \frac{x-y}{\phi(vdb)} + O(\phi(f_\chi)x^{1-\delta_0}) \\ &= (x-y) \sum_{\substack{s < d \leq t \\ \gcd(d, anf_0)=1}} \frac{\chi(d)}{\phi(vdb)} + O(\phi(f_\chi)x^{1-\delta_0}) \\ &= (x-y)(\Phi_{anf_0, vb}(t) - \Phi_{anf_0, vb}(s)) + O(\phi(f_\chi)x^{1-\delta_0}) \\ &\ll (x-y) \frac{(anf_0s)^\varepsilon}{s} + \phi(f_\chi)x^{1-\delta_0}. \end{aligned}$$

Also, if $\mathbf{1}_\chi$ is the principal character of conductor f_χ ,

$$\begin{aligned} T_{\mathbf{1}_\chi}(s, t, y, x) &= \sum_{\substack{s < d \leq t \\ (d, anv_b)=1}} \sum_{\substack{y < \ell \leq x \\ \ell \equiv u \bmod v \\ \ell \equiv n/a \bmod db}} \Lambda(\ell) \\ &= \sum_{\substack{s < d \leq t \\ \gcd(d, anv_b)=1}} \frac{x-y}{\phi(vdb)} + O(x^{1-\delta_0}) \\ &\ll (x-y) \log(t/s) + x^{1-\delta_0}. \end{aligned}$$

Let $\Delta = n^{-\delta_0/2}$. We now split the ℓ -sum in S_3 into k intervals of the form $(L_i, L_{i+1}]$ with $L_1 = Y$, $L_{i+1} = (1-\Delta)L_i + \Delta n/a$, and cropping the last interval so that $(Y, \frac{n-bdD}{a}] = \bigcup_{i=1}^k (L_i, L_{i+1}]$. Let $F(\ell) = \frac{n}{bD} - \ell \frac{a}{bD}$. We have $F(L_{i+1})/F(L_i) = 1-\Delta$. In particular, k is the smallest integer such that $(1-\Delta)^{k-1} \leq F(X)/F(Y)$. We can assume that $X \leq n/a - 1$, and deduce that $k = O(\Delta^{-1} \log(n))$. Now, we have

$$S_3 \ll \sum_i \left(T_\chi \left(\frac{n-Xa}{bD}, F(L_{i+1}), L_i, L_{i+1} \right) + T_{\mathbf{1}_\chi}(F(L_{i+1}), F(L_i), L_i, L_{i+1}) \right).$$

On one hand,

$$\begin{aligned} \sum_{i=1}^k T_{\mathbf{1}_\chi}(F(L_i), F(L_{i-1}), L_{i-1}, L_i) &\ll \sum_{i=1}^k \left((L_i - L_{i-1}) \log(F(L_{i-1})/F(L_i)) + L_i^{1-\delta_0} \right) \\ &= -\log(1-\Delta) \sum_{i=1}^k (L_i - L_{i-1}) + kX^{1-\delta_0} \\ &\ll \Delta(X-Y) + \Delta^{-1} \log(n) X^{1-\delta_0} \\ &\leq X^{1-\delta_0/2} (1 + \log(n)) \\ &\ll n^{1-\delta}. \end{aligned}$$

On the other hand, writing $s = (n - Xa)/bD \geq \frac{a}{bD}n^{1-\delta} \geq \frac{n^{1/2-\delta}}{b}$,

$$\begin{aligned} \sum_{i=1}^k T_\chi(s, F(L_i), L_{i-1}, L_i) &\ll \sum_{i=1}^k \left((L_i - L_{i-1}) \frac{(anf_0s)^\varepsilon}{s} + \phi(f_\chi) L_i^{1-\delta_0} \right) \\ &= (X - Y) \frac{(anf_0s)^\varepsilon}{s} + X^{1-\delta_0/2} \log(n) \phi(f_\chi) \\ &\ll n^{1/2+\delta+2\varepsilon} b (af_0)^\varepsilon + n^{1-\delta_0/2} \log(n) \phi(f_\chi). \end{aligned}$$

This proves that choosing ε and δ appropriately, S_3 is absorbed in the error term. It concludes the proof of Theorem 4.2. \square

5. SOLVING EQUATIONS OF THE FORM $\det(\gamma)^2 f(s, t) + bf^\gamma(x, y) = n$

Let b and n be positive integers, and f a primitive, positive definite, integral, binary quadratic form whose discriminant is fundamental. Let $\gamma \in M_{2 \times 2}(\mathbf{Z})$ be a matrix of rank 2. In this section, we focus on the problem of finding integer solutions of the equation

$$\det(\gamma)^2 f(s, t) + bf^\gamma(x, y) = n.$$

More precisely, we prove the following theorem.

Theorem 5.1 (GRH). *There exists a constant $c > 0$ and an algorithm \mathcal{A} such that the following holds. Let b and n be positive integers, and f a reduced, primitive, positive definite, integral, binary quadratic form whose discriminant is fundamental. Let $\gamma \in M_{2 \times 2}(\mathbf{Z})$ of rank 2 and content 1. Suppose that the factorisation of $\det(\gamma)$ is known, that $\det(\gamma)$, $\text{disc}(f)$ and b are pairwise coprime, and that $\gcd(\det(\gamma)b, n) = 1$. Suppose that $\log n \geq \max(c \cdot \log b, \log(\det(\gamma))^c, \text{disc}(f)^c)$, and either*

- (1) $\log n \geq \omega(n)^c$, or
- (2) *the prime divisors of n are larger than $\text{disc}(f)^c$, $(\log \log b)^c$ and $\log(\det(\gamma))^c$.*

Then $\mathcal{A}(f, \gamma, b, n)$ returns an integer solution $(s, t, x, y) \in \mathbf{Z}^4$ of the equation

$$\det(\gamma)^2 f(s, t) + bf^\gamma(x, y) = n,$$

provided that the equation has a solution modulo $\text{disc}(f^\gamma)$ for which $f(s, t)$ is invertible modulo $\text{disc}(f)$. The algorithm runs in expected polynomial time in $\text{disc}(f)$, $\text{length}(\gamma)$, $\log n$, and the output is random with min-entropy $\Omega(\log n)$.

Cornacchia's algorithm allows to solve equations of the form

$$f(s, t) = z,$$

in time polynomial in $\text{disc}(f)$ and $\log z$ when the factorisation of z is known and a solution exists. We are therefore led to study the solutions (z, x, y) of the equation

$$(2) \quad \det(\gamma)^2 z + bf^\gamma(x, y) = n,$$

where the factorisation of z is known and $z > 0$. Factoring is hard, but primality testing is easy, so we will simply look for solutions where z is prime. Having z prime has another advantage: if χ is the Kronecker symbol of modulus $\text{disc}(f)$, the condition $\chi(z) = 1$ ensures that there is a solution $f'(s, t) = z$ for some f' of same discriminant as f . Replacing this condition by $z \equiv u \pmod{\text{disc}(f)}$ for any u represented by f ensures that z is represented by some f' in the same genus as f . Ensuring that z is represented by f itself will require additional tricks.

5.1. Solutions of $az + bg(x, y) = n$. First, let us lift the delicate primality condition on z . The following proposition allows to sample random solutions of Equation (2) if z is only required to be a positive integer.

Proposition 5.2. *Let g be a primitive, positive definite, integral, binary quadratic form. Let a, b, n be positive integers, and suppose that a divides $\text{disc}(g)$ and $\gcd(a, 2bn) = 1$. Let X be the*

set of integral solutions (z, x, y) of the equation $az + bg(x, y) = n$, with $z > 0$. If there exists a solution modulo a , then X is a disjoint union of $2^{\omega(a)}$ sets X_i , with

$$\#X_i = \frac{\pi n}{ba \operatorname{Vol}(g)} + O\left(\left(\frac{n}{b}\right)^{1/2} + a \operatorname{Vol}(g)\right),$$

and knowing the factorisation of a allows to sample uniformly from any X_i in polynomial time.

Proof. First consider the equation modulo a . Since a divides the discriminant of $g(x, y)$, the latter polynomial splits modulo a and the equation becomes

$$\varepsilon \cdot (\alpha x + \beta y)^2 = n/b \pmod{a},$$

where ε , and at least one of α and β , are invertible. Now, the element $n/b\varepsilon \pmod{a}$ is invertible, so it either has no square root (in which case X is empty), or it has $2^{\omega(a)}$ distinct square roots. Suppose it has square roots. There is a sublattice Λ of index a in \mathbf{Z}^2 such that the space of admissible pairs (x, y) is the disjoint union $\bigsqcup_{i=1}^{\delta} (\Lambda + v_i)$, where the v_i -vectors are representative solutions for the $2^{\omega(a)}$ roots modulo a . Accounting for the condition $z > 0$, it remains to count for each translated lattice $\Lambda + v_i$ the number of points $(x, y) \in \Lambda + v_i$ such that

$$g(x, y) < n/b.$$

From Lemma 3.2, it is equal to

$$\frac{\pi n}{b \operatorname{Vol}(\Lambda) \operatorname{Vol}(g)} + O\left(\left(\frac{n}{b}\right)^{1/2} + \operatorname{Vol}(\Lambda) \operatorname{Vol}(g)\right) = \frac{\pi n}{ba \operatorname{Vol}(g)} + O\left(\left(\frac{n}{b}\right)^{1/2} + a \operatorname{Vol}(g)\right).$$

Let X_i be the solutions stemming from $\Lambda + v_i$. Given the factorisation of a , one can compute all the square roots of $n/b\varepsilon \pmod{a}$. Therefore, to sample uniformly in X_i , apply Lemma 3.3 to sample uniformly a point in the intersection of $\Lambda + v_i$ and the ellipsoid $g(x, y) < n/b$. \square

5.2. Randomisation in the genus. Proposition 5.2 tells us that integer solutions of Equation (2) can be sampled uniformly (up to a small error). We would then be done if a large proportion of these have a z -value which is a prime represented by f . Unfortunately, it is hard to control the primality and representability of these solutions when the form f^γ is fixed. Theorem 4.2 only gives information about the family of equations where f^γ is replaced by any form in its genus. Luckily, we can randomise within the genus thanks to the following two lemmata. Their proofs use the classical correspondence between binary quadratic forms and ideals in quadratic orders; for an account of this theory, we refer the reader to [Cox11, Section 7].

Lemma 5.3 below is useful to deal with forms f^γ of large discriminant.

Lemma 5.3 (GRH). *For any discriminant d and positive integer m , there exists an integer B coprime to md such that any primitive binary quadratic form of discriminant d represents a divisor of B , and $\log B = O_\varepsilon((\log |d|) \cdot ((\log |d|)^{2+\varepsilon} + \omega(m)^{1+\varepsilon}))$. There is an algorithm which samples a form uniformly distributed in the class group, together with a representation by this class of a divisor of B , in time polynomial in $\log |d|$ and $\log m$.*

Proof. We utilise the correspondence between classes of binary quadratic forms and ideal classes in quadratic number fields, and the fact that a form represents $n > 0$ if and only if the corresponding ideal class contains an ideal of norm n . The key is the fact that there is a small bound C such that the ideals of prime norm at most C constitute a generating set of the class group so that the Cayley graph is highly connected: any two vertices are connected by a path of length at most $D = O(\log h(d)) = O(\log |d|)$. From [JMV05, Theorem 1.1] and [JMV05, Corollary 1.3], one can choose $C = O_\varepsilon((\log |d|)^{2+\varepsilon})$ for any $\varepsilon > 0$. However, to construct our B , we wish to consider only prime ideals coprime to m . The same proof as [JMV05] implies that we can choose $C = O_\varepsilon((\log |d|)^{2+\varepsilon} + \omega(m)^{1+\varepsilon})$: replace the estimate (in [JMV05, Equation (2.4)], with $n = 2$)

$$\sum_{N(\mathfrak{p}) \leq x} (\chi(\mathfrak{p}) + \chi(\mathfrak{p})^{-1}) = 2r \frac{x}{\log x} + O\left(x^{1/2} \log(xd)\right)$$

with its simple corollary

$$\sum_{\substack{N(\mathfrak{p}) \leq x \\ (\mathfrak{p}, m) = 1}} (\chi(\mathfrak{p}) + \chi(\mathfrak{p})^{-1}) = 2r \frac{x}{\log x} + O\left(x^{1/2} \log(xd) + \omega(m)\right).$$

We deduce that if p_1, \dots, p_k are all the primes at most C not dividing m , any ideal is equivalent to an ideal of norm $\prod_i p_i^{e_i}$ where $\sum_i e_i \leq D$. The latter is a divisor of $B = \prod_i p_i^D$, and we have

$$\log B = D \sum_i \log p_i \leq D\pi(C) \log(C) = O_\varepsilon\left((\log |d|) \cdot ((\log |d|)^{2+\varepsilon} + \omega(m)^{1+\varepsilon})\right),$$

where $\pi(C) \log(C) = O(C)$ is the prime number theorem.

To sample an ideal with norm dividing B and uniformly distributed in the class group, compute a random walk of length D in the (expander) Cayley graph, as in [JMV05]. \square

The next lemma is similar, but mostly useful when the discriminant is small.

Lemma 5.4 (GRH). *For any discriminant d and positive integer m , there exists an integer B coprime to md such that any primitive binary quadratic form of discriminant d represents a divisor of B , and $\log B = O_\varepsilon(|d|^{1/2+\varepsilon} \log(\omega(m) + 2))$. There is an algorithm which samples a form uniformly distributed in the class group in time polynomial in $|d|$ and $\log m$. Given any class, one can compute a divisor of B together with a representation by this class in time polynomial in $|d|$ and $\log m$.*

Proof. We proceed as above, but instead of doing random walks, we use that class group computations can be done in time polynomial in the discriminant (see [Coh13, Chapter 5]). As already seen in the proof of Lemma 5.3, the class group is generated by the set P_0 of ideal of prime norm at most $C = O_\varepsilon((\log |d|)^{2+\varepsilon} + \omega(m)^{1+\varepsilon})$ not dividing md . In time polynomial in $|d|$, one can compute the class group, and find a minimal subset $P \subseteq P_0$ generating the class group. We have $\#P \leq \log(h)$, where $h = O(|d|^{1/2} \log |d|)$ is the class number. For each class, one can compute a representative that is a product of ideals in P , with exponents at most the class number h . These representatives divide $B = \prod_{\mathfrak{p} \in P} N(\mathfrak{p})^h$, and

$$\log B = h \sum_{\mathfrak{p} \in P} \log N(\mathfrak{p}) \leq h \log(C) \#P = O_\varepsilon(|d|^{1/2+\varepsilon} \log(\omega(m) + 2)).$$

The added 2 avoids the degeneracy at $m = 1$. In polynomial time in $|d|$, one can sample a uniformly random ideal class. Given any class, one can return the corresponding (previously computed) representative that divides B . \square

Let a, b, u, v and n be positive integers, and f and g two primitive, positive definite, integral, binary quadratic forms. We are looking for a solution (z, x, y) to the equation

$$(3) \quad az + bg(x, y) = n,$$

where $z > 0$ is represented by f . A condition of the form $\ell \equiv u \pmod{\text{disc}(f)}$ can ensure that ℓ is represented by the genus of f , but this is not enough for ℓ to be represented by f itself. The following trick deals with this difficulty.

Lemma 5.5. *There is an integer B_0 coprime to $2nb \text{disc}(g)$ such that the following holds. For any $\rho \in M_{2 \times 2}(\mathbf{Z})$ of determinant B_0 and content 1, given an integral solution (ℓ, x, y) of*

$$a \det(\rho)^2 \ell + bg^\rho(x, y) = n,$$

with ℓ a prime represented by the genus of f , one can compute an integral solution (s, t, x, y) of

$$af(s, t) + bg(x, y) = n$$

in expected polynomial time in $\text{disc}(f)$ and the size of the input.

Proof. Let B_0 be the integer from Lemma 5.4, for $d = \text{disc}(f)$ and $m = 2nb \text{disc}(g)$. Let $\rho \in M_{2 \times 2}(\mathbf{Z})$ of determinant B_0 and content 1, and suppose we have a solution (ℓ, x, y) of

$$aB_0^2\ell + bg^\rho(x', y') = n,$$

with ℓ a prime represented by the genus of f . Then, Cornacchia's algorithm allows to find h in the genus of f and integers (s_0, t_0) such that $\ell = h(s_0, t_0)$. Let k such that $[k]^2 = [h]^{-1}[f]$ and (s_1, t_1) such that $k(s_1, t_1) = d \mid B_0$. Then, we can compute¹ (s_2, t_2) such that

$$B_0^2\ell = (B_0/d)^2k(s_1, t_1)^2h(s_0, t_0) = (B_0/d)^2f(s_2, t_2) = f(s_2B_0/d, t_2B_0/d).$$

With $(s, t) = (s_2B_0/d, t_2B_0/d)$ and $(x, y) = \rho(x', y')$, we get $af(s, t) + bg(x, y) = n$. \square

Therefore, it is sufficient to study solutions of Equation (3) where ℓ is a prime represented by the genus of f , up to replacing a with aB_0^2 and g with g^ρ .

For the rest of this section, consider all notation and conditions from Theorem 5.1, and let $a = \det(\rho\gamma)^2$ and $g = f^{\rho\gamma}$, with ρ as in Lemma 5.5. By working carefully at each prime factor of B_0 , one can craft ρ in a way that ensures the local solvability of the equation $a\ell + bg(x_0, y_0) = n$. Observe that in general, the B_0 constructed in Lemma 5.5 satisfies

$$\log(B_0) = O_\varepsilon\left(|\text{disc}(f)|^{1/2+\varepsilon} \log(2 + \log(nb \text{disc}(g)))\right),$$

but if Condition (2) holds, we can obtain

$$\log(B_0) = O_\varepsilon\left(|\text{disc}(f)|^{1/2+\varepsilon} \log(2 + \log(b \text{disc}(g)))\right)$$

by choosing $m = 2b \text{disc}(g)$ in the application of Lemma 5.4; the condition that $(B_0, n) = 1$ is then enforced by the fact that the prime divisors of n are larger than those of B_0 .

Now, let B be the integer from Lemma 5.3 for $d = \text{disc}(g)$, and $m = n$ if Condition (1) holds, or $m = 1$ if Condition (2) holds. In either case, Lemma 5.3 ensures that $\log(B) = O((\log n)^{\frac{9+3\varepsilon}{2c}})$, either by bounding $\omega(n)$ with Condition (1), or with $\omega(1) = 0$ in the other case. Also, even in the case $m = 1$, we have $(B, n) = 1$, as Condition (2) ensures that the prime factors of n are all larger than those of B . Consider the sets

$$\mathcal{X} = \{(z, x_1, y_1, h, x_0, y_0, k) \mid az + B^2bh(x_1, y_1) = n, z > 0, [k^2h] = [g], \text{ and } k(x_0, y_0) \text{ divides } B\},$$

$$X = \{(z, x, y, [h]) \mid az + B^2bh(x, y) = n, z > 0, \text{ and } h \text{ is in the genus of } g\},$$

$$\mathcal{S} = \{(\ell, x, y, [h]) \in X \mid \ell \text{ is prime and } \ell \equiv u \pmod{\text{disc}(f)}\},$$

where $[h]$ denotes the class of h in the class group. There is a natural surjection $\pi : \mathcal{X} \rightarrow X$.

Lemma 5.6. *Given a tuple $T \in \mathcal{X}$ such that $\pi(T) \in \mathcal{S}$, one can compute a solution (ℓ, x, y) of Equation (3) in polynomial time, where $\ell > 0$ is a prime number such that $\ell \equiv u \pmod{v}$.*

Proof. Let $T = (\ell, x_1, y_1, h, x_0, y_0, k) \in \mathcal{X}$ such that $\pi(T) \in \mathcal{S}$. We then have

$$a\ell + B^2bh(x_1, y_1) = n.$$

Since $[k^2h] = [g]$, one can compute x_2 and y_2 such that $k(x_0, y_0)^2h(x_1, y_1) = g(x_2, y_2)$ (see the footnote from the proof of Lemma 5.5). We obtain

$$B^2h(x_1, y_1) = (B/d)^2k(x_0, y_0)^2h(x_1, y_1) = (B/d)^2g(x_2, y_2) = g(x_2B/d, y_2B/d),$$

where $d = k(x_0, y_0)$. Then, $(\ell, x_2B/d, y_2B/d)$ is a solution of Equation (3). \square

Lemma 5.7. *Suppose a is an odd prime power dividing $\text{disc}(g)$ and coprime to b . There is an algorithm that samples elements T in \mathcal{X} such that $\pi(T)$ is close to uniformly distributed in X , and runs in expected polynomial time. More precisely, the probability of any $x \in X$ is between $1/(2\#X)$ and $3/(2\#X)$.*

¹Using the Gauss composition law, find S, T such that $(k^2h)(S, T) = k(s_0, t_0)^2h(s_1, t_1)$. Then, find $\gamma, \gamma' \in \text{SL}_2(\mathbf{Z})$ such that $(k^2h)^\gamma$ and $g^{\gamma'}$ are reduced (see [Coh13, Algorithm 5.4.2]). Since they are in the same class and reduced, we actually have $(k^2h)^\gamma = g^{\gamma'}$. Finally, let $(s_2, t_2) = \gamma'\gamma^{-1}(S, T)$.

Proof. Generate a uniformly random class $[k]$ in $\text{Cl}(\text{disc}(g))$ together with a divisor d of B which it represents as $k(z_0, t_0) = d$. Let $h \in [k^{-2}g]$. Following Proposition 5.2, sample a uniformly random integer solution (z, x_1, y_1) of $az + bB^2h(x_1, y_1) = n$, with $z > 0$. Note that the local solvability at a in Proposition 5.2 is satisfied for any k (because h is always in the same genus), so the deviation from uniformity only comes from the error term. \square

5.3. Proof of Theorem 5.1. Recall that $a = \det(\rho\gamma)^2$ and $g = f^{\rho\gamma}$. From the assumption on the solutions modulo $\text{disc}(f)$, there exists an invertible $u \bmod \text{disc}(f)$ represented by the genus of f and such that $\chi\left(\frac{n-ua}{b}\right) \neq -1$ (an exhaustive search finds it in time polynomial in $\text{disc}(f)$). From Lemma 5.7, we can efficiently sample $T \in \mathcal{X}$ such that $\pi(T)$ is uniform in X . From Lemma 5.6, we are done as soon as $\pi(T) \in \mathcal{S}$. Indeed, such a T gives a solution of

$$n = az + bg(x_0, x_0) = \det(\gamma)^2 \det(\rho)^2 \ell + b(f^\gamma)^\rho(x_0, y_0),$$

giving rise via Lemma 5.5 to a solution of $\det(\gamma)f(s, t) + bf^\gamma(x, y) = n$. Then, it only remains to prove that $\#X/\#\mathcal{S}$ is small. On one hand, from Proposition 5.2, we have

$$\#X \ll \frac{n2^{\omega(a)}h}{B^2ba|\text{disc}(g)|^{1/2}} \ll \frac{n \log |a \text{disc}(f)|}{B^2ba},$$

where $h \ll 2^{-\omega(a)}|\text{disc}(g)|^{1/2} \log |\text{disc}(g)|$ is the number of classes in the genus of g . On the other hand, the local solvability of the equation modulo $\text{disc}(g)$ together with Corollary 4.3 implies that there is a constant c' such that

$$\#\mathcal{S} \gg \sum_{\substack{u' \bmod \text{disc}(f) \det(\rho\gamma) \\ u' \equiv u \bmod \text{disc}(f) \\ (u', \det(\rho\gamma))=1}} \frac{n \left(1 + \chi\left(\frac{n-ua}{b}\right)\right)}{B^2ba\phi(\text{disc}(f) \det(\rho\gamma))(\log n)^{c'}} = \frac{n \left(1 + \chi\left(\frac{n-ua}{b}\right)\right)}{B^2ba\phi(\text{disc}(f))(\log n)^{c'}}.$$

Note that since $a = \det(\rho\gamma)^2$ is coprime to n , the condition $\gcd(n - au', \det(\rho\gamma)) = 1$ (required for Corollary 4.3) is satisfied for any u' . The theorem follows. \square

5.4. Representing integers in special orders. Theorem 5.1 has an immediate, but important corollary. Recall that for any prime p , we denote by \mathcal{O}_0 the special order in $B_{p,\infty}$ defined in Lemma 2.3.

Corollary 5.8 (GRH). *There is a constant c and an algorithm \mathcal{A} such that the following holds. For any prime p and integer n with $\log n > (\log p)^c$, if either $\log n \geq \omega(n)^c$, or the prime divisors of n are larger than $(\log p)^c$, then the algorithm \mathcal{A} finds an element $\alpha \in \mathcal{O}_0$ of reduced norm n , and runs in expected polynomial time in $\log p$ and $\log n$. The output α is random with min-entropy at least $\Omega(\log n)$.*

Proof. With notations as in Lemma 2.3, the order \mathcal{O}_0 contains the elements $1, \omega, j$ and ωj , and

$$\text{Nrd}(s + t\omega + xj + y\omega j) = f(s, t) + pf(x, y).$$

When $\gcd(n, p) = 1$, the result follows from Theorem 5.1 with $b = p$ and γ the identity matrix. In the case where prime divisors of n are at least $(\log p)^c$, we use that $\text{disc}(f) = O((\log p)^2)$, which allows to satisfy Condition (2). Since $\text{Nrd}(j) = p$, the general result follows by multiplicativity of the reduced norm. \square

6. SOLVING THE QUATERNION PATH PROBLEM

For the rest of the article we consider the quaternion algebra $B_{p,\infty}$, with basis $1, i, j, ij$, as defined in Section 2.2, and \mathcal{O}_0 is the special maximal order defined in Lemma 2.3. In this section, we consider the QUATERNIONPATH problem. Since computing connecting ideals between two maximal orders is easy (see [KV10, Algorithm 3.5]), it is sufficient to consider the following problem: given a maximal order \mathcal{O} in $B_{p,\infty}$, a left \mathcal{O} -ideal I , and an integer N , find an ideal J equivalent to I such that $\text{Nrd}(J) = N$. As noted in [KLPT14, Section 4.6], the general case reduces to the case $\mathcal{O} = \mathcal{O}_0$, so we focus on this special maximal order.

6.1. Random walks between ideal classes. As a first step towards a rigorous algorithm, we start by randomising the input ideal, thereby avoiding pathological cases.

Definition 6.1 (Brandt graph). Let p be a prime number, and \mathcal{O} a maximal order in $B_{p,\infty}$. Let I, J be two left \mathcal{O} -ideals. We say J is an ℓ -neighbor of I if $J \subseteq I$ and $\text{Nrd}(J) = \ell \text{Nrd}(I)$. The ℓ -Brandt graph is the graph with vertices $\text{Cls}(\mathcal{O})$ and an edge from $[I_i]$ to $[J]$ for each ℓ -neighbor $J \subseteq I_i$, where $(I_i)_{i=1}^{\#\text{Cls}(\mathcal{O})}$ is a list of ideal class representatives.

Through the Deuring correspondence, the ℓ -Brandt graph is isomorphic to the ℓ -isogeny graph (up to the action of $\text{Gal}(\mathbf{F}_{p^2}/\mathbf{F}_p)$). It is common in isogeny-based cryptography to compute random walks on these graphs: at each step on the walk, the current vertex is an elliptic curve E , one chooses uniformly at random one of the $\ell + 1$ outgoing isogenies, and the next vertex is its target. Equivalently, given a left \mathcal{O} -ideal I , one can choose uniformly at random one of the $\ell + 1$ left submodules $M \subset I/\ell I$, and the next vertex is $M + \ell I$.

Theorem 6.2 ([GPS20, Theorem 1]). Let p be a prime number, and \mathcal{O} a maximal order in $B_{p,\infty}$. Let N_p be the size of the ideal class set of \mathcal{O} . Let I be the ideal obtain from a random walk of norm $n = \prod_i \ell_i^{e_i}$. Then, for any ideal class C , we have

$$\left| \Pr[I \in C] - \frac{1}{N_p} \right| \leq \prod_i \left(\frac{2\sqrt{\ell_i}}{\ell_i + 1} \right)^{e_i}.$$

Proof. This is precisely [GPS20, Theorem 1], translated from isogenies to quaternions through the Deuring correspondence. It is a consequence of the fact that each ℓ_i -Brandt graph (or ℓ_i -isogeny graph) has the Ramanujan property. \square

6.2. Solving the quaternion analog of the isogeny-path problem. The main result of this section is the following theorem.

Theorem 6.3 (GRH). There exists an integer c such that Algorithm 2 is correct and runs in expected polynomial time in $\log p$, $\log n_i$, $\log \text{Nrd}(I)$ and ℓ for all inputs satisfying $\log n_i \geq (\log p)^c$, and $n_2 \ell^e \not\equiv 2, 4 \pmod{8}$ for $e \in \{0, 1\}$, and either $\log n_2 \geq \omega(n_2)^c$, or all prime divisors of $n_2 \ell$ are larger than $(\log p)^c$.

Algorithm 2 $\text{EQUIVIDEAL}_c(I, n_1, n_2)$

Require: A left ideal I in the special maximal order \mathcal{O}_0 , positive integers n_1, n_2 , and a prime ℓ .

Ensure: An equivalent ideal J of norm $n_1 n_2$ or $n_1 n_2 \ell$.

- 1: Define R , ω and f as in Lemma 2.3.
 - 2: **while** β has not been found **do**
 - 3: $I' \leftarrow$ the endpoint $I' \subset I$ of a random walk of norm n_1 in the Brandt graph; {Theorem 6.2}
 - 4: $(I'', \rho) \leftarrow$ an ideal I'' equivalent to I' , of prime norm $N \in [p^c, p^{2c}]$ such that ℓ is a non-quadratic residue modulo N , and the element $\rho \in I'$ such that $I'' = I' \bar{\rho} / \text{Nrd}(I')$; {Proposition 3.8}
 - 5: $\gamma \leftarrow$ an element $\gamma \in \mathcal{O}_0$ such that $\text{Nrd}(\gamma) = N$; {Corollary 5.8}
 - 6: $\beta \leftarrow$ an element $\beta \in R$ such that $I'' = \mathcal{O}_0 N + \mathcal{O}_0 \gamma \beta j$ if it exists;
 - 7: **end while**
 - 8: $\Gamma \leftarrow$ a matrix in $M_{2 \times 2}(\mathbf{Z})$ such that $\mathbf{Z} \beta + R N = \{x + y \omega \mid (x, y) \in \Gamma \mathbf{Z}^2\}$;
 - 9: $(s, t, x, y) \leftarrow$ an integral solution of $N^2 f(s, t) + p f^\Gamma(x, y) = n_2 \ell^e$ for some $e \in \{0, 1\}$; {Theorem 5.1}
 - 10: $(x', y') \leftarrow \Gamma(x, y)$;
 - 11: $\alpha \leftarrow (s + t \omega) N + (x' + y' \omega) j$;
 - 12: $\delta \leftarrow \rho \gamma \alpha / N \in I' \subset I$;
 - 13: **return** $J = I \bar{\delta} / \text{Nrd}(I)$.
-

Proof. The efficiency and correctness of most steps are already justified by the various results referred to in the comments of Algorithm 2. The constraints on n_2 come from Theorem 5.1.

Step 3. From Theorem 6.2, there is a constant c such that if $\log n_1 \geq (\log p)^c$, then I' is in any given class set with probability between $1/2N_p$ and $3/2N_p$, with N_p the number of classes.

Step 5. Corollary 5.8 requires $N = \text{Nrd}(I'')$ to be large enough; therefore, in constructing I'' , we resort to Proposition 3.8 rather than Theorem 3.7. This issue is dealt with differently in [KLPT14]: they solve an equation of the form $\text{Nrd}(\gamma) = Nn_3$ for some large enough n_3 . Our approach has a theoretical advantage: Corollary 5.8 ensures that γ has large entropy, which allows to avoid corner cases in Step 6. More precisely, let us prove that $\mathcal{O}_0\gamma/\mathcal{O}_0N$ has large entropy. It is sufficient to prove that the map $\gamma \mapsto \mathcal{O}_0\gamma/\mathcal{O}_0N$, for $\gamma \in \mathcal{O}_0$ of norm N , has small fibre. Suppose that $\mathcal{O}_0\gamma/\mathcal{O}_0N = \mathcal{O}_0\gamma'/\mathcal{O}_0N$. Then, there exists $x, y \in \mathcal{O}_0$ such that $\gamma' = x\gamma + yN$. Then,

$$\gamma' = x\gamma + yN = (x + y\bar{\gamma})\gamma.$$

Comparing norms, we deduce $\text{Nrd}(x + y\bar{\gamma}) = 1$, hence $\gamma' \in \mathcal{O}_0^\times \gamma$. Since $\#\mathcal{O}_0^\times \leq 6$, the map $\gamma \mapsto \mathcal{O}_0\gamma/\mathcal{O}_0N$ is $O(1)$ -to-1, which proves that $\mathcal{O}_0\gamma/\mathcal{O}_0N$ has large entropy.

Step 6. This step is solved with elementary linear algebra, as described in [KLPT14, Section 4.3]. The method of [KLPT14, Section 4.3] succeeds under the assumption that $I''/N\mathcal{O}_0$ and $\mathcal{O}_0\gamma/N\mathcal{O}_0$ are distinct from the (at most two) fixed points for the action of $(R/NR)^\times$. This is heuristically assumed in [KLPT14], but with our new methods, we can prove it. The large entropy of $\mathcal{O}_0\gamma/N\mathcal{O}_0$ ensures that with good probability, it is not one of the two fixed points. From Step 3, I' is close to uniformly distributed in the class set, so with overwhelming probability it is not equivalent to an ideal induced by an R -ideal (i.e., to an ideal of the form $\mathcal{O}_0\mathfrak{a}$ for some R -ideal \mathfrak{a}). It is then also the case of I'' , so with good probability, it is not a fixed point either (the ideals of norm N that are fixed points are induced by the R -ideals above N).

Step 9. Most conditions for Theorem 5.1 are already met. The value $e \in \{0, 1\}$ is determined by the constraint that the equation must have a solution modulo N . It remains to justify that the equation does have a solution in $G = \mathbf{Z}/\text{disc}(f)\mathbf{Z}$ for which $f(s, t)$ is invertible. Suppose $p \equiv 1 \pmod{8}$, so from Lemma 2.2, $\text{disc}(f)$ is a negative odd prime. From [Cox11, Theorem 3.15], there is only one genus of forms of discriminant $\text{disc}(f)$, so f represents all quadratic residues in G . Since $(N, \text{disc}(f)) = 1$, the form f^Γ also represents all the quadratic residues in G . Lemma 2.2 implies that p is not a quadratic residue. Any element in G is the sum of a quadratic residue and a quadratic non-residue, so $n_2\ell^e$ also is, and we are done. Similarly, if $\text{disc}(f) = 4$ there is a solution when $n_2\ell^e \not\equiv 2 \pmod{4}$, and if $\text{disc}(f) = 8$ there is a solution when $n_2\ell^e \not\equiv 4 \pmod{8}$. \square

Remark 3. Given a prime ℓ , one can choose n_1 and n_2 to be large enough powers of ℓ , so Algorithm 2 straightforwardly specialises to the power-of- ℓ variant ℓ -QUATERNIONPATH. We deal with the powersmooth variant B -PSQUATERNIONPATH in the next section.

6.3. Finding power-smooth paths. In Theorem 6.3, the integers n_i either have very few prime factors, or the prime factors are not too small. This seems to come at odds with a major application of [KLPT14]: constructing ideals of powersmooth norm. We now prove that it is not an issue, and we can indeed solve the B -PSQUATERNIONPATH variant.

Theorem 6.4 (GRH). *There exists an integer c and an algorithm \mathcal{A} such that the following holds. On input a left \mathcal{O}_0 -ideal I , the algorithm outputs an equivalent ideal J whose norm is $(\log p)^c$ -powersmooth, and runs in expected polynomial time in $\log p$ and $\log \text{Nrd}(I)$.*

Proof. It is sufficient to prove that one can find suitable powersmooth integers n_i to apply Theorem 6.3. Let c_0 be the constant from Theorem 6.3, and let $\delta > 0$ be some parameter to be adjusted. Let $c = 2(c_0 + \delta)$. We need to construct two $(\log p)^{c_0+\delta}$ -powersmooth integers n_i

such that $\log n_i > (\log p)^{c_0}$ and prime divisors of n_2 are larger than $(\log p)^{c_0}$. We choose

$$n_1 = n_2 = \prod_{(\log p)^{c_0} < \ell < (\log p)^{c_0 + \delta}} \ell^{\frac{(c_0 + \delta) \log p}{\log \ell}}.$$

Then, from the prime number theorem (with Riemann's hypothesis),

$$\log n_2 \geq \sum_{(\log p)^{c_0} < \ell < (\log p)^{c_0 + \delta}} \log \ell = (\log p)^{c_0} \left((\log p)^\delta - 1 \right) + O \left((\log p)^{\frac{c_0 + \delta}{2}} (\log \log p)^2 \right).$$

Choosing δ large enough ensures that $\log n_2 > (\log p)^{c_0}$, which concludes the proof. \square

7. MAXIMAL ORDER AND ISOGENY PATH ARE EQUIVALENT

In this section and the next, we prove that ℓ -ISOGENYPATH, ENDRING and MAXORDER are all equivalent, under probabilistic polynomial-time reductions. We start in this section by showing that ℓ -ISOGENYPATH is equivalent to MAXORDER.

7.1. Maximal Order reduces to Isogeny Path. From Lemma 2.6, we know how to translate powersmooth \mathcal{O}_0 -ideals into isogenies. The following lemma deals with the converse direction.

Lemma 7.1. *Let \mathcal{O}_0 and E_0 as in Lemmata 2.3 and 2.5. There exists an algorithm which, given an isogeny $\varphi : E_0 \rightarrow E$ of degree $\prod_i \ell_i^{e_i}$, returns the corresponding left \mathcal{O}_0 -ideal I_φ . The complexity of this algorithm is polynomial in $\log p$ and $\max_i(\ell_i^{e_i})$ (if $p \equiv 1 \pmod{8}$, we assume GRH).*

Proof. A proof of this lemma was first given in [Kri20], building upon the heuristic result [GPS20, Lemma 6]. It can also be seen as a consequence of Lemma 2.6: for each i ,

- (1) Enumerate the set S_i of all left \mathcal{O}_0 -ideals of norm $\ell_i^{e_i}$ (see [KV10]);
- (2) For each $J \in S_i$, compute the corresponding isogeny φ_J with Lemma 2.6, and if $\ker(\varphi_J) = \ker(\varphi) \cap E_0[\ell_i^{e_i}]$, let $I_i = J$.

Finally, return $I_\varphi = \bigcap_i I_i$. Of course, this guessing approach is not as efficient as the method proposed in [GPS20, Lemma 6], but it is still polynomial in $\max_i(\ell_i^{e_i})$. \square

To prove that MAXORDER reduces to ℓ -ISOGENYPATH, we show that an isogeny between E and the special curve E_0 (of known endomorphism ring) allows to recover the endomorphism ring of E . A heuristic version of this approach was described in [DMPS19].

Algorithm 3 Reducing MAXORDER to ℓ -ISOGENYPATH

Require: A supersingular elliptic curves E/\mathbf{F}_{p^2} , with $p \neq \ell$. We suppose there is an algorithm $\mathcal{A}_{\ell\text{-ISOGENYPATH}}$ that solves the ℓ -ISOGENYPATH problem.

Ensure: A basis of an order in $B_{p,\infty}$ isomorphic to $\text{End}(E)$.

- 1: $c \leftarrow$ the constant from Theorem 6.4;
 - 2: $(\mathcal{O}_0, E_0) \leftarrow$ the special order and curve from Lemmata 2.3 and 2.5;
 - 3: $\varphi \leftarrow \mathcal{A}_{\ell\text{-ISOGENYPATH}}(E_0, E)$, with $\varphi = \varphi_e \circ \dots \circ \varphi_1$, and $\deg(\varphi_i) = \ell$;
 - 4: $\psi_0 \leftarrow$ the identity isogeny $E_0 \rightarrow E_0$;
 - 5: **for** $i = 1, \dots, e$ **do**
 - 6: $I_i \leftarrow$ the ideal corresponding to $\varphi_i \circ \psi_{i-1}$; {Lemma 7.1}
 - 7: $J_i \leftarrow$ an ideal equivalent to I_i , with $(\log p)^c$ -powersmooth norm; {Theorem 6.4}
 - 8: $\psi_i \leftarrow$ the isogeny corresponding to J_i ; {Lemma 2.6}
 - 9: **end for**
 - 10: $\mathcal{O} \leftarrow \mathcal{O}_R(J_e)$; {[Rón92, Theorem 3.2]}
 - 11: **return** A basis of \mathcal{O} .
-

Theorem 7.2 (GRH). *The reduction in Algorithm 3 is correct and runs in expected polynomial time in $\log p$ and the output size of $\mathcal{A}_{\ell\text{-ISOGENYPATH}}$, plus one call to $\mathcal{A}_{\ell\text{-ISOGENYPATH}}$.*

Proof. Write $\varphi_i : E_{i-1} \rightarrow E_i$, with $E_e = E$. At each step of the loop, we have that J_i is equivalent to the ideal corresponding to $\varphi_i \circ \dots \circ \varphi_1$, hence $\mathcal{O}_R(J_i) \cong \text{End}(E_i)$. This proves the correctness. The running time follows from the results cited at each step of Algorithm 3. \square

7.2. Isogeny Path reduces to Maximal Order. To reduce ℓ -ISOGENYPATH to MAXORDER, we prove that one can translate left \mathcal{O}_0 -ideals of norm a power of ℓ to the corresponding isogeny.

Algorithm 4 Translating a left \mathcal{O}_0 -ideal of prime-power norm to an isogeny

Require: A left \mathcal{O}_0 -ideal I of norm ℓ^e , with $\ell \neq p$ prime, and $\ell \nmid I$.

Ensure: The corresponding isogeny φ_I .

- 1: $c \leftarrow$ the constant from Theorem 6.4;
 - 2: **for** $i = 1, \dots, e$ **do**
 - 3: $I_i \leftarrow I + \mathcal{O}_0 \ell^i$;
 - 4: $J_i \leftarrow$ an ideal equivalent to I_i , with $(\log p)^c$ -powersmooth norm; {Theorem 6.4}
 - 5: $\psi_i \leftarrow$ the isogeny corresponding to J_i ; {Lemma 2.6}
 - 6: $E_i \leftarrow$ target(ψ_i);
 - 7: $\varphi_i \leftarrow$ the ℓ -isogeny from E_{i-1} to E_i ; {see [Vél71]}
 - 8: **end for**
 - 9: **return** $\varphi_e \circ \dots \circ \varphi_1$.
-

Lemma 7.3. *Algorithm 4 is correct and runs in expected polynomial time in $\log p$, ℓ and e (if $p \equiv 1 \pmod 8$, we assume GRH).*

Proof. Heuristic versions of this strategy have already appeared in the literature (for instance as a part of [EHL⁺18, Algorithm 7]). Using Theorem 6.4 instead of [KLPT14] makes it rigorous. \square

Algorithm 5 Reducing ℓ -ISOGENYPATH to MAXORDER

Require: Two supersingular elliptic curves E_1 and E_2 over \mathbf{F}_{p^2} . We suppose we are given the two MAXORDER-solutions \mathcal{O}_1 and \mathcal{O}_2 , maximal orders in $B_{p,\infty}$ isomorphic to $\text{End}(E_1)$ and $\text{End}(E_2)$ respectively.

Ensure: An ℓ -isogeny path from E_1 to E_2 .

- 1: $c \leftarrow$ the constant from Theorems 6.3;
 - 2: $e \leftarrow \left\lceil \frac{(\log p)^c}{\log \ell} \right\rceil$;
 - 3: $(\mathcal{O}_0, E_0) \leftarrow$ the special order and curve from Lemmata 2.3 and 2.5;
 - 4: **for** $i = 1, 2$ **do**
 - 5: $I_i \leftarrow I(\mathcal{O}_0, \mathcal{O}_i)$ the ideal connecting \mathcal{O}_0 and \mathcal{O}_i ; {[KV10, Algorithm 3.5]}
 - 6: $J_i \leftarrow \text{EQUIVIDEAL}_c(I_i, \ell^e, \ell^e, \ell)$; {Theorem 6.3}
 - 7: $\varphi_i \leftarrow$ the isogeny corresponding to J_i ; {Lemma 7.3}
 - 8: **end for**
 - 9: **return** $\varphi_2 \circ \hat{\varphi}_1$.
-

Theorem 7.4 (GRH). *Algorithm 5 is correct and runs in expected polynomial time in $\log p$, ℓ and in the length of the two provided MAXORDER-solutions \mathcal{O}_1 and \mathcal{O}_2 .*

Proof. The reduction is almost the same as [EHL⁺18, Algorithm 7], but using Theorems 6.3 and 6.4 instead of [KLPT14]. Note that we reduce to MAXORDER whereas [EHL⁺18, Algorithm 7] reduces to ENDRING. However, in [EHL⁺18, Algorithm 7], the algorithm solving ENDRING is only used to recover the maximal orders \mathcal{O}_1 and \mathcal{O}_2 via the reduction from MAXORDER to ENDRING. We simply short-circuit the chain of reductions. \square

8. ENDOMORPHISM RING IS EQUIVALENT TO MAXIMAL ORDER

We finally prove the equivalence between ENDRING and MAXORDER.

8.1. Endomorphism Ring reduces to Maximal Order. We start with the simplest direction, which can readily be adapted from previous heuristic reductions with our new tools.

Algorithm 6 Reducing ENDRING to MAXORDER, with a parameter $\delta > 0$

Require: A supersingular elliptic curve E/\mathbf{F}_{p^2} , with $p \neq \ell$. We suppose we are given the MAXORDER-solution \mathcal{O} , a maximal order in $B_{p,\infty}$ isomorphic to $\text{End}(E)$.

Ensure: Four endomorphisms of E that generate $\text{End}(E)$.

- 1: $I \leftarrow I(\mathcal{O}_0, \mathcal{O})$, the ideal connecting \mathcal{O}_0 to \mathcal{O} ; {[KV10, Algorithm 3.5]}
 - 2: $c \leftarrow$ the constant from Theorem 6.4;
 - 3: $J \leftarrow$ an ideal equivalent to I , with $(\log p)^c$ -powersmooth norm; {Theorem 6.4}
 - 4: $\mathcal{O}' \leftarrow \mathcal{O}_R(J)$ the right-order of J ; {[Rón92, Theorem 3.2]}
 - 5: $(\beta_i)_{i=1}^4 \leftarrow$ a basis of \mathcal{O}' ;
 - 6: $(\alpha_i)_{i=1}^4, (\phi_i)_{i=1}^4 \leftarrow$ the special basis $(\alpha_i)_{i=1}^4$ of \mathcal{O}_0 from Lemma 2.3, with the corresponding endomorphisms $\phi_i \in \text{End}(E_0)$ from Lemma 2.5;
 - 7: $(c_{ij})_{i,j=1}^4 \leftarrow$ integers such that $\text{Nrd}(J)\beta_i = \sum_{j=1}^4 c_{ij}\alpha_j$ for $i = 1, \dots, 4$;
 - 8: $\varphi \leftarrow$ the isogeny corresponding to J ; {Lemma 2.6}
 - 9: **return** $(N, \varphi, (c_{ij})_{i,j})$, which represents the endomorphisms $\frac{1}{N} \sum_{j=1}^4 c_{ij}\varphi\phi_j\hat{\varphi}$.
-

Theorem 8.1 (GRH). *Algorithm 6 is correct and runs in expected polynomial time in $\log p$ and in the length of the provided MAXORDER-solution \mathcal{O} .*

Proof. The algorithm is the same as [EHL⁺18, Algorithm 4], but using Theorem 6.3 instead of [KLPT14]. In particular, it is proven in [EHL⁺18, Lemma 3] that $(N, \varphi, (c_{ij})_{i,j})$ is an efficient representation of the basis. \square

8.2. Maximal Order reduces to Endomorphism Ring. Finally, we prove that MAXORDER reduces to ENDRING. The most delicate issue is that the corresponding heuristic reduction [EHL⁺18, Algorithm 6] requires the factorisation of large integers, a task that in the worst case cannot be solved in polynomial time (to the best of our knowledge). We modify the reduction to provably avoid all hard factorisations. To do so, we force the corresponding integers to be prime, by leveraging Proposition 3.5 and an explicit parameterisation of solutions of quadratic forms. We start with a lemma, and introducing some handy notation.

Lemma 8.2. *Given two endomorphisms α and β in an efficient representation, one can compute $\langle \alpha, \beta \rangle$ in time polynomial in the length of the representation of α and β , and in $\log p$.*

Proof. This is proven in [EHL⁺18, Lemma 4]. Recall that an efficient representation means that there is an algorithm that evaluates $\alpha(P)$ for any $P \in E(\mathbf{F}_{p^k})$ in time polynomial in the length of the representation of α and in $k \log p$. Also, the length of an efficient representation of α is $\Omega(\log(\deg(\alpha)))$ (which rules out exotic representations where the number of bits of $\langle \alpha, \beta \rangle$ would be exponential in the length of the input). \square

Notation 1. Given two quadratic forms f and g , we write $f \oplus g$ their orthogonal sum, defined as

$$(f \oplus g)(x, y) = f(x) + g(y).$$

We extend this notation naturally to $U \oplus V$ or $G \oplus H$ for quadratic spaces U and V or Gram matrices G and H .

Notation 2. We write $\langle a_1, \dots, a_r \rangle$ the quadratic form whose Gram matrix is $\text{diag}(a_1, \dots, a_r)$.

Theorem 8.3 (GRH). *Algorithm 7 is correct and runs in expected polynomial time in $\log p$, and in the length of the provided ENDRING-solution $(\beta_i)_{i=1}^4$.*

Proof. Let us go through the reduction step by step.

Algorithm 7 Reducing MAXORDER to ENDRING

Require: A supersingular elliptic curve E/\mathbf{F}_{p^2} . We suppose we are given the ENDRING-solution $(\beta_i)_{i=1}^4$, a list of four endomorphisms that generate $\text{End}(E)$.

Ensure: A basis of an order in $B_{p,\infty}$ isomorphic to $\text{End}(E)$.

- 1: $G_0 \leftarrow (\langle \beta_i, \beta_j \rangle)_{i,j=1}^4$ the Gram matrix of $(\beta_i)_{i=1}^4$;
- 2: Find a change of basis such that $A^t G_0 A = \langle 1 \rangle \oplus G$, where G is integral and $\text{disc}(G)$ is only divisible by p and 2;
- 3: Solve $x^t G x = q(a\ell)^2$, where $x \in \mathbf{Z}^3$ is primitive, ℓ is prime (or $\ell = 1$) and a may only be divisible by the primes 2 and p ;
- 4: Find a change of basis $B = \langle 1 \rangle \oplus B'$ such that $B^t(\langle 1 \rangle \oplus G)B = \langle 1, q \rangle \oplus H$, where H is integral, and $\text{disc}(H)$ is only divisible by 2, p , q and ℓ .
- 5: Solve $y^t H y = p$ with $y \in \mathbf{Q}^2$;
- 6: $\iota \leftarrow (A(0, \frac{x}{a\ell}))^t (\beta_i)_{i=1}^4$;
- 7: $\pi \leftarrow (AB(0, 0, y))^t (\beta_i)_{i=1}^4$;
- 8: $\kappa \leftarrow \iota \circ \pi$;
- 9: $\Phi : \text{End}(E) \otimes \mathbf{Q} \rightarrow B_{p,\infty}$, the isomorphism sending $1, \iota, \pi, \kappa$ to $1, i, j, ij$;
- 10: **return** $(\Phi(\beta_i))_{i=1}^4$.

Step 1. The Gram matrix G_0 of $(\beta_1, \beta_2, \beta_3, \beta_4)$ can be computed via Lemma 8.2.

Step 2. First recall that $\text{disc}(G_0) = \text{disc}(\text{End}(E)) = p^2$. This step follows from the fact that the endomorphisms $(2\beta_i - \text{tr}(\beta_i))_{i=1}^4$ generate the (rank 3) orthogonal complement of 1 in $\mathbf{Z} + 2\text{End}(E)$ (an order of discriminant only divisible by p and 2).

Step 3. This step calls for more extensive explanations. First note that the norm form on $B_{p,\infty}$ is \mathbf{Q} -equivalent to $\langle 1, q, p, qp \rangle$, so by the cancellation theorem, $G \simeq_{\mathbf{Q}} \langle q, p, qp \rangle$. Let $Q = G \oplus \langle -q \rangle$. The factorisation of $\text{disc}(G)$ (hence $\text{disc}(Q)$) being known, we can find a solution $X_0^t Q X_0 = 0$ with $X_0 = (x_0, \ell_0)$, where $x_0 \in \mathbf{Z}^3$ is primitive and $\ell_0 \in \mathbf{Z}_{>0}$ using [Sim06]. Yet, ℓ_0 is not necessarily prime. From [Coh08, Proposition 6.3.2], the general solution X is given by

$$X = d((R^t Q R)X_0 - 2(R^t Q X_0)R),$$

for arbitrary $R \in \mathbf{Q}^4$ and $d \in \mathbf{Q}^*$. Fix $d = 1$. Write $R = (r_x, r_\ell)$ with $r_x \in \mathbf{Z}^3$ and $r_\ell \in \mathbf{Z}$. The last coordinate of X is given by the integral quadratic form

$$r_x^t G r_x \ell_0 - 2r_x^t G x_0 r_\ell + q \ell_0 r_\ell^2 = \frac{(r_x \ell_0 - x_0 r_\ell)^t G (r_x \ell_0 - x_0 r_\ell)}{\ell_0}.$$

It is of rank 3, so let $M \in M_{3 \times 3}(\mathbf{Z})$ be a matrix whose columns generate $\Lambda = \ell_0 \mathbf{Z}^3 + x_0 \mathbf{Z}$, and

$$g(z) = \frac{z^t (M^t G M) z}{\ell_0}.$$

It is positive definite, since G is and $\ell_0 > 0$. Let us show that g is (almost) primitive. If s is a prime that does not divide ℓ_0 , both M and ℓ_0 are invertible modulo s , so g is primitive at s because G is. Now suppose $s \mid \ell_0$. Then, writing $Mz = r_x \ell_0 - x_0 r_\ell$, we have

$$g(z) \equiv -2r_x^t G x_0 r_\ell \pmod{s}.$$

Therefore, if $s \neq 2$ and $Gx_0 \not\equiv 0 \pmod{s}$, then g is primitive at s . If $Gx_0 \equiv 0 \pmod{s}$, since x_0 is primitive, s must divide $\text{disc}(G)$, so s is 2 or p . This proves that the only primes where g might not be primitive are 2 and p . We can then write $g = g'/a$ where g' is primitive and a may only be divisible by the primes 2 and p . Applying Proposition 3.5, we can find in polynomial time a z such that $\ell = g'(z)$ is prime, hence a solution of the form $x^t G x = q(a\ell)^2$. In this solution, we can assume that x is primitive: if c divides the content of x , then c^2 divides $q(a\ell)^2$, so c divides $a\ell$.

Step 4. We are looking for B' such that $(B')^tGB' = \langle q \rangle \oplus H$. Let $(x \mid \Gamma)$ be a unimodular integral matrix with first column equal to x (it can be found because x is primitive). Let

$$P = I_3 - \left(\begin{array}{c|c|c} e_1^t Gx & e_2^t Gx & e_3^t Gx \\ \hline x^t Gx & x^t Gx & x^t Gx \end{array} \right)$$

be the 3×3 matrix projecting orthogonally along x . With $B' = (x/(a\ell) \mid (x^t Gx)P\Gamma)$, we obtain $(B')^tGB'$ of the desired form.

Step 5. It can be solved efficiently with [Sim05] since the factorisation of $\text{disc}(H)$ is known.

Steps 6 to 9. In $\text{End}(E) \otimes \mathbf{Q}$, we have $\text{Nrd}(\iota) = q$ and $\text{tr}(\iota) = 0$ so $\iota^2 = -q$. Similarly, $\pi^2 = -p$. Therefore Φ is indeed an isomorphism.

Step 10. All we need to do is express each β_i in the basis $1, \iota, \pi, \kappa$ (allowing to evaluate $\Phi(\beta_i)$). We already know how to express $1, \iota, \pi$ in terms of $(\beta_i)_{i=1}^4$; if we can also express κ , then we obtain a change of basis between $1, \iota, \pi, \kappa$ and $(\beta_i)_{i=1}^4$ and we are done. Without loss of generality, β_4 is not in $\text{span}(1, \iota, \pi)$. Let

$$\gamma = \beta_4 - \langle \beta_4, 1 \rangle - \langle \beta_4, \iota \rangle \iota - \langle \beta_4, \pi \rangle \pi.$$

Then, γ is orthogonal to $\text{span}(1, \iota, \pi)$, so it belongs to $\text{span}(\kappa)$. Renormalising, we obtain κ as a combination of $(\beta_i)_{i=1}^4$. \square

9. ACKNOWLEDGEMENTS

The author wishes to thank Corentin Perret-Gentil and Léo Ducas for their help and feedback on several aspects of this work. This work was supported by the Agence Nationale de la Recherche under grants ANR MELODIA (ANR-20-CE40-0013) and ANR CIAO (ANR-19-CE48-0008).

REFERENCES

- [ABL20] Edgar Assing, Valentin Blomer, and Junxian Li. Uniform Titchmarsh divisor problems. Preprint arXiv:2005.13915, 2020. <https://arxiv.org/abs/2005.13915>.
- [BKV19] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class group computations. In Steven D. Galbraith and Shihō Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security*, volume 11921 of *Lecture Notes in Computer Science*, pages 227–247. Springer, 2019.
- [BKW20] Dan Boneh, Dmitry Kogan, and Katharine Woo. Oblivious pseudorandom functions from isogenies. In Shihō Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security*, volume 12492 of *Lecture Notes in Computer Science*, pages 520–550. Springer, 2020.
- [CLG09] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, Jan 2009.
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: an efficient post-quantum commutative group action. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018.
- [Coh08] Henri Cohen. *Number theory: Volume I: Tools and diophantine equations*, volume 239. Springer Science & Business Media, 2008.
- [Coh13] Henri Cohen. *A course in computational algebraic number theory*, volume 138. Springer Science & Business Media, 2013.
- [Cos20] Craig Costello. B-SIDH: supersingular isogeny diffie-hellman using twisted torsion. In Shihō Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security*, volume 12492 of *Lecture Notes in Computer Science*, pages 440–463. Springer, 2020.
- [Cou06] Jean Marc Couveignes. Hard homogeneous spaces. IACR Cryptology ePrint Archive, Report 2006/291, 2006. <https://eprint.iacr.org/2006/291>.
- [Cox11] David A Cox. *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication*, volume 34. John Wiley & Sons, 2011.

- [CPV20] Wouter Castryck, Lorenz Panny, and Frederik Vercauteren. Rational isogenies from irrational endomorphisms. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 12106 of *Lecture Notes in Computer Science*, pages 523–548. Springer, 2020.
- [Deu41] Max Deuring. Die typen der multiplikatorenringe elliptischer funktionenkörper. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 14(1):197–272, 1941.
- [DG19] Luca De Feo and Steven D. Galbraith. Seasign: Compact isogeny signatures from class group actions. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 11478 of *Lecture Notes in Computer Science*, pages 759–789. Springer, 2019.
- [DKL⁺20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. Sqisign: Compact post-quantum signatures from quaternions and isogenies. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security*, volume 12491 of *Lecture Notes in Computer Science*, pages 64–93. Springer, 2020.
- [DKPS19] Cyprien Delpèch de Saint Guilhem, Péter Kutas, Christophe Petit, and Javier Silva. Séta: Supersingular encryption from torsion attacks. IACR Cryptology ePrint Archive, Report 2019/1291, 2019. <https://eprint.iacr.org/2019/1291>.
- [DMPS19] Luca De Feo, Simon Masson, Christophe Petit, and Antonio Sanso. Verifiable delay functions from supersingular isogenies and pairings. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security*, volume 11921 of *Lecture Notes in Computer Science*, pages 248–277. Springer, 2019.
- [EHL⁺18] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018*, pages 329–368, Cham, 2018. Springer International Publishing.
- [EHM17] Kirsten Eisenträger, Sean Hallgren, and Travis Morrison. On the hardness of computing endomorphism rings of supersingular elliptic curves. IACR Cryptology ePrint Archive, Report 2017/986, 2017. <https://eprint.iacr.org/2017/986>.
- [FKM21] Tako Boris Fouotsa, Péter Kutas, and Simon-Philipp Merz. On the isogeny problem with torsion point information. IACR Cryptology ePrint Archive, Report 2021/153, 2021. <https://eprint.iacr.org/2021/153>.
- [GPS20] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. *Journal of Cryptology*, 33(1):130–175, 2020.
- [GPST16] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security*, volume 10031 of *Lecture Notes in Computer Science*, pages 63–91, 2016.
- [HL23] Godfrey H. Hardy and John E. Littlewood. Some problems of ‘partitio numerorum’; III: On the expression of a number as a sum of primes. *Acta Mathematica*, 44:1–70, 1923.
- [Hoo57] Christopher Hooley. On the representation of a number as the sum of two squares and a prime. *Acta Mathematica*, 97:189–210, 1957.
- [IK04] H. Iwaniec and E. Kowalski. *Analytic Number Theory*. Number v. 53 in American Mathematical Society Colloquium Publications. American Mathematical Society, 2004.
- [JAC⁺17] David Jao, Reza Azarderakhsh, Matt Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalili, Brian Koziel, Brian Lamacchia, Patrick Longa, et al. Sike: Supersingular isogeny key encapsulation. 2017.
- [JD11] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *International Workshop on Post-Quantum Cryptography - PQCrypto 2011*, pages 19–34, 2011.
- [JMV05] D. Jao, S. D. Miller, and R. Venkatesan. Do all elliptic curves of the same order have the same difficulty of discrete log? In Bimal K. Roy, editor, *Advances in Cryptology - ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 21–40. Springer, 2005.
- [KLPT14] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ -isogeny path problem. *LMS Journal of Computation and Mathematics*, 17(A):418–432, 2014.
- [Koh96] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.
- [Kri20] Alexander Krigsman. Efficient Deuring correspondence computation. Master’s thesis, Universiteit Leiden, The Netherlands, 2020.
- [KV10] Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM Journal on Computing*, 39(5):1714–1747, 2010.

- [LLL82] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
- [LO77] J.C. Lagarias and A.M. Odlyzko. Effective versions of the Chebotarev density theorem. In *Algebraic number fields: L-functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, pages 409–464. Academic Press, London, 1977.
- [Mes86] Jean-Francois Mestre. La méthode des graphes. exemples et applications. In *Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata)*, pages 217–242, 1986.
- [Pal33] Gordon Pall. The structure of the number of representations function in a positive binary quadratic form. *Mathematische Zeitschrift*, 36(1):321–343, 1933.
- [Piz80] Arnold Pizer. An algorithm for computing modular forms on $\gamma_0(n)$. *Journal of algebra*, 64(2):340–390, 1980.
- [Piz90] Arnold K Pizer. Ramanujan graphs and Hecke operators. *Bulletin of the American Mathematical Society*, 23(1):127–137, 1990.
- [PL17] Christophe Petit and Kristin Lauter. Hard and easy problems for supersingular isogeny graphs. IACR Cryptology ePrint Archive, Report 2017/962, 2017. <https://eprint.iacr.org/2017/962>.
- [Rón92] Lajos Rónyai. Algorithmic properties of maximal orders in simple algebras over \mathbf{Q} . *Computational Complexity*, 2(3):225–243, 1992.
- [Sil86] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, 1986.
- [Sim05] Denis Simon. Solving quadratic equations using reduced unimodular quadratic forms. *Mathematics of Computation*, 74:1531–1543, 2005.
- [Sim06] Denis Simon. Quadratic equations in dimensions 4, 5 and more. Preprint, 2006. See [Wat13] for a published review.
- [Vél71] J. Vélu. Isogénies entre courbes elliptiques. *Comptes rendus de l'Académie des Sciences, Séries A-B*, 273:A238–A241, 1971.
- [Voi21] John Voight. *Quaternion Algebras*. Springer International Publishing, 2021. Graduate Texts in Mathematics, No. 288.
- [Wat13] Mark Watkins. Some comments about indefinite LLL. *Diophantine Methods, Lattices, and Arithmetic Theory of Quadratic Forms*, 587(233):32, 2013.

UNIV. BORDEAUX, CNRS, BORDEAUX INP, IMB, UMR 5251, F-33400, TALENCE, FRANCE, INRIA, IMB, UMR 5251, F-33400, TALENCE, FRANCE

E-mail address: benjamin.wesolowski@math.u-bordeaux.fr