On the Ring-LWE and Polynomial-LWE Problems

Miruna Rosca[1,2], Damien Stehlé[1], and Alexandre Wallet[1(✉)]

[1] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL),
Lyon, France
damien.stehle@gmail.com, wallet.alexandre@gmail.com
[2] Bitdefender, Bucharest, Romania

Abstract. The Ring Learning With Errors problem (RLWE) comes in various forms. Vanilla RLWE is the decision dual-RLWE variant, consisting in distinguishing from uniform a distribution depending on a secret belonging to the dual $\mathcal{O}_K^\vee$ of the ring of integers $\mathcal{O}_K$ of a specified number field K. In primal-RLWE, the secret instead belongs to $\mathcal{O}_K$. Both decision dual-RLWE and primal-RLWE enjoy search counterparts. Also widely used is (search/decision) Polynomial Learning With Errors (PLWE), which is not defined using a ring of integers $\mathcal{O}_K$ of a number field K but a polynomial ring $\mathbb{Z}[x]/f$ for a monic irreducible $f \in \mathbb{Z}[x]$. We show that there exist reductions between all of these six problems that incur limited parameter losses. More precisely: we prove that the (decision/search) dual to primal reduction from Lyubashevsky *et al.* [EUROCRYPT 2010] and Peikert [SCN 2016] can be implemented with a small error rate growth for all rings (the resulting reduction is non-uniform polynomial time); we extend it to polynomial-time reductions between (decision/search) primal RLWE and PLWE that work for a family of polynomials f that is exponentially large as a function of $\deg f$ (the resulting reduction is also non-uniform polynomial time); and we exploit the recent technique from Peikert *et al.* [STOC 2017] to obtain a search to decision reduction for RLWE for arbitrary number fields. The reductions incur error rate increases that depend on intrinsic quantities related to K and f.

1 Introduction

DIFFERENT SHADES OF RLWE. Ring Learning With Errors (RLWE) was introduced by Lyubashevsky *et al.* in [LPR10], as a means of speeding up cryptographic constructions based on the Learning With Errors problem (LWE) [Reg05]. Let K be a number field, $\mathcal{O}_K$ its ring of integers and $q \geq 2$ a rational integer. The search variant of RLWE with parameters K and q consists in recovering a secret $s \in \mathcal{O}_K^\vee / q\mathcal{O}_K^\vee$ with $\mathcal{O}_K^\vee$ denoting the dual of $\mathcal{O}_K$, from arbitrarily many samples $(a_i, a_i \cdot s + e_i)$. Here each a_i is uniformly sampled in $\mathcal{O}_K / q\mathcal{O}_K$ and each e_i is a small random element of $K_\mathbb{R} := K \otimes_\mathbb{Q} \mathbb{R}$. The noise term e_i is sampled such that its Minkowski embedding vector follows a Gaussian distribution with a small

covariance matrix (relative to $q\mathcal{O}_K^\vee$). The decision variant consists in distinguishing arbitrarily many such pairs for a common s chosen uniformly in $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$, from uniform samples in $\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K^\vee$. More formal definitions are provided in Sect. 2, but these suffice for describing our contributions.

Lyubashevsky *et al.* backed in [LPR10] the conjectured hardness of RLWE with a quantum polynomial-time reduction from the (worst-case) Approximate Shortest Vector Problem (ApproxSVP) restricted to the class of Euclidean lattices corresponding to ideals of $\mathcal{O}_K$, with geometry inherited from the Minkowski embeddings. They showed its usefulness by describing a public-key encryption with quasi-optimal efficiency: the bit-sizes of the keys and the run-times of all involved algorithms are quasi-linear in the security parameter. A central technical contribution was a reduction from search RLWE to decision RLWE, when K is cyclotomic, and decision RLWE for cyclotomic fields is now pervasive in lattice-based cryptography, including in practice [ADPS16,BDK+18,DLL+18]. The search-to-decision reduction from [LPR10] was later extended to the case of general Galois rings in [EHL14,CLS17].

Prior to RLWE, Stehlé *et al.* [SSTX09] introduced what is now referred to as Polynomial Ring Learning With Errors (PLWE), for cyclotomic polynomials of degree a power of 2. PLWE is parametrized by a monic irreducible $f \in \mathbb{Z}[x]$ and an integer $q \geq 2$, and consists in recovering a secret $s \in \mathbb{Z}_q[x]/f$ from arbitrarily many samples $(a_i, a_i \cdot s + e_i)$ where each a_i is uniformly sampled in $\mathbb{Z}_q[x]/f$ and each e_i is a small random element of $\mathbb{R}[x]/f$. The decision variant consists in distinguishing arbitrarily many such samples for a common s sampled uniformly in $\mathbb{Z}_q[x]/f$, from uniform samples. Here the noise term e_i is sampled such that its coefficient vector follows a Gaussian distribution with a small covariance matrix. Stehlé *et al.* gave a reduction from the restriction of ApproxSVP to the class of lattices corresponding to ideals of $\mathbb{Z}[x]/f$, to search PLWE, for f a power-of-2 cyclotomic polynomial.

Finally, a variant of RLWE with $s \in \mathcal{O}_K/q\mathcal{O}_K$ rather than $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ was also considered (see, e.g., [DD12] among others), to avoid the complication of having to deal with the dual $\mathcal{O}_K^\vee$ of $\mathcal{O}_K$. In the rest of this paper, we will refer to the latter as primal-RLWE and to standard RLWE as dual-RLWE.

THE CASE OF CYCLOTOMICS. Even though [LPR10] defined RLWE for arbitrary number fields, the problem was mostly studied in the literature for K cyclotomic. This specialization had three justifications:

- it leads to very efficient cryptographic primitives, in particular if q totally splits over K;
- the hardness result from [LPR10] holds for cyclotomics;
- no particular weakness was known for these fields.

Among cyclotomics, those of order a power of 2 are a popular choice. In the case of a field K defined by the cyclotomic polynomial f, we have that $\mathcal{O}_K = \mathbb{Z}[\alpha]$ for α a root of f. Further, in the case of power-of-2 cyclotomics, mapping the coefficient vector of a polynomial in $\mathbb{Z}[x]/f$ to its Minkowski embedding is a scaled isometry. This makes primal-RLWE and PLWE collapse into a single

problem. Still in the case of power-of-2 cyclotomics, the dual $\mathcal{O}_K^\vee$ is a scaling of $\mathcal{O}_K$, implying that dual and primal-RLWE are equivalent. Apart from the monogenicity property, these facts do not hold for all cyclotomics. Nevertheless, Ducas and Durmus [DD12] showed it is still possible to reduce dual-RLWE to primal-RLWE.

LOOKING AT OTHER FIELDS. The RLWE hardness proof holds with respect to a fixed field: the reduction in [LPR10] maps ApproxSVP for lattices corresponding to $\mathcal{O}_K$-ideals with small approximation factors, to decision/search dual-RLWE on K. Apart from the very specific case of field extensions [GHPS12], hardness on K seems unrelated to hardness on another field K'. One may then wonder if RLWE is easier for some fields. The attacks presented in [EHL14, ELOS15, CLS17, CLS16] were used to identify weak generating polynomials f of a number field K, but they only work for error distributions with small width relative to the geometry of the corresponding ring [CIV16b, CIV16a, Pei16]. At this occasion, the relationships between the RLWE and PLWE variants were more closely investigated.

Building upon [CGS14, CDPR16], Cramer *et al.* [CDW17] gave a quantum polynomial-time ApproxSVP algorithm for ideals of $\mathcal{O}_K$ when K is a cyclotomic field of prime-power conductor, when the ApproxSVP approximation factor is $2^{\widetilde{O}(\sqrt{\deg K})}$. For general lattices, the best known algorithm [SE94] runs in time $2^{\widetilde{O}(\sqrt{n})}$ for such an approximation factor, where n is the lattice dimension (here $n = \deg K$). We note that the result from [CGS14, CDPR16] was partly extended in [BBdV+17] to principal ideals generated by a short element in a completely different family of fields. These results show that all fields are not equal in terms of ApproxSVP hardness (unless they turn out to be all weak!). So far, there is no such result for RLWE.

On the constructive front, Bernstein *et al.* [BCLvV16] showed that some non-cyclotomic polynomials f also enjoy practical arithmetic over $\mathbb{Z}_q[x]/f$ and lead to efficient cryptographic design (though the concrete scheme relies on the presumed hardness of another problem than RLWE).

HEDGING AGAINST THE WEAK FIELD RISK. Two recent works propose complementary approaches to hedge against the risk of a weakness of RLWE for specific fields. First, in [PRS17], Peikert *et al.* give a new (quantum) reduction from ApproxSVP for $\mathcal{O}_K$-ideals to decision dual-RLWE for the corresponding field K. All fields support a (quantum) reduction from ApproxSVP, and hence, from this respect, one is not restricted to cyclotomics. Second, following an analogous result by Lyubashevsky for the Small Integer Solution problem [Lyu16], Roşca *et al.* [RSSS17] introduced the Middle-Product LWE problem and showed that it is at least as hard as PLWE for any f in an exponentially large family of f's (as a function of their degree). Neither result is fully satisfactory. In the first case, it could be that ApproxSVP is easy for lattices corresponding to ideals of $\mathcal{O}_K$ for any K: this would make the result vacuous. In the second case, the result of [RSSS17] focuses on PLWE rather than the more studied RLWE problem.

OUR RESULTS. The focus on the RLWE hardness for non-cyclotomic fields makes the discrepancies between the RLWE and PLWE variants more critical. In this

article, we show that the six problems considered above — dual-RLWE, primal-RLWE and PLWE, all in both decision and search forms — reduce to one another in polynomial time with limited error rate increases, for huge classes of rings. More precisely, these reductions are obtained with the following three results.

- We show that for every field K, it is possible to implement the reduction from decision (resp. search) dual-RLWE to decision (resp. search) primal-RLWE from [LPR10, Le. 2.15] and [Pei16, Se. 2.3.2], with a limited error growth. Note that there exists a trivial converse reduction from primal-RLWE to dual-RLWE.
- We show that the reduction mentioned above can be extended to a reduction from decision (resp. search) primal-RLWE in K to decision (resp. search) PLWE for f, where K is the field generated by the polynomial f. The analysis is significantly more involved. It requires the introduction of the so-called conductor ideal, to handle the transformation from the ideal $\mathcal{O}_K$ to the order $\mathbb{Z}[x]/f$, and upper bounds on the condition number of the map that sends the coefficient embeddings to the Minkowski embeddings, to show that the noise increases are limited. Our conditioning upper bound is polynomial in n only for limited (but still huge) classes of polynomials that include those of the form $x^n + x \cdot P(x) - a$, with $\deg P < n/2$ and a prime that is $\geq 25 \cdot \|P\|_1^2$ and $\leq \mathrm{poly}(n)$. A trivial converse reduction goes through for the same f's.
- We exploit the recent technique from [PRS17] to obtain a search to decision reduction for dual-RLWE.

Concretely, the error rate increases are polynomial in $n = \deg K$, the root discriminant $|\Delta_K|^{1/n}$ and, for the reduction to PLWE, in the root algebraic norm $\mathcal{N}(\mathcal{C}_{\mathbb{Z}[\alpha]})^{1/n}$ of the conductor ideal $\mathcal{C}_{\mathbb{Z}[\alpha]}$ of $\mathbb{Z}[\alpha]$, where α is a root of f defining K. We note that in many cases of interest, all these quantities are polynomially bounded in n. To enjoy these limited error rate growths, the first two reductions require knowledge of specific data related to K, namely, a short element (with respect to the Minkowski embeddings) in the different ideal $(\mathcal{O}_K^\vee)^{-1}$ and a short element in $\mathcal{C}_{\mathbb{Z}[\alpha]}$. In general, these are hard to compute.

TECHNIQUES. The first reduction is derived from [LPR10, Le. 2.15] and [Pei16, Se. 2.3.2]: if it satisfies some arithmetic properties, a multiplication by an element $t \in \mathcal{O}_K$ induces an $\mathcal{O}_K$-module isomorphism from $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ to $\mathcal{O}_K/q\mathcal{O}_K$. For the reduction to be meaningful, we need t to have small Minkowski embeddings. We prove the existence of such a small t satisfying the appropriate arithmetic conditions, by generalizing the inclusion-exclusion technique developed in [SS13] to study the key generation algorithm of the NTRU signature scheme [HHPW10].

The Lyubashevsky *et al.* bijection works with $\mathcal{O}_K^\vee$ and $\mathcal{O}_K$ replaced by arbitrary ideals of K, but this does not provide a bijection from $\mathcal{O}_K/q\mathcal{O}_K$ to $\mathbb{Z}[\alpha]/q\mathbb{Z}[\alpha]$, as $\mathbb{Z}[\alpha]$ may only be an order of $\mathcal{O}_K$ (and not necessarily an ideal). We circumvent this difficulty by using the conductor ideal of $\mathbb{Z}[\alpha]$. Intuitively, the conductor ideal describes the relationship between $\mathcal{O}_K$ and $\mathbb{Z}[\alpha]$. As far as we

are aware, this is the first time the conductor ideal is used in the RLWE context. This bijection and the existence of an appropriate multiplier t as above provide a (non-uniform) reduction from primal-RLWE to a variant of PLWE for which the noise terms have small Minkowski embeddings (instead of small polynomial coefficients).

We show that for many number fields, the linear map between polynomial coefficients and Minkowski embeddings has a condition number that is polynomially bounded in n, i.e., the map has bounded distortion and behaves not too noticeably differently from a scaling. This implies that the latter reduction is also a reduction from primal-RLWE to standard PLWE for these rings. We were able to show condition number bounds that are polynomial in n only for restricted families of polynomials f, yet exponentially large as n increases. These include in particular those of the form mentioned above. Note that the primality condition on the constant coefficient is used only to ensure that f is irreducible and hence defines a number field. For these f's, we use Rouché's theorem to prove that the roots are close to the scaled n-th roots of unity $(a^{1/n} \cdot \alpha_n^k)_{0 \leq k < n}$, and then that f "behaves" as $x^n - a$ in terms of geometric distortion.

Our search-to-decision reduction for dual-RLWE relies on techniques developed in [PRS17]. In that article, Peikert *et al.* consider the following 'oracle hidden center' problem (OHCP). In this problem, we are given access to an oracle $\mathcal{O}$ taking as inputs a vector $\boldsymbol{z} \in \mathbb{R}^k$ and a scalar $t \in \mathbb{R}^{\geq 0}$, and outputting a bit. The probability that the oracle outputs 1 (over its internal randomness) is assumed to depend only on $\exp(t) \cdot \|\boldsymbol{z} - \boldsymbol{x}\|$, for some vector $\boldsymbol{x}$. The goal is to recover $\mathcal{O}$'s center $\boldsymbol{x}$. On the one hand, Peikert *et al.* give a polynomial-time algorithm for this problem, assuming the oracle is 'well-behaved' ([PRS17, Prop. 4.4]). On the other hand, they show how to map a Bounded Distance Decoding (BDD) instance to such an OHCP instance if they have access to Gaussian samples in the dual of the BDD lattice, where the engine of the oracle is the decision dual-RLWE oracle [PRS17, Se. 6.1]. We construct the OHCP instance from the decision RLWE oracle in a different manner. We use our input search dual-RLWE samples and take small Gaussian combinations of them. By re-randomizing the secret and adding some noise, we can obtain arbitrarily many dual-RLWE samples. Subtracting from the input samples well-chosen z_i's in $K_\mathbb{R}$ and setting the standard deviation of the Gaussian combination appropriately leads to a valid OHCP instance. The main technical hurdle is to show that a Gaussian combination of elements of $\mathcal{O}_K^\vee / q\mathcal{O}_K^\vee$ is close to uniform. For this, we generalize a ring Leftover Hash Lemma proved for specific pairs $(\mathcal{O}_K, q)$ in [SS11].

RELATED WORKS. The reductions studied in this work can be combined with those from ApproxSVP for $\mathcal{O}_K$-ideals to dual-RLWE [LPR10,PRS17]. Recently, Albrecht and Deo [AD17] built upon [BLP+13] to obtain a reduction from Module-LWE to RLWE. This can be both combined with our reductions and the quantum reductions from ApproxSVP for $\mathcal{O}_K$-modules to Module-LWE[1] [LS15,PRS17]. Downstream, the reductions can be combined with the

---

[1] The reduction from [LS15] is limited to cyclotomic fields, but [PRS17] readily extends to module lattices.

reduction from PLWE to Middle-Product LWE from [RSSS17]. The latter was showed to involve an error rate growth that is linearly bounded by the so-called the expansion factor of f: it turns out that those f's for which we could bound the condition number of the Minkowski map by a polynomial function of $\deg f$ also have polynomially bounded expansion factor. These reductions and those considered in the present work are pictorially described in Fig. 1.
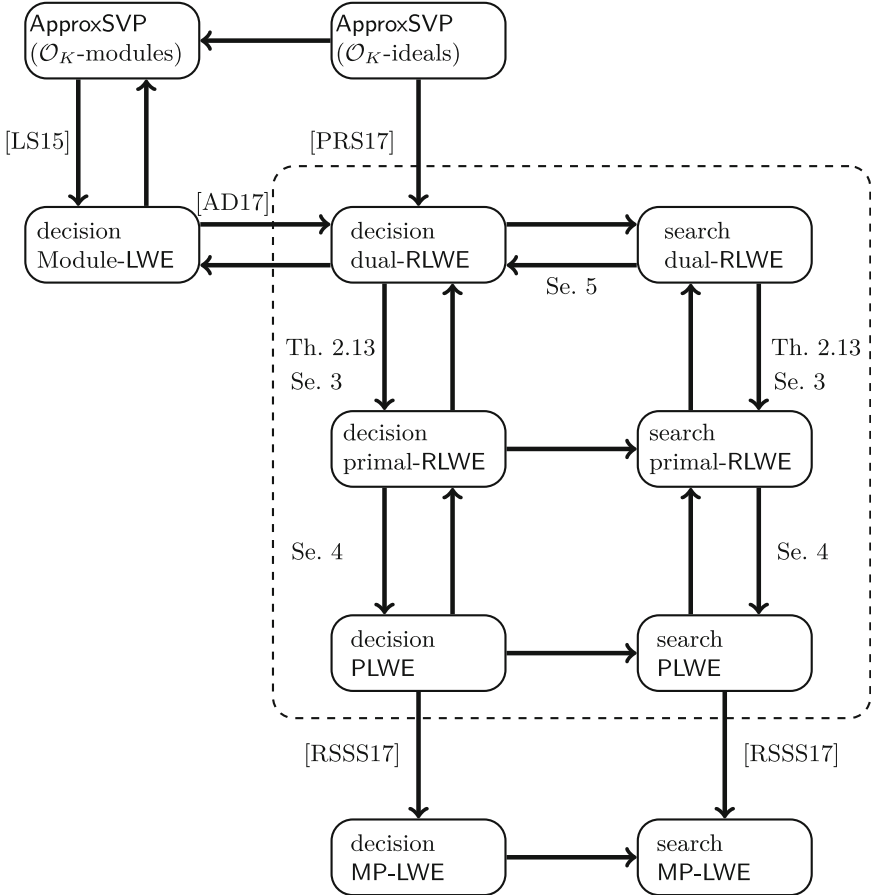
Fig. 1. Relationships between variants of RLWE and PLWE. The dotted box contains the problems studied in this work. Each arrow may hide a noise rate degradation (and module rank - modulus magnitude transfer in the case of [AD17]). The top to bottom arrows in the dotted box correspond to non-uniform reductions. The reductions involving PLWE are analyzed for limited family of defining polynomials. The arrows without references correspond to trivial reductions.

The ideal-changing scaling element t and the distortion of the Minkowski map were closely studied in [CIV16b, CIV16a, Pei16] for a few precise polynomials and fields. We use the same objects, but provide bounds that work for all (or many) fields.

IMPACT. As it is standard for the hardness foundations of lattice-based cryptography, our reductions *should not* be considered for setting practical parameters. They should rather be viewed as a strong evidence that the six problems under scope are essentially equivalent and do not suffer from a design flaw (unless they all do). We hope they will prove useful towards understanding the plausibility of weak fields for RLWE.

Our first result shows that there exists a way of reducing dual-RLWE to primal-RLWE while controlling the noise growth. Even though the reduction is non-uniform, it gives evidence that these problems are qualitatively equivalent. Our second result shows that RLWE and PLWE are essentially equivalent for a large class of polynomials/fields. In particular, the transformation map between the Minkowski embeddings and the coefficient embeddings has a bounded distortion. Finally, our search to decision fills an important gap. On the one hand, it precludes the possibility that search RLWE could be harder than decision RLWE. On the other hand, it gives further evidence of the decision RLWE hardness. In [PRS17], the authors give a reduction from ApproxSVP for $\mathcal{O}_K$-ideals to decision RLWE. But in the current state of affairs, ApproxSVP for this special class of lattices seems easier than RLWE, at least for some parameters. Indeed, Cramer *et al.* [CDW17] gave quantum algorithms that outperform generic lattice algorithms for some range of approximation factors in the context of ideal lattices. On the opposite, RLWE is qualitatively equivalent to ApproxSVP for $\mathcal{O}_K$-modules [LS15, AD17].

As the studied problems reduce to one another, one may then wonder which one to use for cryptographic design. Using dual-RLWE requires knowledge of $\mathcal{O}_K$, which is notoriously hard to compute for an arbitrary field K. This may look as an incentive to use the corresponding PLWE problem instead, as it does not require the knowledge of $\mathcal{O}_K$. Yet, for it to be useful in cryptographic design, one must be able to decode the noise from its representative modulo a scaled version of the lattice corresponding to $\mathbb{Z}[\alpha]$. This seems to require the knowledge of a good basis of that lattice, which may not be easy to obtain either, depending on the considered polynomial f.

NOTATIONS. If D is a distribution, we write $x \hookleftarrow D$ to say that we sample x from D. If D_1, D_2 are continuous distributions over the same measurable set Ω, we let $\Delta(D_1, D_2) = \int_\Omega |D_1(x) - D_2(x)| \mathrm{d}x$ denote their statistical distance. Similarly, we let $R(D_1 \| D_2) = \int_\Omega D_1(x)^2 / D_2(x) \mathrm{d}x$ denote their Rényi divergence. If E is a set of finite measure, we let $U(E)$ denote the uniform distribution over E. For a matrix $V = (v_{ij})$, we let $\|V\| = \sqrt{\sum_{1 \le i,j \le n} |v_{ij}|^2}$ denote its Frobenius norm.

This is the proceedings' version. The full version contains additional appendices and it is available on the IACR eprint archive.

2   Preliminaries

In this section, we give some background on algebraic number theory used in lattice-based cryptography, recall properties of Euclidean lattices, and state the

precise definitions of the RLWE variants we will consider. More details on standard tools of algebraic number theory can be found in the full version. Useful references include [Ste17, Cona].

2.1   Some Algebraic Number Theory

RINGS AND IDEALS IN NUMBER FIELDS. In this article, we call any subring of K a number ring. For a number ring R, an (integral) R-ideal is an additive subgroup $I \subseteq R$ which is closed by multiplication in R, i.e., such that $IR = I$. A more compact definition is to say that I is an R-module. If $a_1, \ldots, a_k$ are elements in R, we let $\langle a_1, \ldots, a_k \rangle = a_1 R + \ldots + a_k R$ and call it the ideal generated by the a_i's. The product of two ideals I, J is the ideal generated by all elements xy with $x \in I$ and $y \in J$. The sum, product and intersection of two R-ideals are again R-ideals.

Two integral R-ideals I, J are said to be coprime if $I + J = R$, and, in this case, we have $I \cap J = IJ$. Any non-zero ideal in a number ring has finite index, i.e., the quotient ring R/I is always finite when I is a non-zero R-ideal. An R-ideal $\mathfrak{p}$ is said to be prime if whenever $\mathfrak{p} = IJ$ for some R-ideals I, J, then either $I = \mathfrak{p}$ or $J = \mathfrak{p}$. In a number ring, any prime ideal $\mathfrak{p}$ is maximal [Ste17, p. 19], i.e., R is the only R-ideal containing it. It also means that the quotient ring $R/\mathfrak{p}$ is a finite field. It is well-known that any $\mathcal{O}_K$-ideal admits a unique factorization into prime $\mathcal{O}_K$-ideals, i.e., it can be written $I = \mathfrak{p}_1^{e_1} \ldots \mathfrak{p}_k^{e_k}$ with all $\mathfrak{p}_i$'s distinct prime ideals. It fails to hold in general number rings and orders, but we describe later in Lemma 2.1 how the result can be extended in certain cases.

A fractional R-ideal I is an R-module such that $xI \subseteq R$ for some $x \in K^\times$. An integral ideal is a fractional ideal, and so are the sum, the product and the intersection of two fractional ideals. A fractional R-ideal I is said to be invertible if there exists a fractional R-ideal J such that $IJ = R$. In this case, the (unique) inverse is the integral ideal $I^{-1} = \{x \in K : xI \subseteq R\}$. Any $\mathcal{O}_K$-ideal is invertible, but it is again false for a general number ring.

The algebraic norm of a non-zero integral R-ideal I is defined as $\mathcal{N}_R(I) = |R/I|$, and we will omit the subscript when $R = \mathcal{O}_K$. It satisfies $\mathcal{N}_R(IJ) = \mathcal{N}_R(I)\mathcal{N}_R(J)$ for every R-ideals I, J.

The dual of a fractional R-ideal I is $I^\vee = \{\alpha \in K : \text{Tr}(\alpha I) \subseteq \mathbb{Z}\}$, which is also a fractional R-ideal. We always have $II^\vee = R^\vee$, so that $I^\vee = I^{-1}R^\vee$ when I is invertible. We also have $I^{\vee\vee} = I$ for any R-ideal I.

A particularly interesting dual is $\mathcal{O}_K^\vee$, whose inverse $(\mathcal{O}_K^\vee)^{-1}$ is called the different ideal. The different ideal is an integral ideal, whose norm $\Delta_K = \mathcal{N}((\mathcal{O}_K^\vee)^{-1})$ is called the discriminant of the number field. We note that, for every f defining K, the field discriminant Δ_K is a factor of the discriminant of f. The latter is denoted Δ_f and is defined as $\Delta_f = \prod_{i \neq j}(\alpha_i - \alpha_j)$, where $\alpha_1, \ldots, \alpha_n$ are the roots of f. This provides an upper bound on Δ_K in terms of the defining polynomial f.

ORDERS IN NUMBER FIELDS. An order $\mathcal{O}$ in K is a number ring which is a finite index subring of $\mathcal{O}_K$. In particular, the ring of integers $\mathcal{O}_K$ is the maximal order

in K. Number rings such as $\mathbb{Z}[\alpha]$, with α a root of a defining polynomial f, are of particular interest. The conductor of an order $\mathcal{O}$ is defined as the set $\mathcal{C}_\mathcal{O} = \{x \in K : x\mathcal{O}_K \subseteq \mathcal{O}\}$. It is contained in $\mathcal{O}$, and it is both an $\mathcal{O}$-ideal and an $\mathcal{O}_K$-ideal: it is in fact the largest ideal with this property. It is never empty, as it contains the index $[\mathcal{O}_K : \mathcal{O}]$.

If it is coprime with the conductor, an ideal in $\mathcal{O}_K$ can be naturally considered as an ideal in $\mathcal{O}$, and reciprocally. This is made precise in the following lemma.

Lemma 2.1 ([Cona, **Th. 3.8**]). *Let $\mathcal{O}$ be an order in K.*

1. *Let I be an $\mathcal{O}_K$-ideal coprime to $\mathcal{C}_\mathcal{O}$. Then $I \cap \mathcal{O}$ is an $\mathcal{O}$-ideal coprime to $\mathcal{C}_\mathcal{O}$ and the natural map $\mathcal{O}/I \cap \mathcal{O} \longrightarrow \mathcal{O}_K/I$ is a ring isomorphism.*
2. *Let J be an $\mathcal{O}$-ideal coprime to $\mathcal{C}_\mathcal{O}$. Then $J\mathcal{O}_K$ is an $\mathcal{O}_K$-ideal coprime to $\mathcal{C}_\mathcal{O}$ and the natural map $\mathcal{O}/J \longrightarrow \mathcal{O}_K/J\mathcal{O}_K$ is a ring isomorphism.*
3. *The set of $\mathcal{O}_K$-ideals coprime to $\mathcal{C}_\mathcal{O}$ and the set of $\mathcal{O}$-ideals coprime to $\mathcal{C}_\mathcal{O}$ are in multiplicative bijection by $I \longmapsto I \cap \mathcal{O}$ and $J \longmapsto J\mathcal{O}_K$.*

The above description does not tell how to "invert" the isomorphisms. This can be done by a combination of the following lemmas and passing through the conductor, as we will show in the next section.

Lemma 2.2. *Let $\mathcal{O}$ be an order in K and I an $\mathcal{O}_K$-ideal coprime to the conductor $\mathcal{C}_\mathcal{O}$. Then the inclusions $\mathcal{C}_\mathcal{O} \subseteq \mathcal{O}$ and $\mathcal{C}_\mathcal{O} \subseteq \mathcal{O}_K$ induce isomorphisms $\mathcal{C}_\mathcal{O}/I \cap \mathcal{C}_\mathcal{O} \simeq \mathcal{O}/I \cap \mathcal{O}$ and $\mathcal{C}_\mathcal{O}/I \cap \mathcal{C}_\mathcal{O} \simeq \mathcal{O}_K/I$.*

Proof. By assumption we have $\mathcal{C}_\mathcal{O} + I = \mathcal{O}_K$, so that the homomorphism $\mathcal{C}_\mathcal{O} \to \mathcal{O}_K/I$ is surjective. By Lemma 2.1, the set $I \cap \mathcal{O}$ is an $\mathcal{O}$-ideal coprime to $\mathcal{C}_\mathcal{O}$ so that $\mathcal{C}_\mathcal{O} + I \cap \mathcal{O} = \mathcal{O}$. This implies that the homomorphism $\mathcal{C}_\mathcal{O} \to \mathcal{O}/I \cap \mathcal{O}$ is surjective too. Both homomorphisms have kernel $I \cap \mathcal{C}_\mathcal{O}$.   □

Lemma 2.3 ([Cona, **Cor. 3.10**]). *Let $\mathcal{O}$ be an order in K and $\beta \in \mathcal{O}$ such that $\beta\mathcal{O}_K$ is coprime to $\mathcal{C}_\mathcal{O}$. Then $\beta\mathcal{O}_K \cap \mathcal{O} = \beta\mathcal{O}$.*

QUOTIENTS OF IDEALS. We will use the following result.

Lemma 2.4 ([LPR10, **Le. 2.14**]). *Let I and J two $\mathcal{O}_K$-ideals. Let $t \in I$ such that the ideals $t \cdot I^{-1}$ and J are coprime and let $\mathcal{M}$ be any fractional $\mathcal{O}_K$-ideal. Then the function $\theta_t : \mathcal{M} \to \mathcal{M}$ defined as $\theta_t(x) = t \cdot x$ induces an $\mathcal{O}_K$-module isomorphism from $\mathcal{M}/J\mathcal{M}$ to $I\mathcal{M}/IJ\mathcal{M}$.*

The authors of [LPR10] also gave an explicit way to obtain a suitable t by solving a set of conditions stemming from the Chinese Remainder Theorem. However, this construction does not give good control on the magnitudes of the Minkowski embeddings of t.

2.2   Lattices

For the remainder of this article, a lattice is defined as a full-rank discrete additive subgroup of an $\mathbb{R}$-vector space V which is a Cartesian power H^m (for $m \geq 1$) of $H := \{\boldsymbol{x} \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} : \forall i \leq s_2 : x_{s_1+s_2+i} = \overline{x_{s_1+i}}\}$. This space H is sometimes called the "canonical" space. A given lattice $\mathcal{L}$ can be thought as the set of $\mathbb{Z}$-linear combinations $(\boldsymbol{b}_i)_i$ of some linearly independent vectors of V. These vectors are said to form a lattice basis, and we define the lattice determinant as $\det \mathcal{L} = (\det(\langle \boldsymbol{b}_i, \boldsymbol{b}_j \rangle)_{i,j})^{1/2}$ (it does not depend on the choice of the basis of $\mathcal{L}$). For $v \in V$, let $\|v\| = (\sum_{i \leq \dim V} |v_i|^2)^{1/2}$ denote the standard Hermitian norm on V and $\|v\|_\infty = \max_{i \leq \dim V} |v_i|$ denote the infinity norm. The minimum $\lambda_1(\mathcal{L})$ is the Hermitian norm of a shortest non-zero element in $\mathcal{L}$. We define $\lambda_1^\infty(\mathcal{L})$ similarly. If $\mathcal{L}$ is a lattice, then we define its dual as $\mathcal{L}^* = \{\boldsymbol{y} \in V : \boldsymbol{y}^T \mathcal{L} \subseteq \mathbb{Z}\}$.

IDEAL LATTICES. While it is possible to associate lattices with fractional ideals of a number ring, we will not need it. Any fractional $\mathcal{O}_K$-ideal I is a free $\mathbb{Z}$-module of rank $n = \deg(K)$, i.e., it can be written as $\mathbb{Z}u_1 + \cdots + \mathbb{Z}u_n$ for some u_i's in K. Its canonical embedding $\sigma(I)$ is a lattice of dimension n in the $\mathbb{R}$-vector space $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$. Such a lattice is called an ideal lattice (for $\mathcal{O}_K$). For the sake of readability, we will abuse notations and often identify I and $\sigma(I)$. It is possible to look at the coefficient embedding of such lattices as well, but we will not need it in this work. The lattice corresponding to $I^\vee$ is $\overline{I^*}$. The discriminant of K satisfies $\Delta_K = (\det \mathcal{O}_K)^2$. In the following lemma, the upper bounds follow from Minkowski's theorem whereas the lower bounds are a consequence of the algebraic structure underlying ideal lattices.

Lemma 2.5 (Adapted from [PR07, Se. 6.1]). *Let K be a number field of degree n. For any fractional $\mathcal{O}_K$-ideal I, we have:*

$$\sqrt{n} \cdot \mathcal{N}(I)^{1/n} \leq \lambda_1(I) \leq \sqrt{n} \cdot (\mathcal{N}(I)\sqrt{\Delta_K})^{1/n},$$
$$\mathcal{N}(I)^{1/n} \quad \leq \lambda_1^\infty(I) \leq \quad (\mathcal{N}(I)\sqrt{\Delta_K})^{1/n}.$$

GAUSSIANS. It is standard practice in the RLWE setting to consider Gaussian distributions with diagonal covariance matrices. In this work, we will be interested in the behavior of samples after linear transformations that are not necessarily diagonal. As the resulting covariance matrix may not be diagonal, we adopt a more general framework. Let $\boldsymbol{\Sigma} \succ 0$, i.e., a symmetric positive definite matrix. We define the Gaussian function on $\mathbb{R}^n$ of covariance matrix $\boldsymbol{\Sigma}$ as $\rho_{\boldsymbol{\Sigma}}(\mathbf{x}) := \exp(-\pi \cdot \mathbf{x}^T \boldsymbol{\Sigma}^{-1} \mathbf{x})$ for every vector $\mathbf{x} \in \mathbb{R}^n$. The Gaussian distribution $D_{\boldsymbol{\Sigma}}$ is the probability distribution whose density is proportional to $\rho_{\boldsymbol{\Sigma}}$. When $\boldsymbol{\Sigma} = \mathrm{diag}(r_i^2)_i$ for some $\boldsymbol{r} \in \mathbb{R}^n$, we write $\rho_{\boldsymbol{r}}$ and $D_{\boldsymbol{r}}$, respectively.

Let $(\boldsymbol{e}_i)_{i \leq n}$ be the canonical basis of $\mathbb{C}^n$. We define $\boldsymbol{h}_i = \boldsymbol{e}_i$ for $i \leq s_1$, and $\boldsymbol{h}_{s_1+i} = (\boldsymbol{e}_{s_1+i} + \boldsymbol{e}_{s_1+s_2+i})/\sqrt{2}$ and $\boldsymbol{h}_{s_1+s_2+i} = (\boldsymbol{e}_{s_1+i} - \boldsymbol{e}_{s_1+s_2+i})/\sqrt{-2}$ for $i \leq s_2$. The $\boldsymbol{h}_i$'s form an orthonormal $\mathbb{R}$-basis of H. We define the Gaussian distribution $D_{\boldsymbol{\Sigma}}^H$ as the distribution obtained by sampling $x \hookleftarrow D_{\boldsymbol{\Sigma}}$ and returning $\sum_i x_i \boldsymbol{h}_i$. We will repeatedly use the observation that if $\boldsymbol{x}$ is sampled from $D_{\boldsymbol{\Sigma}}^H$ and t belongs to $K_{\mathbb{R}}$, then $t \cdot \boldsymbol{x}$ is distributed as $D_{\boldsymbol{\Sigma}'}^H$ with $\boldsymbol{\Sigma}' = \mathrm{diag}(|\sigma_i(t)|) \cdot \boldsymbol{\Sigma} \cdot \mathrm{diag}(|\sigma_i(t)|)$.

For a lattice $\mathcal{L}$ over $V = H^m$ (for some $m \geq 1$) and a coset $\boldsymbol{c} \in V/\mathcal{L}$, we let $D_{\mathcal{L}+\boldsymbol{c},r}$ denote the discretization of $D_{r\mathbf{I}}^H$ over $\mathcal{L} + \boldsymbol{c}$ (we omit the subscript for $D_{\mathcal{L}+\boldsymbol{c},r}$ as all our lattices are over Cartesian powers of H). For $\varepsilon > 0$, we define the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ as the smallest $r > 0$ such that $\rho_{(1/r)\mathbf{I}}(\mathcal{L}^* \setminus \boldsymbol{0}) \leq \varepsilon$. We have the following upper bounds.

Lemma 2.6 ([MR04, Le. 3.3]). *For any lattice $\mathcal{L}$ over H^m and $\varepsilon \in (0, 1)$, we have $\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\log(2mn(1 + 1/\varepsilon))/\pi}/\lambda_1^\infty(\mathcal{L}^*)$.*

Lemma 2.7 (Adapted from [PR07, Le. 6.5]). *For any $\mathcal{O}_K$-ideal I and $\varepsilon \in (0, 1)$, we have $\eta_\varepsilon(I) \leq \sqrt{\log(2n(1 + 1/\varepsilon))/(\pi n)} \cdot (\mathcal{N}(I)\Delta_K)^{1/n}$.*

The following are standard applications of the smoothing parameter.

Lemma 2.8 ([GPV08, Cor. 2.8]). *Let $\mathcal{L}' \subseteq \mathcal{L}$ be full-rank lattices, $\varepsilon \in (0, 1/2)$ and $r \geq \eta_\varepsilon(\mathcal{L}')$. Then $\Delta(D_{\mathcal{L},r} \bmod \mathcal{L}', U(\mathcal{L}/\mathcal{L}')) \leq 2\varepsilon$.*

Lemma 2.9 ([PR06, Le. 2.11]). *Let $\mathcal{L}$ be an n-dimensional lattice, $\varepsilon \in (0, 1/3)$ and $r \geq 4\eta_\varepsilon(\mathcal{L})$. Then $D_{\mathcal{L},r}(\boldsymbol{0}) \leq 2^{-2n+1}$.*

Lemma 2.10 (Adapted from [MR04, Le. 4.4]). *Let $\mathcal{L}$ be an n-dimensional lattice, $\varepsilon \in (0, 1/3)$ and $r \geq \eta_\varepsilon(\mathcal{L})$. Then $\mathrm{Pr}_{\boldsymbol{x} \hookleftarrow D_{\mathcal{L},r}}[\|\boldsymbol{x}\| \geq 2r\sqrt{n}] \leq 2^{-2n}$.*

2.3   Computational Problems

We now formally define the computational problems we will study.

Definition 2.11 (RLWE and PLWE distributions). *Let K a degree n number field defined by f, $\mathcal{O}_K$ its ring of integers, $\boldsymbol{\Sigma} \succ 0$ and $q \geq 2$.*

For $s \in \mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$, we define the dual-RLWE distribution $\mathcal{A}_{s,\boldsymbol{\Sigma}}^\vee$ as the distribution over $\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K^\vee$ obtained by sampling $a \hookleftarrow U(\mathcal{O}_K/q\mathcal{O}_K)$, $e \hookleftarrow D_{\boldsymbol{\Sigma}}^H$ and returning the pair $(a, a \cdot s + e)$.

For $s \in \mathcal{O}_K/q\mathcal{O}_K$, we define the primal-RLWE distribution $\mathcal{A}_{s,\boldsymbol{\Sigma}}$ as the distribution over $\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K$ obtained by sampling $a \hookleftarrow U(\mathcal{O}_K/q\mathcal{O}_K)$, $e \hookleftarrow D_{\boldsymbol{\Sigma}}^H$ and returning the pair $(a, a \cdot s + e)$.

For $s \in \mathbb{Z}_q[x]/f$, we define the PLWE distribution $\mathcal{B}_{s,\boldsymbol{\Sigma}}$ as the distribution over $\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ obtained by sampling $a \hookleftarrow U(\mathbb{Z}_q[x]/f)$, $e \hookleftarrow D_{\boldsymbol{\Sigma}}$ and returning the pair $(a, a \cdot s + e)$ (with $\mathbb{R}_q = \mathbb{R}/q\mathbb{Z}$).

In the definition above, we identified the support H of $D_{\boldsymbol{\Sigma}}^H$ with $K_\mathbb{R}$, and the support $\mathbb{R}^n$ of $D_{\boldsymbol{\Sigma}}$ with $\mathbb{R}[x]/f$. Note that sampling from $\mathcal{A}_{s,\boldsymbol{\Sigma}}^\vee$ and $\mathcal{A}_{s,\boldsymbol{\Sigma}}$ seems to require the knowledge of a basis of $\mathcal{O}_K$. It is not known to be computable in polynomial-time from a defining polynomial f of an arbitrary K. In this article, we assume that a basis of $\mathcal{O}_K$ is known.

Definition 2.12 (The RLWE and PLWE problems). *We use the same notations as above. Further, we let $\mathcal{E}_\succ$ be a subset of $\Sigma \succ 0$ and $D_\succ$ be a distribution over $\Sigma \succ 0$.*

Search dual-RLWE$_{q,\mathcal{E}_\succ}$ (resp. primal-RLWE and PLWE) consists in finding s from a sampler from $\mathcal{A}_{s,\Sigma}^\vee$ (resp. $\mathcal{A}_{s,\Sigma}$ and $\mathcal{B}_{s,\Sigma}$), where $s \in \mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ (resp. $s \in \mathcal{O}_K/q\mathcal{O}_K$ and $s \in \mathbb{Z}_q[x]/f$) and $\Sigma \in \mathcal{E}_\succ$ are arbitrary.

Decision dual-RLWE$_{q,D_\succ}$ (resp. primal-RLWE and PLWE) consists in distinguishing between a sampler from $\mathcal{A}_{s,\Sigma}^\vee$ (resp. $\mathcal{A}_{s,\Sigma}$ and $\mathcal{B}_{s,\Sigma}$) and a uniform sampler over $\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K^\vee$ (resp. $\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K$ and $\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$), with non-negligible probability over $s \hookleftarrow \mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ (resp. $s \in \mathcal{O}_K/q\mathcal{O}_K$ and $s \in \mathbb{Z}_q[x]/f$) and $\Sigma \hookleftarrow D_\succ$.

The problems above are in fact defined for sequences of number fields of growing degrees n such that the bit-size of the problem description grows at most polynomially in n. The run-times, success probabilities and distinguishing advantages of the algorithms solving the problems are considered asymptotically as functions of n.

The following reduction from dual-RLWE to primal-RLWE is a consequence of Lemma 2.4. A proof is given in the full version.

Theorem 2.13 (Adapted from [Pei16, Se. 2.3.2]). *Let $\Sigma \succ 0$ and $s \in \mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$. Let $t \in (\mathcal{O}_K^\vee)^{-1}$ such that $t(\mathcal{O}_K^\vee) + q\mathcal{O}_K = \mathcal{O}_K$. Then the map $(a,b) \mapsto (a, t \cdot b)$ transforms $\mathcal{A}_{s,\Sigma}^\vee$ to $\mathcal{A}_{t\cdot s,\Sigma'}$ and $U(\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K^\vee)$ into $U(\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K)$, with $\Sigma' = \mathrm{diag}(|\sigma_i(t)|) \cdot \Sigma \cdot \mathrm{diag}(|\sigma_i(t)|)$. The natural inclusion $\mathcal{O}_K \to \mathcal{O}_K^\vee$ induces a map that transforms $U(\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K)$ to $U(\mathcal{O}_K/q\mathcal{O}_K \times K_\mathbb{R}/q\mathcal{O}_K^\vee)$, and $\mathcal{A}_{s,\Sigma}$ to $\mathcal{A}_{s,\Sigma}^\vee$.*

We will consider variants of the decision problems for which the distinguishing must occur for all $s \in \mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ (resp. $s \in \mathcal{O}_K/q\mathcal{O}_K$ and $s \in \mathbb{Z}_q[x]/f$) and all $\Sigma \succ 0$ rather than with non-negligible probability over s. We call this variant worst-case decision dual-RLWE (resp. primal-RLWE and PLWE). Under some conditions on $D_\succ$ and $\mathcal{E}_\succ$, these variants are computationally equivalent.

Lemma 2.14 (Adapted from [LPR10, Se. 5.2]). *We use the same notations as above. If $\Pr_{\Sigma \hookleftarrow D_\succ}[\Sigma \notin \mathcal{E}_\succ] \leq 2^{-n}$, then decision dual-RLWE$_{q,D_\succ}$ (resp. primal-RLWE and PLWE) reduces to worst-case decision dual-RLWE$_{q,\mathcal{E}_\succ}$ (resp. primal-RLWE and PLWE).*

Assume further that $D_\succ$ can be sampled from in polynomial-time. If $\max_{\Sigma \in \mathcal{E}_\succ} R(D_\succ \| D_\succ + \Sigma) \leq \mathrm{poly}(n)$, then worst-case decision dual-RLWE$_{q,\mathcal{E}_\succ}$ (resp. primal-RLWE and PLWE) reduces to decision dual-RLWE$_{q,D_\succ}$ (resp. primal-RLWE and PLWE).

Note that it is permissible to use the Rényi divergence here even though we are considering decision problems. Indeed, the argument is applied to the random choice of the noise distribution and not to the distinguishing advantage. The same argument has been previously used in [LPR10, Se. 5.2].

Proof. The first statement is direct. We prove the second statement only for dual-RLWE, as the proofs for primal-RLWE and PLWE are direct adaptations. Assume we are given a sampler that outputs (a_i, b_i) with $a_i \hookleftarrow U(\mathcal{O}_K/q\mathcal{O}_K)$ and b_i either uniform in $K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ or of the form $b_i = a_i s + e_i$ with $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ and $e_i \hookleftarrow D_{\mathbf{\Sigma}}^H$. The reduction proceeds by sampling $s' \hookleftarrow U(\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee})$ and $\mathbf{\Sigma}' \hookleftarrow D_{\succ}$, and mapping all input (a_i, b_i)'s to $(a_i', b_i') = (a_i, b_i + a_i s' + e_i')$ with $e_i' \hookleftarrow D_{\mathbf{\Sigma}'}^H$. This transformation maps the uniform distribution to itself, and $\mathcal{A}_{s,\mathbf{\Sigma}}^{\vee}$ to $\mathcal{A}_{s+s',\mathbf{\Sigma}''}^{\vee}$ with $\mathbf{\Sigma}_{ij}'' = \mathbf{\Sigma}_{ij} + \mathbf{\Sigma}_{ij}'$ for all i, j. If the success probability (success being enjoying a non-negligible distinguishing advantage) over the error parameter sampled from $D_{\succ}$ is non-negligible, then so is it for the error parameter sampled $D_{\succ} + \mathbf{\Sigma}$, as, by assumption, the Rényi divergence $R(D_{\succ} \| D_{\succ} + \mathbf{\Sigma})$ is polynomially bounded. $\square$

Many choices of $D_{\succ}$ and $\mathcal{E}_{\succ}$ satisfy the conditions of Lemma 2.14. The following is inspired from [LPR10, Se. 5.2]. We define the distribution $\mathcal{E}_{\succ}$ as follows, for an arbitrary r: Let $s_{ij} = r^2(1 + nx_{ij})$ for all $i > j$, $s_{ii} = r^2(1 + n^3 x_{ii})$ for all i and $s_{ij} = s_{ji}$ for all $i < j$, where the x_{ij}'s are independent samples from the $\Gamma(2, 1)$ distribution (of density function $x \mapsto x \exp(-x)$); the output matrix is $(s_{ij})_{ij}$. Note that it is symmetric and strictly diagonally dominant (and hence $\succ 0$) with probability $1 - 2^{-\Omega(n)}$. Then the set of all $\Sigma \succ 0$ with coefficients of magnitudes $\leq r^2 n^4$ satisfies the first condition of Lemma 2.14, and the set of all $\Sigma \succ 0$ with coefficients of magnitudes $\leq r^2$ satisfies the second condition of Lemma 2.14. We can hence switch from one variant to the other while incurring an error rate increase that is $\leq \mathrm{poly}(n)$.

3   Controlling Noise Growth in Dual to Primal Reduction

The reduction of Theorem 2.13 is built upon the existence of t as in Lemma 2.4. While this existence is guaranteed constructively by [LPR10], the size is not controlled by the construction. Another t that satisfies the conditions is $t = f'(\alpha)$, where f' is the derivative of f defining $K = \mathbb{Q}[\alpha]$. Indeed, from [Conb, Rem. 4.5], we know that $f'(\alpha) \in (\mathcal{O}_K^{\vee})^{-1}$. However, the noise growth incurred by multiplication by $f'(\alpha)$ may be rather large in general: we have $N(f'(\alpha)) = \Delta_f = [\mathcal{O}_K : \mathbb{Z}[\alpha]]^2 \cdot \mathcal{N}((\mathcal{O}_K^{\vee})^{-1})$.

In this section, we give a probabilistic proof that adequate t's with controlled size can be found by Gaussian sampling.

Let I and J be integral ideals of $\mathcal{O}_K$. Theorem 3.1 below states that a Gaussian sample t in I is such that $t \cdot I^{-1} + J = \mathcal{O}_K$ with non-negligible probability. The main technical hurdle is to show that the sample is not trapped in IJ' with J' a non-trivial factor of J. We handle this probability in different ways depending on the algebraic norm of J', extending an idea used in [SS13, Se. 4].

- For small-norm factors J' of J, the Gaussian folded modulo IJ' is essentially uniform over I/IJ', by Lemma 2.8. This requires the standard deviation parameter s to be above the smoothing parameter of IJ'. We use the smoothing parameter bound from Lemma 2.7.

- For large-norm factors J', we argue that the non-zero points of IJ' are very unlikely to be hit, thanks to the Gaussian tail bound given in Lemma 2.10 and the fact that the lattice minimum of IJ' is large, by Lemma 2.5.
- For middle-norm factors J', neither of the arguments above applies. Instead, we bound the probability that t belongs to IJ' by the probability that t belongs to IJ'', where J'' is a non-trivial factor of J', and use the first argument above. The factor J'' must be significantly denser than J' so that we have smoothing. But it should also be significantly sparser than $\mathcal{O}_K$ so that the upper bound is not too large.

Setting the standard deviation parameter of the discrete Gaussian so that at least one of the three arguments above applies is non-trivial. In particular, this highly depends on how the ideal J factors into primes (whether the pieces are numerous, balanced, unbalanced, etc.). The choice we make below works in all cases while still providing a reasonably readable proof and still being sufficient for our needs, from an asymptotic perspective. In many cases, better choices can be made. If J is prime, we can take a very small s and use only the second argument. If all factors of J are small, there is good enough 'granularity' in the factorization to use the third argument, and again s can be chosen very small.

Theorem 3.1. *Let I and J be integral $\mathcal{O}_K$-ideals, and write $J = \mathfrak{p}_1^{e_1} \ldots \mathfrak{p}_k^{e_k}$ for some prime ideals $\mathfrak{p}_i$. We sort the $\mathfrak{p}_i$'s by non-decreasing algebraic norms. Assume that we can take $\delta \in [\frac{4n + \log_2 \Delta_K}{\log_2 \mathcal{N}(J)}, 1]$.[2] We define:*

$$s = \begin{cases} \left(\mathcal{N}(J)^{1/2} \mathcal{N}(I) \Delta_K \right)^{1/n} & \text{if } \mathcal{N}(\mathfrak{p}_k) \geq \mathcal{N}(J)^{1/2+\delta}, \\ \left(\mathcal{N}(J)^{1/2+2\delta} \mathcal{N}(I) \Delta_K \right)^{1/n} & \text{else.} \end{cases}$$

Then we have

$$\Pr_{t \leftarrow D_{I,s}} [tI^{-1} + J = \mathcal{O}_K] \geq 1 - \frac{k}{\mathcal{N}(\mathfrak{p}_1)} - 2^{-n+4}.$$

Proof. We bound the probability P of the negation, from above. We have

$$P = \Pr_{t \leftarrow D_{I,s}} \left[t \in \bigcup_{i \in [k]} I\mathfrak{p}_i \right] = \sum_{S \subseteq [k], S \neq \emptyset} (-1)^{|S|+1} \cdot \Pr_{t \leftarrow D_{I,s}} \left[t \in I \cdot \prod_{i \in S} \mathfrak{p}_i \right].$$

We rewrite it as $P = P_1 + P_2$ with

$$P_1 = \sum_{S \subseteq [k], S \neq \emptyset} (-1)^{|S|+1} \frac{1}{\prod_{i \in S} \mathcal{N}(\mathfrak{p}_i)} = 1 - \prod_{i \in [k]} \left(1 - \frac{1}{\mathcal{N}(\mathfrak{p}_i)} \right),$$

$$P_2 = \sum_{S \subseteq [k], S \neq \emptyset} (-1)^{|S|+1} \left(\Pr_{t \leftarrow D_{I,s}} \left[t \in I \cdot \prod_{i \in S} \mathfrak{p}_i \right] - \prod_{i \in S} \frac{1}{\mathcal{N}(\mathfrak{p}_i)} \right).$$

---

[2] The parameter δ should be thought as near 0. It can actually be chosen such if $\mathcal{N}(J)$ is sufficiently large.

We have $P_1 \leq 1 - (1 - 1/\mathcal{N}(\mathfrak{p}_1))^k \leq k/\mathcal{N}(\mathfrak{p}_1)$. Our task is now to bound P_2.

Assume first that $\mathcal{N}(\mathfrak{p}_k) \geq \mathcal{N}(J)^{1/2+\delta}$. This implies that $\prod_{i \in S} \mathcal{N}(\mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2-\delta}$ for all $S \subseteq [k]$ not containing k. By Lemma 2.7, we have $s \geq \eta_\varepsilon(I \prod_{i \in S} \mathfrak{p}_i)$ for all such S's, with $\varepsilon = 2^{-2n}$. We "smooth" out those ideals, i.e., we use Lemma 2.8 to obtain, for all $S \subseteq [k] \setminus \{k\}$:

$$\left| \Pr_{t \leftarrow D_{I,s}}[t \in I \cdot \prod_{i \in S} \mathfrak{p}_i] - \prod_{i \in S} \frac{1}{\mathcal{N}(\mathfrak{p}_i)} \right| \leq 2\varepsilon.$$

Now if S is a subset containing k, then we have $\mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \geq \mathcal{N}(J)^{1/2+\delta}$. By Lemma 2.5, we have $\lambda_1(I \prod_{i \in S} \mathfrak{p}_i) \geq \sqrt{n} \cdot \mathcal{N}(I)^{1/n} \mathcal{N}(J)^{(1/2+\delta)/n}$. On the other hand, by Lemma 2.10, we have $\Pr_{t \leftarrow D_{I,s}}[\|t\| \geq 2s\sqrt{n}] \leq 2^{-2n}$. Thanks to our choice of s, the assumption on δ and Lemma 2.9, we obtain

$$\Pr_{t \leftarrow D_{I,s}}[t \in I \prod_{i \in S} \mathfrak{p}_i] \leq \Pr_{t \leftarrow D_{I,s}}[t = 0] + 2^{-2n} \leq 2^{-2n+2}.$$

This allows us to bound P_2 as follows:

$$P_2 \leq 2^k \cdot \left(\varepsilon + 2^{-2n+2} + \mathcal{N}(J)^{-(1/2+\delta)} \right).$$

By assumption on δ, we have $\mathcal{N}(J) \geq 2^{2n}$ and $P_2 \leq 2^{-n+3}$. This completes the proof for the large $\mathcal{N}(\mathfrak{p}_k)$ case.

Now, assume that $\mathcal{N}(\mathfrak{p}_k) < \mathcal{N}(J)^{1/2+2\delta}$. Then, as above, the definition of s implies that, for any $S \subseteq [k]$ with $\mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2+\delta}$, we have $|\Pr[t \in I \prod_{i \in S} \mathfrak{p}_i] - 1/\prod_{i \in S} \mathcal{N}(\mathfrak{p}_i)| \leq 2^{-2n+1}$. Also as above, if we have $\mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \geq \mathcal{N}(J)^{1/2+3\delta}$, then $\lambda_1(I \prod_{i \in S} \mathfrak{p}_i)$ is too large for a non-zero element of $I \prod_{i \in S} \mathfrak{p}_i$ to be hit with significant probability. Assume finally that

$$\mathcal{N}(J)^{1/2+2\delta} \leq \mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2+3\delta}.$$

As $\mathcal{N}(\mathfrak{p}_k) < \mathcal{N}(J)^{1/2+\delta}$, there exists $S' \subseteq S$ such that

$$\mathcal{N}(J)^{\delta} \leq \mathcal{N}(\prod_{i \in S'} \mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2+2\delta}.$$

By inclusion, we have that $\Pr[t \in I \prod_{i \in S} \mathfrak{p}_i] \leq \Pr[t \in I \prod_{i \in S'} \mathfrak{p}_i]$. Now, as the norm of $\prod_{i \in S'} \mathfrak{p}_i$ is small enough, we can use the smoothing argument above to claim that

$$\Pr_{t \leftarrow D_{I,s}}[t \in I \prod_{i \in S} \mathfrak{p}_i] \leq 2^{-2n+1} + \frac{1}{\mathcal{N}(\prod_{i \in S'} \mathfrak{p}_i)} \leq 2^{-2n+1} + \frac{1}{\mathcal{N}(J)^{\delta}}.$$

By assumption on δ, the latter is $\leq 2^{-n+2}$. Collecting terms allows to complete the proof. $\qquad\square$

The next corollary shows that the needed t can be found with non-negligible probability.

Corollary 3.2. *Let I be an integral $\mathcal{O}_K$-ideal. Let $q \geq \max(2n, 2^{16} \cdot \Delta_K^{8/n})$ be a prime rational integer and $\mathfrak{p}_k$ a prime factor of $q\mathcal{O}_K$ with largest norm. We define:*

$$s = \begin{cases} q^{1/2} \cdot (\mathcal{N}(I)\Delta_K)^{1/n} & \text{if } \mathcal{N}(\mathfrak{p}_k) \geq q^{(5/8)\cdot n}, \\ q^{3/4} \cdot (\mathcal{N}(I)\Delta_K)^{1/n} & \text{else.} \end{cases}$$

Then, for sufficiently large n, we have

$$\Pr_{t \hookleftarrow D_{I,s}}[tI^{-1} + q\mathcal{O}_K = \mathcal{O}_K] \geq 1/2.$$

Proof. The result follows from applying Theorem 3.1 with $J = q\mathcal{O}_K$ and $\delta = 1/8$. The first lower bound on q ensures that $k/\mathcal{N}(\mathfrak{p}_1) \leq 1/2$, where $k \leq n$ denotes the number of prime factors of $q\mathcal{O}_K$ and $\mathfrak{p}_1$ denotes a factor with smallest algebraic norm. The second lower bound on q ensures that we can indeed set $\delta = 1/8$. □

We insist again on the fact that the required lower bounds on s can be much improved under specific assumptions on the factorization of q. For example, one could choose a q such that all the factors of $q\mathcal{O}_K$ have large norms, by sampling q randomly and checking its primality and the factorization of the defining polynomial f modulo q. In that case, the factors $q^{1/2}$ and $q^{3/4}$ can be decreased drastically.

We note that if the noise increase incurred by a reduction from an LWE-type problem to another is bounded as $n_1^c \cdot q_2^c$ for some $c_1 < 1$ and some $c_2 > 0$, then one may set the working modulus q so that the starting LWE problem has a sufficient amount of noise to not be trivially easy to solve, and the ending LWE problem has not enough noise to be information-theoretically impossible to solve (else the reduction would be vacuous). Indeed, it suffices to set q sufficiently larger than $n^{c_1/(1-c_2)}$.

4   From Primal-**RLWE** to **PLWE**

The goal of this section is to describe a reduction from primal-RLWE to PLWE. As an intermediate step, we first consider a reduction from primal-RLWE to a variant PLWE$^\sigma$ of PLWE where the noise is small with respect to the Minkowski embedding rather than the coefficient embedding. Then, we assess the noise distortion when looking at its Minkowski embedding versus its coefficient embedding.

If $K = \mathbb{Q}[x]/f$ for some $f = \prod_{j \leq n}(x - \alpha_j)$, the associated Vandermonde matrix V_f has jth row $(1, \alpha_j, \ldots, \alpha_j^{n-1})$ and corresponds to the linear map between the coefficient and Minkowski embedding spaces. Thus a good approximation of the distortion is given by the condition number $\text{Cond}(V_f) = s_n/s_1$, where the s_i's refer to the largest/smallest singular values of V_f. As we also have $\text{Cond}(V_f) = \|V_f\| \cdot \|V_f^{-1}\|$, these matrix norms also quantify how much V_f distorts the space. For a restricted, yet exponentially large, family of polynomials

defining number fields, we show that both $\|V_f\|$ and $\|V_f^{-1}\|$ are polynomially bounded.

To do this, we start from $f_{n,a} = x^n - a$ whose distortion is easily computable. Then we add a "small perturbation" to this polynomial. Intuitively, the roots of the resulting polynomial should not move much, so that the norms of the "perturbed" Vandermonde matrices should be essentially the same. We formalize this intuition in Sect. 4.2 and locate the roots of the perturbed polynomial using Rouché's theorem.

Mapping a sample of PLWE^σ to a sample of the corresponding PLWE simply consists in changing the geometry of the noise distribution. A noise distribution with covariance matrix $\mathbf{\Sigma}$ in the Minkowski embedding corresponds to a noise distribution of covariance matrice $(V_f^{-1})^T \mathbf{\Sigma} V_f^{-1}$ in the coefficient space. The converse is also true, replacing V_f^{-1} by V_f. Moreover, the noise growths incurred by the reductions remain limited whenever $\|V_f\|$ and $\|V_f^{-1}\|$ are small.

Overall, reductions between primal-RLWE to PLWE can be obtained by combining Theorems 4.2 and 4.7 below (with Lemma 2.14 to randomize the noise distributions).

4.1   Reducing Primal-RLWE to PLWE$^\sigma$

We keep the notations of the previous section, and let $\mathbb{Z}[x]/(f) = \mathcal{O}$.

Definition 4.1 (The PLWE^σ problem). *Let also $\mathbf{\Sigma}$ be a positive definite matrix, and $q \geq 2$. For $s \in \mathcal{O}/q\mathcal{O}$, we define the PLWE^σ distribution $\mathcal{B}_{s,\mathbf{\Sigma}}^\sigma$ as the distribution over $\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O}$ obtained by sampling $a \hookleftarrow U(\mathcal{O}/q\mathcal{O})$, $e \hookleftarrow D_{\mathbf{\Sigma}}^H$ and returning the pair $(a, a \cdot s + e)$*

Let $D_\succ$ be a distribution over $\Sigma \succ 0$. Decision PLWE^σ consists in distinguishing between a sampler from $\mathcal{B}_{s,\mathbf{\Sigma}}^\sigma$ and a uniform sampler over $\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O}$, with non-negligible probability over $s \hookleftarrow \mathcal{O}/q\mathcal{O}$ and $\mathbf{\Sigma} \hookleftarrow D_\succ$.

Theorem 4.2. *Assume that $q\mathcal{O}_K + \mathcal{C}_{\mathcal{O}} = \mathcal{O}_K$. Let $\mathbf{\Sigma}$ be a positive definite matrix and $s \in \mathcal{O}_K/q\mathcal{O}_K$. Let $t \in \mathcal{C}_{\mathcal{O}}$ such that $t\mathcal{C}_{\mathcal{O}}^{-1} + q\mathcal{O}_K = \mathcal{O}_K$. Then the map $(a, b) \mapsto (t \cdot a, t^2 \cdot b)$ transforms $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$ to $U(\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O})$ and $\mathcal{A}_{s,\mathbf{\Sigma}}$ to $\mathcal{B}_{t \cdot s,\mathbf{\Sigma}'}^\sigma$, where the new covariance is $\mathbf{\Sigma}' = \mathrm{diag}(|\sigma(t_i)|^2) \cdot \mathbf{\Sigma} \cdot \mathrm{diag}(|\sigma_i(t)|^2)$.*

Let $\mathcal{B}_{s,\mathbf{\Sigma}}^\sigma$ be a PLWE^σ distribution. The natural inclusion $\mathcal{O} \to \mathcal{O}_K$ induces a map that transforms $U(\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O})$ to $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$ and $\mathcal{B}_{s,\mathbf{\Sigma}}^\sigma$ to $\mathcal{A}_{s,\mathbf{\Sigma}}$.

Proof. Let $(a, b = a \cdot s + e)$ be distributed as $\mathcal{A}_{s,\mathbf{\Sigma}}$. Let $a' = t \cdot a$ and $b' = t^2 \cdot b = a' \cdot (t \cdot s) + e'$, with $e' = t^2 \cdot e$. Then a' is uniformly distributed in $\mathcal{C}_{\mathcal{O}}/q\mathcal{C}_{\mathcal{O}}$ by applying Lemma 2.4 for $I = \mathcal{C}_{\mathcal{O}}$, $J = q\mathcal{O}_K$ and $\mathcal{M} = \mathcal{O}_K$. It is also uniformly distributed in $\mathcal{O}/q\mathcal{O}$ by combining Lemmas 2.2 and 2.3. The noise follows the claimed distribution, see the observation in Sect. 2.2. The fact that $t \cdot s \in \mathcal{O}/q\mathcal{O}$ completes the proof that $\mathcal{A}_{s,\mathbf{\Sigma}}$ is mapped to $\mathcal{B}_{t \cdot s,\mathbf{\Sigma}'}^\sigma$.

Now, let (a, b) be uniform in $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$. We already know that a' is uniformly distributed in $\mathcal{O}/q\mathcal{O}$. Let us now consider the distribution of b'. Thanks

to the assumption on $q\mathcal{O}_K$, we also have $t^2 \mathcal{C}_{\mathcal{O}}^{-1} + q\mathcal{O}_K = \mathcal{O}_K$. Therefore, by Lemma 2.4, multiplication by t^2 induces an isomorphism $\mathcal{O}_K/q\mathcal{O}_K \simeq \mathcal{C}_{\mathcal{O}}/q\mathcal{C}_{\mathcal{O}}$, and hence, by Lemmas 2.2 and 2.3, an isomorphism $\mathcal{O}_K/q\mathcal{O}_K \simeq \mathcal{O}/q\mathcal{O}$. This gives the first reduction.

We now turn to the converse reduction. By coprimality and Lemmas 2.2 and 2.4, we have $|\mathcal{O}/q\mathcal{O}| = |\mathcal{O}_K/q\mathcal{O}_K|$. This implies that any (a, b) uniform in $\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O}$ is also uniform in $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$. The inclusion $\mathcal{O} \subseteq \mathcal{O}_K$ allows to conclude.   □

As Theorem 2.13, Theorem 4.2 relies on a the existence of a good multiplier. Writing $K = \mathbb{Q}[x]/(f) = \mathbb{Q}[\alpha]$ and $\mathcal{O} = \mathbb{Z}[\alpha]$, the element $f'(\alpha)$ again satisfies the constraints. Indeed, we know that $\mathcal{O}^{\vee} = \frac{1}{f'(\alpha)}\mathcal{O}$ (see [Conb, Th. 3.7]), and we have the inclusion $\mathcal{O}_K \subseteq \mathcal{O}^{\vee}$. Multiplying by $f'(\alpha)$, we obtain $f'(\alpha)\mathcal{O}_K \subseteq \mathcal{O}$. By definition, this means that $f'(\alpha) \in \mathcal{C}_{\mathcal{O}}$, as claimed. While a large $f'(\alpha)$ would mean a large noise growth in the primal-RLWE to PLWE$^{\sigma}$ reduction, we described in Sect. 3 how to find a smaller adequate multiplier if needed.

We have $\mathcal{N}(f'(\alpha)) = [\mathcal{O}_K : \mathbb{Z}[\alpha]]^2 \cdot \Delta_K$, and, from [Ste17, p. 48], the prime factors of $[\mathcal{O}_K : \mathbb{Z}[\alpha]]$ are exactly those of $\mathcal{N}(\mathcal{C}_{\mathcal{O}})$. Provided the valuations are not too high, there should be smaller elements in $\mathcal{C}_{\mathcal{O}}$ than $f'(\alpha)$. We provide in the full version concrete examples of number fields with defining polynomials f such that the norm of $f'(\alpha)$ is considerably larger than both the norms of $\mathcal{C}_{\mathcal{O}}$ and $(\mathcal{O}_K^{\vee})^{-1}$.

4.2   Distortion Between Embeddings

To bound the norms of a Vandermonde matrix associated to a polynomial and its inverse, we study the magnitude of the roots and their pairwise distances. It is known that $\|V\|^2 = \mathrm{Tr}(V^*V)$, where $*$ denotes the transpose-conjugate operator. For Vandermonde matrices, this gives

$$\|V_f\|^2 = \sum_{j \in [n]} \sum_{k \in [n]} |\alpha_j|^{2(k-1)}, \tag{1}$$

which can be handled when the magnitudes of the α_j's are known. The entries of $V_f^{-1} = (w_{ij})$ have well-known expressions as:

$$w_{ij} = (-1)^{n-i} \frac{e_{n-i}(\overline{\boldsymbol{\alpha}}^j)}{\prod_{k \neq j}(\alpha_j - \alpha_k)}, \tag{2}$$

where $e_0 = 1$, e_j for $j > 0$ stands for the elementary symmetric polynomial of total degree j in $n-1$ variables, and $\overline{\boldsymbol{\alpha}}^j = (\alpha_1, \ldots, \alpha_{j-1}, \alpha_{j+1}, \ldots, \alpha_n)$, the vector of all roots but α_j. We have the following useful relations with the symmetric functions E_i of all the roots (for all j):

$$\begin{cases} E_1(\boldsymbol{\alpha}) = & \alpha_j + e_1(\overline{\boldsymbol{\alpha}}^j), \\ E_i(\boldsymbol{\alpha}) = & \alpha_j e_{i-1}(\overline{\boldsymbol{\alpha}}^j) + e_i(\overline{\boldsymbol{\alpha}}^j) \text{ for } 2 \leq i \leq n-1, \\ E_n(\boldsymbol{\alpha}) = & \alpha_j e_{n-1}(\overline{\boldsymbol{\alpha}}^j). \end{cases} \tag{3}$$

Combining (3) with Vieta's formulas, bounds on the magnitudes of the roots leads to bounds on the numerators of the w_{ij}'s. The denominators encode the separation of the roots, and deriving a precise lower bound turns out to be the main difficulty. By differentiating $f(x) = \prod_{j\in[n]}(x-\alpha_j)$, we note that $\prod_{k\neq j}|\alpha_j - \alpha_k| = |f'(\alpha_j)|$.

A FAMILY OF POLYNOMIALS WITH EASILY COMPUTABLE DISTORTION. We first introduce a family of polynomials for which $\|V_f\|$ and $\|V_f^{-1}\|$ are both simple to estimate. For $n \geq 2$ and $a \geq 1$, we define $f_{n,a} = x^n - a$. The roots can be written[3] as $\alpha_j = a^{1/n}e^{2i\pi\frac{j}{n}}$, for $0 \leq j < n$. As these are scalings of the roots of unity, both their magnitude and separation are well-understood. With (1), we obtain $\|V_{f_{n,a}}\| \leq na^{\frac{n-1}{n}} \leq na$.

For any j, we readily compute $|f'_{n,a}(\alpha_j)| = na^{\frac{n-1}{n}}$. Using (3), we observe that $|e_i(\overline{\boldsymbol{\alpha}}^j)| = |\alpha_j|^i$ for $1 \leq i < n$. We obtain that the row norm of $V_{f_{n,a}}^{-1}$ is given by its first row as

$$\sum_{j\in[n]} |w_{1j}| = \frac{1}{na^{\frac{n-1}{n}}} \cdot \sum_{j\in[n]} |\alpha_j|^{n-1} = 1,$$

from which it follows that $\|V_{f_{n,a}}^{-1}\| \leq \sqrt{n}$.

SMALL PERTURBATIONS OF $f_{n,a}$. Let $P(x) = \sum_{1\leq j\leq\rho\cdot n} p_j x^j$ for some constant $\rho \in (0,1)$, where the p_j's are a priori complex numbers. Locating the roots of $g_{n,a} = f_{n,a} + P$ is our first step towards estimating $\|V_{g_{n,a}}\|$ and $\|V_{g_{n,a}}^{-1}\|$. We will use the following version of Rouché's theorem.

Theorem 4.3 (Rouché, adapted from [Con95, pp. 125–126]). *Let f, P be complex polynomials, and let D be a disk in the complex plane. If for any z on the boundary ∂D we have $|P(z)| < |f(z)|$, then f and $f + P$ have the same number of zeros inside D, where each zero is counted as many times as its multiplicity.*

The lemma below allows to determine sufficient conditions on the parameters such that the assumptions of Theorem 4.3 hold. We consider small disks $D_k = D(\alpha_k, 1/n)$ of radius $1/n$ around the roots $\alpha_1, \ldots, \alpha_n$ of $f_{n,a}$, and we let ∂D_k denote their respective boundaries. We let $\|P\|_1 = \sum_j |p_j|$ denote the 1-norm of P.

Lemma 4.4. *We have, for all $k \leq n$ and $z \in \partial D_k$:*

$$|P(z)| \leq (ae)^\rho \cdot \|P\|_1 \quad and \quad |f_{n,a}(z)| \geq a\left(1 - \cos(a^{-1/n}) - \frac{2e^{a^{-1/n}}}{na^{2/n}}\right).$$

Proof. Write $z = \alpha_k + \frac{e^{it}}{n}$ for some $t \in [0, 2\pi)$. We have $|z| \leq a^{1/n} + 1/n$, and hence $|z|^{\rho n} \leq a^\rho\left(1 + \frac{1}{na^{1/n}}\right)^{\rho n}$. The first claim follows from the inequality $|P(z)| \leq \max(1, |z|^{\rho n}) \cdot \|P\|_1$.

---

[3] For the rest of this section, 'i' will refer to the imaginary unit.

Next, we have $|f_{n,a}(z)| = a|(1 + \frac{e^{it'}}{na^{1/n}})^n - 1|$, where $t' = t - 2k\pi/n$. W.l.o.g., we assume that $k = 0$. Let Log denote the complex logarithm, defined on $\mathbb{C} \setminus \mathbb{R}^-$. Since the power series $\sum_{k \geq 1}(-1)^{k-1}u^k/k$ converges to $\text{Log}(1 + u)$ on the unit disk, we have $\text{Log}(1 + \frac{e^{it}}{na^{1/n}}) = \frac{e^{it}}{na^{1/n}} + \delta$, for some δ satisfying $|\delta| \leq |u| \cdot \sum_{k \geq 1}|u|^k/(k+1) \leq |u|^2$ for $u = \frac{e^{it}}{na^{1/n}}$ (note that it has modulus $\leq 1/n \leq 1/2$). Similarly, we can write $\exp(n\delta) = 1 + \varepsilon$ for some ε satisfying $|\varepsilon| \leq 2n|\delta| \leq 2/(na^{2/n})$. We hence have:

$$|f_{n,a}(z)| = a \cdot |A \cdot (1 + \varepsilon) - 1| \geq a \cdot ||A - 1| - |\varepsilon \cdot A||,$$

with $A = \exp(e^{it}a^{-1/n})$. Elementary calculus leads to the inequalities $|A - 1| > 1 - \cos(a^{-1/n})$ and $|A| \leq e^{a^{-1/n}}$ for all $t \in [0, 2\pi)$. Details can be found in the full version. The second claim follows.   □

We note that when $a = 2^{o(n)}$ and n is sufficiently large, then the lower bound on $|f_{n,a}(z)|$ may be replaced by $|f_{n,a}(z)| > a/3$. To use Rouché's theorem, it is then enough that a, ρ and $\|P\|_1$ satisfy $a > (3e^\rho\|P\|_1)^{\frac{1}{1-\rho}}$. We can now derive upper bounds on the norms of $V_{g_{n,a}}$ and its inverse.

Lemma 4.5. *For any* $a > (\|P\|_1 \cdot C^{-1} \cdot e^\rho)^{\frac{1}{1-\rho}}$ *with* $C = |1 - \cos(a^{-1/n}) - \frac{2e^{a^{-1/n}}}{na^{2/n}}|$, *we have:*

$$\|V_{g_{n,a}}\| \leq ane \quad and \quad \|V_{g_{n,a}}^{-1}\| \leq n^{5/2}(\|P\|_1 + 1)a^{1/n}e^2.$$

Proof. Let $\alpha_j = a^{1/n}e^{2i\pi j/n}$ be the roots of $f_{n,a}$ (for $0 \leq j < n$). Thanks to the assumptions and Lemma 4.4, Theorem 4.3 allows us to locate the roots $(\beta_j)_{0 \leq j < n}$ of $g_{n,a}$ within distance $1/n$ from the α_j's. Up to renumbering, we have $|\alpha_j - \beta_j| \leq 1/n$ for all j. In particular, this implies that $|\beta_j| \leq a^{1/n} + 1/n$ for all j. The first claim follows from (1).

Another consequence is that any power less than n of any $|\beta_j|$ is $\leq ae$. We start the estimation of $\|V_{g_{n,a}}^{-1}\|$ by considering the numerators in (2). Let $k_0 = 1 + \lfloor n(1 - \rho)\rfloor$. For any $k < k_0$, we know that $E_k(\boldsymbol{\beta}) = 0$. Using (3), we obtain $|e_k(\overline{\boldsymbol{\beta}}^j)| = |\beta_j|^k \leq ae$ for $k < k_0$ and that $e_{k_0-1}(\overline{\boldsymbol{\beta}}^j) = (-1)^{k_0-1}\beta_j^{k_0-1}$. Then (3) gives $E_{k_0}(\boldsymbol{\beta}) = (-1)^{k_0}p_{n-k_0} = (-1)^{k_0-1}\beta_j^{k_0} + e_{k_0}(\overline{\boldsymbol{\beta}}^j)$, which implies that $|e_{k_0}(\overline{\boldsymbol{\beta}}^j)| \leq |\beta_j|^{k_0} + |p_{n-k_0}|$. By induction, we obtain, for all $k < n - k_0$:

$$|e_{k_0+k}(\overline{\boldsymbol{\beta}}^j)| \leq |p_{n-k_0-k}| + |p_{n-k_0-k+1}\beta_j| + \cdots + |p_{n-k_0}\beta_j^k| + |\beta_j|^{k_0+k}$$
$$\leq (\|P\|_1 + 1)\max(1, |\beta_j|^n),$$

so that $|e_k(\overline{\boldsymbol{\beta}}^j)| \leq (\|P\|_1 + 1)ae$ for $k \geq 1$.

We now derive a lower bound on the denominators in (2). The separation of the β_j's is close to that of the α_j's. Concretely: $|\beta_j - \beta_k| \geq |\alpha_j - \alpha_k| - 2/n$ for all j, k. Therefore, we have $\prod_{k \neq j}|\beta_j - \beta_k| \geq \prod_{k \neq j}(|\alpha_j - \alpha_k| - 2/n)$. Using the identity $|\alpha_j - \alpha_k| = 2a^{1/n}\sin(|k - j|\pi/n)$ and elementary calculus,

we obtain $\prod_{k \neq j} |\beta_j - \beta_k| \geq a^{\frac{n-1}{n}}/(ne)$. Details can be found in the full version. Thus any coefficient w_{ij} of $V_{g_{n,a}}^{-1}$ satisfies $|w_{ij}| \leq n(\|P\|_1 + 1)a^{1/n}e^2$. The claim follows from equivalence between the row and Frobenius norms.   □

We now assume that the p_j's and a are integers. The following lemma states that, for a prime and sufficiently large, the polynomial $g_{n,a}$ is irreducible, and thus defines a number field.

Lemma 4.6. *Assume that P is an integer polynomial. For any prime $a >$ $\|P\|_1 + 1$, the polynomial $g_{n,a}$ is irreducible over $\mathbb{Q}$.*

Proof. Let β be a root of $g_{n,a}$. Then we have $a = |\beta^n + P(\beta)| \leq |\beta|^n + \|P\|_1 \max(1, |\beta|^n)$. The assumption on a implies that $|\beta| > 1$. In other words, all the roots of $g_{n,a}$ have a magnitude >1. Now, assume by contradiction that $g_{n,a} = h_1 h_2$ for some rational polynomials h_1, h_2. Since $g_{n,a}$ is monic, it is primitive and we can choose h_1, h_2 as integer polynomials. The product of their constant coefficients is then the prime a. Hence the constant coefficient of h_1 or h_2 is ± 1, which contradicts the fact that the roots of $g_{n,a}$ have magnitude >1.   □

Overall, we have proved the following result.

Theorem 4.7. *Let $\rho \in (0,1)$ and $p_j \in \mathbb{Z}$ for $1 \leq j \leq \rho \cdot n$. Then for $a \geq (3e^\rho \|P\|_1)^{1/(1-\rho)}$ smaller than $2^{o(n)}$ and prime, and n sufficiently large, the polynomial $g_{n,a} = x^n + \sum_{1 \leq j \leq \rho \cdot n} p_j x^j + a$ is irreducible over $\mathbb{Q}$ and satisfies:*

$$\|V_{g_{n,a}}\| \leq ane \quad and \quad \|V_{g_{n,a}}^{-1}\| \leq n^{5/2}(\|P\|_1 + 1)a^{1/n}e^2.$$

In particular, if a and $\|P\|_1$ are polynomial in n, then both $\|V_{g_{n,a}}\|$ and $\|V_{g_{n,a}}^{-1}\|$ are polynomial in n.

In the full version of this article, we give another family of well-behaved polynomials.

5   Search to Decision Dual-RLWE

The reduction relies on the recent technique of [PRS17]. To leverage it, we use a generalized Leftover Hash Lemma over rings. The proof generalizes a technique used in [SS11] to the case where the irreducible factors of the defining polynomial (of K) reduced modulo q do not share the same degree. Alternatively, a generalization of the regularity lemma from [LPR13, Se. 7] to arbitrary number fields could be used. Such a generalization may go through and improve our results a little.

5.1   A Ring-Based Leftover Hash Lemma

Let $m \geq 2$. We identify any rank m $\mathcal{O}_K$-module $M \subseteq K^m$ with the lattice $\sigma(M) \subseteq H^m$. For such modules, the dual may be defined as

$$\widehat{M} = \{\mathbf{t} \in K^m \, : \, \forall \mathbf{x} \in M, \mathrm{Tr}(\langle \mathbf{t}, \mathbf{x} \rangle) \in \mathbb{Z}\}.$$

Here $\langle \cdot, \cdot \rangle$ is the K-bilinear map defined by $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^m x_i y_i$. We have $\sigma(\widehat{M}) = \overline{\sigma(M)^*}$ in H^m. For some $q \geq 2$ and a fixed $\mathbf{a} \in (\mathcal{O}_K/q\mathcal{O}_K)^m$, we focus on the modules:

$$L(\mathbf{a}) = \frac{\mathbf{a}}{q}\mathcal{O}_K^\vee + (\mathcal{O}_K^\vee)^m \ \text{ and } \ \mathbf{a}^\perp = \{\mathbf{t} \in \mathcal{O}_K^m \, : \, \langle \mathbf{t}, \mathbf{a} \rangle = 0 \bmod q\mathcal{O}_K\}.$$

To prove our Leftover Hash Lemma variant, the main argument relies on an estimation of $\lambda_1^\infty(\widehat{\mathbf{a}^\perp})$, which is obtained by combining the following two lemmas. The first one was stated in [LS15, Se. 5] without a proof, for the case of cyclotomic fields (this restriction is unnecessary). We give a proof of the general case in the full version of this article.

Lemma 5.1. *Let $q \geq 2$ and $\mathbf{a} \in (\mathcal{O}_K/q\mathcal{O}_K)^m$. Then we have $\widehat{\mathbf{a}^\perp} = L(\mathbf{a})$.*

We now obtain a probabilistic lower bound on $\lambda_1^\infty(\widehat{\mathbf{a}^\perp}) = \lambda_1^\infty(L(\mathbf{a}))$. In full generality, it should depend on the ramification of the selected prime integer q, i.e., the exponents appearing in the factorization of $q\mathcal{O}_K$ in prime ideals. It is a classical fact that the ramified prime integers are exactly the primes dividing the discriminant of the field, so that there are only finitely many such q's. Moreover, it is always possible to use modulus switching techniques [BLP+13,LS15] if q ramifies. Therefore, we consider only the non-ramified case.

Lemma 5.2. *Let $q \geq 2$ a prime that does not divide Δ_K. For any $m \geq 2$ and $\delta > 0$, and except with a probability $\leq 2^{3n(m+1)}q^{-mn\delta}$ over the uniform choice of $\mathbf{a} \in ((\mathcal{O}_K/q\mathcal{O}_K)^\times)^m$, we have:*

$$\lambda_1^\infty(L(\mathbf{a})) \geq \Delta_K^{-1/n} \cdot q^{-\frac{1}{m}-\delta}.$$

Proof. Thanks to the assumption on q, we can write $q\mathcal{O}_K = \mathfrak{p}_1 \ldots \mathfrak{p}_k$ for distinct prime ideals $\mathfrak{p}_i$. By Lemma 2.4 and the Chinese Remainder Theorem, we have $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee \simeq \mathcal{O}_K/q\mathcal{O}_K \simeq \bigoplus_{i=1}^k \mathbb{F}_{q^{d_i}}$, where $q^{d_i} = \mathcal{N}(\mathfrak{p}_i)$.

Let $\mathbf{a} = (a_1, \ldots, a_m)$ sampled uniformly in $((\mathcal{O}_K/q\mathcal{O}_K)^\times)^m$. Fix some bound $B > 0$ and let p_B be the probability that $qL(\mathbf{a}) = \mathbf{a}\mathcal{O}_K^\vee + q(\mathcal{O}_K^\vee)^m$ contains a $\mathbf{t} = (t_1, \ldots, t_m)$ such that $0 < \|\mathbf{t}\|_\infty < B$. Our goal is to bound p_B from above. By the union bound, we have that

$$p_B \leq \sum_{s \in \mathcal{O}_K^\vee/q\mathcal{O}_K^\vee} \sum_{\substack{\mathbf{t} \in (\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee)^m \\ 0 < \|\mathbf{t}\|_\infty < B}} p(\mathbf{t}, s),$$

with $p(\mathbf{t}, s) = \Pr_{\mathbf{a}}[\forall j, t_j = a_j s \bmod q\mathcal{O}_K^\vee]$ for any s and $\mathbf{t}$ over $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$. By independance of the a_j's, we can write $p(\mathbf{t}, s) = \prod_{j\in[m]} p(t_j, s)$ with $p(t_j, s) = \Pr_{a_j}[t_j = a_j s \bmod q\mathcal{O}_K^\vee]$. As $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ and $\mathcal{O}_K/q\mathcal{O}_K$ are isomorphic, estimating this probability amounts to studying the solutions in $(\mathcal{O}_K/q\mathcal{O}_K)^\times$ of the equation $t = as \bmod q\mathcal{O}_K$, for all $t, s \in \mathcal{O}_K/q\mathcal{O}_K$.

Note that if there is an i such that $t = 0 \bmod \mathfrak{p}_i$ and $s \neq 0 \bmod \mathfrak{p}_i$, or viceversa, then there is no solution, so that $p(t, s) = 0$. Now, assume that s and t are 0 modulo the same $\mathfrak{p}_i$'s. Let $S \subseteq [k]$ denote the set of their indices, and let d_S be such that $q^{d_S} = \mathcal{N}(\prod_{i\in S} \mathfrak{p}_i)$. On the one hand, for all $i \in [k] \setminus S$, both t and s are invertible modulo $\mathfrak{p}_i$ so there is exactly one solution modulo those i's. On the other hand, for all $i \in S$, all the elements of $\mathbb{F}_{q^{d_i}}^\times$ are solutions. This gives $\prod_{i\in S}(q^{d_i}-1)$ possibilities out of the $\prod_i (q^{d_i}-1)$ elements of $(\mathcal{O}_K/q\mathcal{O}_K)^\times$. Overall, we obtain that $p(t, s) = \prod_{i\in[k]\setminus S}(q^{d_i}-1)^{-1}$. Hence, for $\mathbf{t}$ with coordinates t_j such that s and all t_j's are 0 modulo the same $\mathfrak{p}_i$'s, we have:

$$p(\mathbf{t}, s) = q^{-m(n-d_S)} \prod_{i\in[k]\setminus S} (1 - \frac{1}{q^{d_i}})^{-m} \leq q^{-m(n-d_S)} \cdot 2^{mk},$$

the last inequality coming from the fact that $1 - 1/q^{d_i} \geq 1/2$ for all i.

Let τ denote the isomorphism mapping $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ to $\mathcal{O}_K/q\mathcal{O}_K$. The probability to bound is now

$$p_B \leq 2^{mk} \cdot \sum_{S\subseteq[k]} \sum_{\substack{\tau(s)\in\mathcal{O}_K/q\mathcal{O}_K \\ \forall i\in S:\mathfrak{p}_i\,|\,\tau(s)}} \sum_{\substack{\tau(\mathbf{t})\in(\mathcal{O}_K/q\mathcal{O}_K)^m \\ 0<\|\mathbf{t}\|_\infty<B \\ \forall j,\forall i\in S:\mathfrak{p}_i\,|\,\tau(t_j)}} q^{-m(n-d_S)}.$$

For any $r > 0$, we let $\mathcal{B}(r)$ denote the (open) ball in H of center 0 and radius r, with respect to the infinity norm. Such a ball has a volume $\mathrm{Vol}(\mathcal{B}(r)) = (2r)^n$. For any $S \subseteq [k]$, we define $N(B, S) = |\mathcal{B}(B) \cap \mathcal{L}(\tau^{-1}(\prod_{i\in S}\mathfrak{p}_i))| - 1$. Since there are 2^k subsets in $[k]$ and q^{n-d_S} elements $\tau(s) \in \mathcal{O}_K/q\mathcal{O}_K$ such that $\mathfrak{p}_i|s$ for all $i \in S$, we have

$$p_B \leq 2^{k(m+1)} \cdot \max_{S\subseteq[k]} \frac{N(B,S)^m}{q^{(n-d_S)(m-1)}}. \tag{4}$$

We now give an upper bound for $N(B, S)$, from which we will obtain the result. Let $I_S = \prod_{i\in S} \mathfrak{p}_i$ and $\lambda_S = \lambda_1^\infty(\tau^{-1}(I_S))$. Observe that any two distinct balls of radius $\lambda_S/2$ and centered around elements of $\mathcal{B}(B) \cap \mathcal{L}(\tau^{-1}(I_S))$ do not intersect. Moreover, all of them are contained in $\mathcal{B}(B + \lambda_S/2)$. This implies that

$$N(B, S) \leq \frac{\mathrm{Vol}(\mathcal{B}(B+\lambda_S/2))}{\mathrm{Vol}(\mathcal{B}(\lambda_S/2))} = \left(\frac{2B}{\lambda_S} + 1\right)^n.$$

It remains to give a lower bound on λ_S. As $\tau^{-1}(I_S) = I_S\mathcal{O}_K^\vee$, we have $\mathcal{N}(\tau^{-1}(I_S)) = q^{d_S}/\Delta_K$. With Lemma 2.5, this gives $\Delta_K^{-1/n} q^{d_S/n} \leq \lambda_S$. If we set $B = \Delta_K^{-1/n} q^\beta$, then $n\beta < d_S$ leads to $N(B, S) = 0$ and $n\beta \geq d_S$ implies the

upper bound $N(B, S) \leq 2^{2n} q^{n\beta - d_S}$. With (4), this gives

$$p_B \leq 2^{(m+1)(k+2n)} \cdot \max_{\substack{S \subseteq [k] \\ d_S \leq n\beta}} q^{m(\beta-1)n + (n - d_S)}.$$

The maximum is reached for $d_S = 0$ (i.e., when $S = \emptyset$). In this case, the exponent of q is $-mn\delta$ for $\beta = 1 - \frac{1}{m} - \delta$. We obtain that $\lambda_1^\infty(qL(\mathbf{a})) \geq \Delta_K^{-1/n} q^{1 - \frac{1}{m} - \delta}$ except with probability $\leq 2^{3n(m+1)} q^{-mn\delta}$.   □

We are now ready to state the variant of the Leftover Hash Lemma.

Theorem 5.3. *Let $q \geq 2$ prime that does not divide Δ_K. Let $\delta > 0, \varepsilon \in (0, 1/2)$ and $m \geq 2$. For a given $\mathbf{a}$ in $((\mathcal{O}_K/q\mathcal{O}_K)^\times)^m$, let $U_{\mathbf{a}}$ be the distribution of $\sum_{i \leq m} t_i a_i$ where the vector $\boldsymbol{t} = (t_1, \ldots, t_m)$ is sampled from $D_{\mathcal{O}_K, s}$ with $s \geq \sqrt{\log(2mn(1 + 1/\varepsilon))/\pi} \cdot \Delta_K^{1/n} q^{1/m + \delta}$. Then, except for $\leq 2^{3n(m+1)} q^{-mn\delta}$ of $\mathbf{a}$'s, the distance to uniformity of $U_{\mathbf{a}}$ is $\leq 2\varepsilon$.*

Proof. First we note that the map $\mathbf{t} \mapsto \sum_{i \leq m} t_i a_i$ is a well-defined surjective $\mathcal{O}_K$-module homomorphism from $\mathcal{O}_K^m$ to $\mathcal{O}_K/q\mathcal{O}_K$, with kernel $\mathbf{a}^\perp$. The distance to uniformity of $U_{\mathbf{a}}$ is hence the same as the distance to uniformity of $\mathbf{t} \bmod \mathbf{a}^\perp$. By Lemma 2.8, the claim follows whenever $s \geq \eta_\varepsilon(\mathbf{a}^\perp)$. By Lemma 2.6, t it suffices to find an appropriate lower bound on $\lambda_1^\infty(L(\mathbf{a}))$. Lemma 5.2 allows to complete the proof.   □

Corollary 5.4 (Leftover Hash lemma). *If $\mathbf{t}$ is sampled from $D_{\mathcal{O}_K, s}$ with $s \geq \sqrt{\log(2mn(1 + 1/\epsilon))/\pi} \cdot \Delta_K^{1/n} q^{1/m + \delta}$, and the a_i's are sampled from $U((\mathcal{O}_K/q\mathcal{O}_K)^\times)$, then:*

$$\Delta\left[\left(a_1, \ldots, a_m, \sum_{i \leq m} t_i a_i\right), U\left(((\mathcal{O}_K/q\mathcal{O}_K)^\times)^m \times \mathcal{O}_K/q\mathcal{O}_K\right)\right]$$
$$\leq 2\varepsilon + 2^{3n(m+1)} \cdot q^{-mn\delta}.$$

5.2 Search RLWE to Decision RLWE

We now give the reduction from search to decision. As all proofs can be done similarly, we focus on the dual-RLWE version of the problems. For the sake of simplicity, we consider only the case of diagonal covariance matrices. The proof readily extends to general covariance matrices. To obtain the reduction, we need to generate suitable new samples from a starting set of samples from search dual-RLWE.

The lemma below is adapted from [LS15, Le. 4.15]. We will use it to analyze the error distribution we get when generating new samples.

Lemma 5.5. *Let $\alpha > 0$, $\mathcal{L}$ a rank-m $\mathcal{O}_K$-module, $\varepsilon \in (0, 1/2)$, a vector $\mathbf{t} \in D_{\mathcal{L}+\mathbf{c}, \mathbf{r}}$ for some $\mathbf{c} \in H^m$, and $e' \in K_\mathbb{R}$ chosen according to D_α^H. If $r_i \geq \eta_\varepsilon(\mathcal{L})$ and $\frac{\alpha}{\delta_i} \geq \eta_\varepsilon(\mathcal{L})$ for all i, then $\Delta(\langle \mathbf{t}, \mathbf{e} \rangle + e', D_{\mathbf{x}}^H) \leq 4\varepsilon$ with $x_i = \sqrt{(r_i \delta_i)^2 + \alpha^2}$ and $\delta_i = (\sum_{k \in [m]} |\sigma_i(e_k)|^2)^{1/2}$ for all i.*

We can now give a reduction from search dual-RLWE to worst-case decision dual-RLWE. It may be combined with the worst-case decision dual-RLWE to decision dual-RLWE from Lemma 2.14.

Theorem 5.6. *Let $r \in (\mathbb{R}^{\geq 0})^n$ be such that $r_i = r_{i+s_2}$ for any $i > s_1$ and $r_i \leq r$ for some $r > 0$. Let $d = \sqrt{n} \cdot \Delta_K^{1/n} q^{1/m+1/n}$, and consider $\Sigma = \{r' : r'_i \leq \sqrt{d^2 \cdot r^2 \cdot m + d^2}\}$. Then there exists a probabilistic polynomial-time reduction from search dual-RLWE$_{q,D_r}$ with $m \leq q/(2n)$ input samples to worst-case decision dual-RLWE$_{q,\Sigma}$.*

Proof. We have m samples $(a_i, b_i = a_i s + e_i) \in \mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ from the dual-RLWE distribution $\mathcal{A}_{s,\mathbf{r}}^{\vee}$, for a uniform $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ that we want to find. This is equivalent to finding the error term $\mathbf{e} = (e_1, \ldots, e_m)$. By assumption on m, the a_i's are all invertible with non-negligible probability. If it is not the case, the reduction aborts. From now on, we hence assume that they are uniformly distributed in $(\mathcal{O}_K/q\mathcal{O}_K)^{\times}$.

We use the same technique as in [PRS17], in that we find the ith embeddings $\sigma_i(e_1), \ldots, \sigma_i(e_m)$ of the error terms by constructing an m-dimensional instance of the Oracle Hidden Center Problem (OHCP). The only difference consists in the way we create the samples that we give to the decision oracle. The reduction uses the dual-RLWE decision oracle to build the oracles $\mathcal{O}_i : \mathbb{R}^m \times \mathbb{R}^{\geq 0} \to \{0,1\}$ for $i \leq s_1$ and $\mathcal{O}_i : \mathbb{C}^m \times \mathbb{R}^{\geq 0} \to \{0,1\}$ for $s_1 < i \leq s_1 + s_2$.

For $i \leq s_1$, we define $k_i : \mathbb{R} \to K_{\mathbb{R}}$ as $k_i(x) = \sigma^{-1}(x \cdot \mathbf{v}_i)$ and for $s_1 < i \leq s_1 + s_2$, we define $k_i : \mathbb{C} \to K_{\mathbb{R}}$ as $k_i(x) = \sigma^{-1}(x \cdot \mathbf{v}_i + \overline{x} \cdot \mathbf{v}_{i+s_2})$, where the $\mathbf{v}_i$'s form the canonical basis of H.

On input $(z_1, \ldots, z_m, \alpha)$, oracle $\mathcal{O}_i$ will output 1 with probability depending on $\exp(\alpha)\|\mathbf{e} - \overline{\mathbf{z}}\|$, where $\overline{\mathbf{z}} = (k_i(z_1), \ldots, k_i(z_m))$. It works as follows. It first chooses a uniform $s' \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$. On input $(z_1, \ldots, z_m, \alpha)$, it samples $\mathbf{t} = (t_1, \ldots, t_m) \in \mathcal{O}_K^m$ Gaussian with parameter $\exp(\alpha) \cdot \sqrt{n} \cdot \Delta_K^{1/n} q^{1/m+1/n}$ and some e' from D_d. The oracle then creates $(a', b') = (\langle \mathbf{t}, \mathbf{a} \rangle, \langle \mathbf{t}, \mathbf{b} - \overline{\mathbf{z}} \rangle + a's' + e')$, where $\mathbf{b} = (b_1, \ldots, b_m)$.

By Corollary 5.4, the distribution of $(\mathbf{a}, \langle \mathbf{t}, \mathbf{a} \rangle)$ is exponentially close to $U(((\mathcal{O}_K/q\mathcal{O}_K)^{\times})^m \times \mathcal{O}_K/q\mathcal{O}_K)$. Since $b_j = a_j s + e_j$ for all j, we get $b' = a'(s + s') + \langle \mathbf{t}, \mathbf{e} - \overline{\mathbf{z}} \rangle + e'$, so oracle $\mathcal{O}_i$ creates RLWE samples for a uniformly distributed $s + s'$, provided the error term follows a suitable distribution. We let $\delta_\ell = (\sum_{j \in [m]} \sigma_\ell(e_j - k_i(z_j))|^2)^{1/2}$ for $\ell \leq n$. In particular, we have $\delta_i = \|\sigma_i(e_1) - z_1, \ldots, \sigma_i(e_m) - z_m\|$. Let us now study the distribution of the error term $\langle \mathbf{t}, \mathbf{e} - \overline{\mathbf{z}} \rangle + e'$. We can see that once the value of $\langle \mathbf{t}, \mathbf{a} \rangle = c$ and the a_i's are known, one can write $\mathbf{t} = (ca_1^{-1}, 0, \ldots, 0) + (-a_1^{-1} \sum_{i \geq 2} t_i a_i, t_2, \ldots, t_m)$, where the second vector belongs to $\mathbf{a}^{\perp}$. This means that the actual support of $\mathbf{t}$ is a shift of the $\mathbf{a}^{\perp}$ lattice by the vector $(ca_1^{-1}, 0, \ldots, 0)$. Using Lemma 5.5, we get that the distribution of the error is $D_{\mathbf{x}}^H$ where $x_j = \sqrt{\exp^2(\alpha) \cdot d^2 \cdot \delta_j^2 + d^2}$.

Let $\mathcal{S}_{i,(z_1,\ldots,z_m,\alpha)}$ be the samples obtained by applying the procedure above many times. Oracle $\mathcal{O}_i$ calls the dual-RLWE decision oracle with these and outputs 1 if and only if the latter accepts. With non-negligible probability over

the choice of the initial errors, the distribution of the samples we get when we call the oracle $\mathcal{O}_i$ on $(\mathbf{0}, 0)$ belongs to the set Σ. One can now show that using the same technique as in [PRS17], it is possible to recover good approximations of the vector $(\sigma_i(e_1), \ldots, \sigma_i(e_m))$. By substracting them from the initial search samples, rounding and then taking the inverses of the a_i's, we obtain s.   □

References

[AD17] Albrecht, M.R., Deo, A.: Large modulus ring-LWE $\geq$ module-LWE. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10624, pp. 267–296. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_10

[ADPS16] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. In: USENIX (2016)

[BBdV+17] Bauch, J., Bernstein, D.J., de Valence, H., Lange, T., van Vredendaal, C.: Short generators without quantum computers: the case of multiquadratics. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 27–59. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_2

[BCLvV16] Bernstein, D.J., Chuengsatiansup, C., Lange, T., van Vredendaal, C.: NTRU Prime (2016). http://eprint.iacr.org/2016/461

[BDK+18] Bos, J.W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Stehlé, D.: CRYSTALS - Kyber: a CCA-secure module-lattice-based KEM. In: EuroS&P (2018)

[BLP+13] Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: STOC (2013)

[CDPR16] Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 559–585. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_20

[CDW17] Cramer, R., Ducas, L., Wesolowski, B.: Short stickelberger class relations and application to ideal-SVP. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 324–348. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_12

[CGS14] Campbell, P., Groves, M., Shepherd, D.: Soliloquy: a cautionary tale. In: ETSI 2nd Quantum-Safe Crypto Workshop (2014). http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf

[CIV16a] Castryck, W., Iliashenko, I., Vercauteren, F.: On the tightness of the error bound in Ring-LWE. LMS J. Comput. Math. 130–145 (2016)

[CIV16b] Castryck, W., Iliashenko, I., Vercauteren, F.: Provably weak instances of ring-LWE revisited. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9665, pp. 147–167. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_6

[CLS17] Chen, H., Lauter, K., Stange, K.E.: Attacks on search RLWE. SIAM J. Appl. Algebra Geom. (SIAGA) **1**, 665–682 (2017)

[CLS16] Chen, H., Lauter, K., Stange, K.E.: Vulnerable Galois RLWE families and improved attacks. In: Proceedings of SAC. Springer (2016)

[Cona] Conrad, K.: The conductor ideal. http://www.math.uconn.edu/~kconrad/blurbs/gradnumthy/conductor.pdf

[Conb] Conrad, K.: The different ideal. http://www.math.uconn.edu/~kconrad/blurbs/gradnumthy/different.pdf

[Con95] Conway, J.B.: Functions of One Complex Variable. Springer, New York (1995). https://doi.org/10.1007/978-1-4612-6313-5

[DD12] Ducas, L., Durmus, A.: Ring-LWE in polynomial rings. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 34–51. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30057-8_3

[DLL+18] Ducas, L., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS - Dilithium: digital signatures from module lattices. In: TCHES (2018)

[EHL14] Eisenträger, K., Hallgren, S., Lauter, K.: Weak instances of PLWE. In: SAC (2014)

[ELOS15] Elias, Y., Lauter, K.E., Ozman, E., Stange, K.E.: Provably weak instances of ring-LWE. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 63–92. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47989-6_4

[GHPS12] Gentry, C., Halevi, S., Peikert, C., Smart, N.P.: Ring switching in BGV-style homomorphic encryption. In: Visconti, I., De Prisco, R. (eds.) SCN 2012. LNCS, vol. 7485, pp. 19–37. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32928-9_2

[GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC (2008)

[HHPW10] Hoffstein, J., Howgrave-Graham, N., Pipher, J., Whyte, W.: Practical lattice-based cryptography: NTRUEncrypt and NTRUSign. In: Nguyen, P., Vallée, B. (eds.) The LLL Algorithm. Information Security and Cryptography, pp. 349–390. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-02295-1_11

[LPR10] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. JACM **60**(6), 43 (2013)

[LPR13] Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 35–54. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_3

[LS15] Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. Des. Codes Cryptogr. **75**(3), 565–599 (2015)

[Lyu16] Lyubashevsky, V.: Digital signatures based on the hardness of ideal lattice problems in all rings. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10032, pp. 196–214. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_7

[MR04]   Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measure. In: Proceedings of FOCS, pp. 371–381. IEEE (2004)

[Pei16]   Peikert, C.: How (not) to instantiate ring-LWE. In: Zikas, V., De Prisco, R. (eds.) SCN 2016. LNCS, vol. 9841, pp. 411–430. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44618-9_22

[PR06]   Peikert, C., Rosen, A.: Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 145–166. Springer, Heidelberg (2006). https://doi.org/10.1007/11681878_8

[PR07]   Peikert, C., Rosen, A.: Lattices that admit logarithmic worst-case to average-case connection factors. In: STOC (2007)

[PRS17]   Peikert, C., Regev, O., Stephens-Davidowitz, N.: Pseudorandomness of Ring-LWE for any ring and modulus. In: STOC (2017)

[Reg05]   Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6), 1–40 (2009)

[RSSS17]   Roşca, M., Sakzad, A., Stehlé, D., Steinfeld, R.: Middle-product learning with errors. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10403, pp. 283–297. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_10

[SE94]   Schnorr, C.-P., Euchner, M.: Lattice basis reduction: improved practical algorithms and solving subset sum problems. Math. Program. **66**, 181–199 (1994)

[SS11]   Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 27–47. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20465-4_4

[SS13]   Stehlé, D., Steinfeld, R.: Making NTRUEncrypt and NTRUSign as secure standard worst-case problems over ideal lattices (2013). http://perso.ens-lyon.fr/damien.stehle/NTRU.html

[SSTX09]   Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public key encryption based on ideal lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 617–635. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_36

[Ste17]   Stevenhagen, P.: Lecture notes on number rings (2017). http://websites.math.leidenuniv.nl/algebra/ant.pdf