Algebraically Structured LWE, Revisited

Chris Peikert[*]   Zachary Pepin[†]

June 9, 2023

Abstract

In recent years, there has been a proliferation of *algebraically structured* Learning With Errors (LWE) variants, including Ring-LWE, Module-LWE, Polynomial-LWE, Order-LWE, and Middle-Product LWE, and a web of reductions to support their hardness, both among these problems themselves and from related worst-case problems on structured lattices. However, these reductions are often difficult to interpret and use, due to the complexity of their parameters and analysis, and most especially their (frequently large) blowup and distortion of the error distributions.

In this paper we unify and simplify this line of work. First, we give a general framework that encompasses *all* proposed LWE variants (over commutative base rings), and in particular unifies all prior "algebraic" LWE variants defined over number fields. We then use this framework to give much simpler, more general, and tighter reductions from Ring-LWE to other algebraic LWE variants, including Module-LWE, Order-LWE, and Middle-Product LWE. In particular, all of our reductions have easy-to-analyze and frequently small error expansion; in most cases they even leave the error unchanged. A main message of our work is that it is straightforward to use the hardness of the original Ring-LWE problem as a foundation for the hardness of all other algebraic LWE problems defined over number fields, via simple and rather tight reductions.

1 Introduction

1.1 Background

Regev's *Learning With Errors* (LWE) problem [Reg05] is a cornerstone of lattice-based cryptography, serving as the basis for countless cryptographic constructions (see, for example, the surveys [Reg10, Pei16a]). One primary attraction of LWE is that it can be supported by worst-case to average-case reductions from conjectured hard problems on general lattices [Reg05, Pei09, BLP+13, PRS17]. But while constructions based on LWE can have reasonably good asymptotic efficiency, they are often not as practically efficient as one might like, especially in terms of key and ciphertext sizes.

Inspired by the early NTRU cryptosystem [HPS98], Micciancio's initial worst-case to average-case reductions for "algebraically structured" lattices over polynomial rings [Mic02], and follow-on works [PR06, LM06], Lyubashevsky, Peikert, and Regev [LPR10] introduced *Ring-LWE* to improve the asymptotic and practical efficiency of LWE (see also the concurrent independent work [SSTX09], which considered a search variant of Ring-LWE). Ring-LWE is parameterized by the ring of integers in a number field, and [LPR10, PRS17] supported the hardness of Ring-LWE by a reduction from conjectured worst-case-hard problems on lattices corresponding to *ideals* in the ring.

Since then, several works have introduced and studied a host of other algebraically structured LWE variants—including Module-LWE [BGV12, LS15, AD17], Polynomial-LWE [SSTX09, RSW18], Order-LWE [BBPS19], and Middle-Product LWE [RSSS17]—relating them to each other and to various worst-case problems on structured lattices. Of particular interest is the work on Middle-Product LWE (MP-LWE) [RSSS17, RSW18], which, building on ideas from [Lyu16], gave a reduction from Ring- or Poly-LWE over a *huge class* of rings to a *single* MP-LWE problem. This means that breaking the MP-LWE problem in question is at least as hard as breaking *all* of large number of Ring-/Poly-LWE problems defined over unrelated rings.

Thanks to the above-described works, we now have a wide assortment of algebraic LWE problems to draw upon, and a thick web of reductions to support their hardness, at least for certain parameters. However, these reductions are often difficult to interpret and use due to the complexity of their parameters, and most especially their effect on the problems' *error distributions*. In particular, some reductions incur a substantial error blowup and distortion, which is often quite complicated to analyze, and bounded loosely by rather large factors. Some desirable reductions, like the one from Ring-LWE to MP-LWE, even require composing multiple hard-to-analyze steps. Finally, some of the reductions require non-uniform advice in the form of special short ring elements that in general do not seem easy to compute. See Figure 1 for a summary.

All this makes it rather challenging to navigate the state of the art, and especially to draw conclusions about which problems and parameters are supported by reductions and proofs. The importance of having clear, precise, and tight reductions is underscored by the fact that certain seemingly reasonable parameters of algebraic LWE problems have turned out to be insecure, but ultimately for prosaic reasons; see, e.g., [CIV16, Pei16b].

1.2 Contributions and Technical Overview

Here we give an overview of our contributions and how they compare to prior works. At a high level, we provide a general framework that encompasses all the previously mentioned LWE variants, and in particular unifies all prior "algebraic" LWE variants defined over number fields. We then use this framework to give much simpler, more general, and tighter reductions from Ring-LWE to other algebraic LWE variants, including Module-LWE, Order-LWE, and Middle-Product LWE. A main message of our work is that it is possible to use the hardness of Ring-LWE as a foundation for the hardness of all prior algebraic LWE problems (and some new ones), via fairly simple and easy-to-analyze reductions.

1.2.1 Generalized (Algebraic) LWE

In Section 3 we define new forms of LWE that unify and strictly generalize all previously mentioned ones.

Defining generalized LWE.   First, in Section 3.1 we describe a general framework that encompasses *all* the previously mentioned forms of LWE, including plain, Ring-, Module-, Poly-, Order-, and Middle-Product LWE (in both "dual" and "primal" forms, where applicable). The key observation is that in all such problems, the secret s, public multipliers a, and their (error-free) products $s \cdot a$ are vectors over the quotient $\mathcal{I}_{\mathcal{Q}} := \mathcal{I}/\mathcal{Q}\mathcal{I}$, where $\mathcal{I} = \mathcal{S}, \mathcal{A}, \mathcal{B} = \mathcal{S}\mathcal{A}$ (respectively) is a fixed *fractional ideal* of an order $\mathcal{O}$ of a number field K, and the "modulus" $\mathcal{Q}$ is some fixed (integral) $\mathcal{O}$-ideal.[1] Moreover, the products $s \cdot a$ are given by some fixed $\mathcal{O}_{\mathcal{Q}}$-*bilinear map* on s and a. By fixing bases, this map can be represented as an *order-three tensor* (i.e., a three-dimensional array) over $\mathcal{O}_{\mathcal{Q}}$, where evaluating the map corresponds to multiplying the tensor by the vectors representing s and a along the corresponding dimensions.

A generalized LWE problem is defined by some fixed choices of the above parameters (order, ideals, and tensor), along with an error distribution. In such a problem, there is a random "secret" s over $\mathcal{S}_{\mathcal{Q}}$, and an instance consists of independent "noisy random products"

$$(a_i \,,\, b = s \cdot a_i + e_i \bmod \mathcal{Q}\mathcal{B}),$$

where each a_i is uniformly random over $\mathcal{A}_{\mathcal{Q}}$ and each e_i is drawn from the error distribution.[2] For example, plain LWE uses n-dimensional vectors over order $\mathcal{O} = \mathcal{S} = \mathcal{A} = \mathbb{Z}$, modulus $\mathcal{Q} = q\mathbb{Z}$ for some integer q, and the vector inner product as the bilinear map, which corresponds to the $n \times n \times 1$ identity-matrix tensor. Ring-LWE uses the ring of integers $\mathcal{O} = \mathcal{O}_K$ of a number field K as its order, with $\mathcal{A} = \mathcal{O}$, $\mathcal{S} = \mathcal{O}^\vee$ (the dual lattice of $\mathcal{O}$, known as the "codifferent" fractional ideal), modulus $\mathcal{Q} = q\mathcal{O}$ (or more generally, any $\mathcal{O}$-ideal), and multiplication in K as the bilinear map, which corresponds to the $1 \times 1 \times 1$ scalar-unity tensor.

We show how Middle-Product LWE also straightforwardly fits into this framework. Interestingly, by a judicious choice of bases, the two-dimensional matrix "slices" $M_{i..}$ of the middle-product tensor M turn out to form the standard basis for the space of all *Hankel* matrices. (In a Hankel matrix, the (j, ℓ)th entry is determined by $j + \ell$.) This formulation is central to our improved reduction from Ring-LWE over a wide class of number fields to Middle-Product LWE, described further in Section 1.2.3 below.[3]

Parameterizing by a single lattice.   Next, in Section 3.2 we define a specialization of generalized LWE that still encompasses all prior "algebraic" LWE variants defined *over number fields*, including Ring-, Module-, Poly-, and Order-LWE. A member $\mathcal{L}$-LWE of this class of problems is parameterized by *any (full-rank) lattice* $\mathcal{L}$ of a number field K, i.e., a free additive subgroup whose $\mathbb{Q}$-span equals K. Its *coefficient ring*

$$\mathcal{O}^{\mathcal{L}} := \{x \in K : x\mathcal{L} \subseteq \mathcal{L}\} \tag{1.1}$$

is the set of field elements by which $\mathcal{L}$ is closed under multiplication. Letting $\mathcal{L}^\vee = \{x \in K : \mathrm{Tr}_{K/\mathbb{Q}}(x\mathcal{L}) \subseteq \mathbb{Z}\}$ denote the *dual lattice* of $\mathcal{L}$, it turns out that $\mathcal{O}^{\mathcal{L}} = (\mathcal{L} \cdot \mathcal{L}^\vee)^\vee$, and it is an order of K; indeed, it is the

---

[1]An *order* $\mathcal{O}$ of K is a subring (with unity) that is also the $\mathbb{Z}$-span of some $\mathbb{Q}$-basis of K; the latter condition is equivalent to being a full-rank lattice in K. An (integral) $\mathcal{O}$-*ideal* is an additive subgroup $\mathcal{I} \subseteq \mathcal{O}$ for which $\mathcal{O}\mathcal{I} = \mathcal{I}$. A *fractional $\mathcal{O}$-ideal* is a set $\mathcal{I} \subseteq K$ for which $d\mathcal{I}$ is an $\mathcal{O}$-ideal for some $d \in \mathcal{O}$. In this work, all ideals are implicitly required to be *nonzero*, i.e., not $\{0\}$.

[2]Observe that the reduction modulo $\mathcal{Q}\mathcal{B}$ is well defined because the (error-free) product $s \cdot a_i$ is over $\mathcal{B}_{\mathcal{Q}}$.

[3]Toeplitz matrices, which are closely related to Hankel matrices, also play an important role in the original reduction from (primal) Poly-LWE to MP-LWE [RSSS17]. Viewing that reduction in our framework (specifically, Theorem 4.8), we observe that the *primal* Poly-LWE tensor typically does *not* have Hankel slices. But at the cost of some relative error growth and distortion (represented by $\mathbf{B}$ in our theorem, and $\mathbf{M}_f$ in [RSSS17]), the tensor's output basis can be transformed to a dual one, which does induce Hankel slices.

maximal order for which $\mathcal{L}$ is a fractional ideal. Notice that $\mathcal{L}$ and $\mathcal{L}^\vee$ have the same coefficient ring; in particular, if $\mathcal{L}$ is an order $\mathcal{O}$ of K or its dual $\mathcal{O}^\vee$, then $\mathcal{O}^\mathcal{L} = \mathcal{O}$.[4]

An $\mathcal{L}$-LWE^k problem is simply the generalized LWE problem, specialized to order $\mathcal{O} = \mathcal{O}^\mathcal{L}$ and fractional $\mathcal{O}$-ideals $\mathcal{A} = \mathcal{O}$, $\mathcal{S} = \mathcal{B} = \mathcal{L}^\vee$, with s and a as k-dimensional vectors over $\mathcal{S}_\mathcal{Q}$ and $\mathcal{A}_\mathcal{Q}$ (respectively) and $s \cdot a$ denoting their inner product; when $k = 1$, we usually omit the superscript. So, an instance has a secret s over $\mathcal{L}_\mathcal{Q}^\vee$, and consists of independent noisy random products

$$(a_i \, , \, b = s \cdot a_i + e_i \bmod \mathcal{Q}\mathcal{L}^\vee),$$

where each a_i is uniformly random over $\mathcal{O}_\mathcal{Q}^\mathcal{L}$ and each e_i is drawn from the error distribution. We show (see Remark 4.6) that under mild conditions, taking the multipliers a_i to be over $\mathcal{O}_\mathcal{Q}$ is without loss of generality, which justifies the focus on this specific formulation of $\mathcal{L}$-LWE.

We now explain how $\mathcal{L}$-LWE generalizes prior algebraic LWE problems over number fields. As already noted, when $\mathcal{L} = \mathcal{O}$ or $\mathcal{L} = \mathcal{O}^\vee$ for an order $\mathcal{O}$ of K, we have $\mathcal{O}^\mathcal{L} = \mathcal{O}$. So, $\mathcal{L}$-LWE specializes to the following problems (and similarly, $\mathcal{L}$-LWE^k specializes to the "Module" variants):

1. Ring-LWE [LPR10] when $\mathcal{L} = \mathcal{O}_K$ is the full ring of integers of K;

2. Poly-LWE [RSW18] when $\mathcal{L} = \mathbb{Z}[\alpha]^\vee$ for some $\alpha \in \mathcal{O}_K$;

3. Order-LWE [BBPS19] (or its "dual" form) when $\mathcal{L} = \mathcal{O}$, for some arbitrary order $\mathcal{O}$ of K.[5]

Notice that in the second case, $\mathcal{L}$ is the *dual* $\mathcal{O}^\vee$ of its coefficient ring $\mathcal{O} = \mathcal{O}^\mathcal{L}$, so both the secret s and product $s \cdot a$ are over $\mathcal{L}^\vee = \mathcal{O}$ itself (modulo $\mathcal{Q}$). But as we shall see, for reductions it turns out to be more natural and advantageous to let $\mathcal{L}$ be an order, not the dual of an order. Furthermore, $\mathcal{L}$-LWE also captures other cases that are not covered by the ones above, namely, those for which $\mathcal{L}$ is neither an order nor its dual. (For $\mathcal{L}$-LWE to be well defined, we just need the $\mathcal{O}^\mathcal{L}$-module structure of $\mathcal{L}^\vee$, not any ring structure.)

As mentioned above, $\mathcal{L}$-LWE is also parameterized by an error distribution. For consistency across problems and with prior work, and without loss of generality, we always define and view the error distribution in terms of the *canonical embedding* of K. For concreteness, and following worst-case hardness theorems for Ring-LWE [LPR10, PRS17], the reader can keep in mind a spherical Gaussian distribution of sufficiently large width $r = \omega(\sqrt{\log n})$ over the canonical embedding, where $n = \deg(K/\mathbb{Q})$. While this differs syntactically from the kind of distribution often considered for Poly-LWE—namely, a spherical Gaussian over the *coefficient vector* of the error polynomial—the two views are interchangeable via some fixed linear transformation. For Gaussians, this transformation just changes the covariance, and if desired we can also add some independent compensating error to obtain a spherical Gaussian. However, our results demonstrate some advantages of working only with spherical Gaussians in the canonical embedding, even for Poly-LWE.

Reductions. In Section 4 we give a modular collection of tight reductions between various parameterizations of generalized LWE. Essentially, each reduction transforms samples of one LWE instantiation (for an unknown secret) to samples of another instantiation (for a related secret), and has the primary effect of changing either the ideals $\mathcal{S}$ and $\mathcal{A}$, the order $\mathcal{O}$ (or the lattice $\mathcal{L}$ defining it, in the case of $\mathcal{L}$-LWE), the tensor T defining

---

[4]We caution that $\mathcal{O}^\mathcal{L}$ is not "monotonic" in $\mathcal{L}$ under set inclusion, i.e., $\mathcal{L} \subseteq \mathcal{L}'$ does not imply any subset relationship between $\mathcal{O}^\mathcal{L}$ and $\mathcal{O}^{\mathcal{L}'}$, in either direction. In particular, $\mathcal{L}'$ and $c\mathcal{L}'$ have the same coefficient ring for any integer $c > 1$, but there can exist $\mathcal{L}$ having a different coefficient ring where $c\mathcal{L}' \subsetneq \mathcal{L} \subsetneq \mathcal{L}'$.

[5]In "dual" Order-LWE as defined in [BBPS19], the domains of the secret s and the multipliers a_i are merely swapped relative to our definition of $\mathcal{O}$-LWE. However, this is merely a syntactic difference, since the problems are tightly equivalent (with no change to the error distribution), under a mild condition that is also needed in other reductions; see [BBPS19, Remark 3.5] or Corollary 4.3.
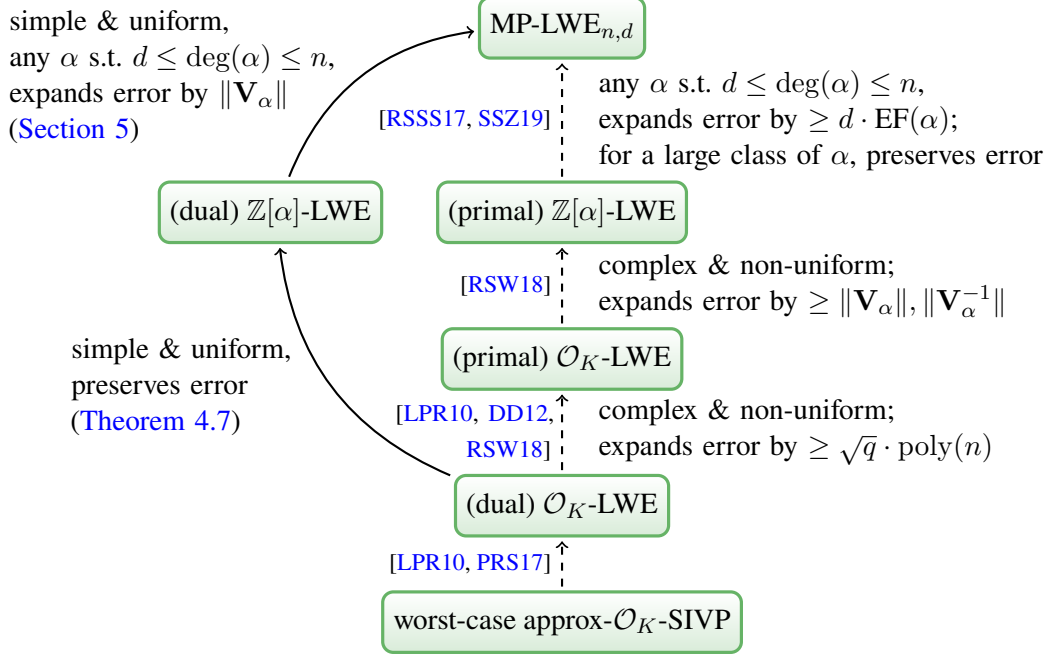
Figure 1: Summary of some known reductions among algebraic lattice and LWE problems. Dashed arrows represent prior reductions, and solid arrows represent some of the reductions given in this work.

the bilinear map, or the number field over which all the other parameters are defined. All of the reductions preserve the number of samples, almost all of them preserve the error distribution, and most of them are even tight polynomial-time *equivalences* (i.e., reductions in both directions). In cases where the error distribution is not preserved, it is changed by an easy-to-analyze linear transformation.

All of the main results of this work are then obtained by invoking or composing the above-described primitive reductions in various ways. We next give an overview of three main theorems that we obtain in this way; it seems likely that other interesting and useful results can be established as well.

1.2.2 Reduction Between $\mathcal{L}$-LWE Instantiations

As a first main result (see Theorem 4.7), we obtain a reduction from $\mathcal{L}'$-LWE to $\mathcal{L}$-LWE for any lattices $\mathcal{L} \subseteq \mathcal{L}'$ of K for which $\mathcal{O} := \mathcal{O}^{\mathcal{L}} \subseteq \mathcal{O}' := \mathcal{O}^{\mathcal{L}'}$, and the *conductor* $\mathcal{O}$-ideal $(\mathcal{L} : \mathcal{L}')$ (see Definition 2.1 and Lemma 2.8) is coprime with the modulus $\mathcal{Q}$ of the target problem; the source problem uses modulus $\mathcal{Q}' = \mathcal{O}'\mathcal{Q}$. Importantly, and unlike prior reductions of a similar flavor, our reduction preserves the error distribution. In particular, it yields a reduction from Ring-LWE to Order-LWE, by taking $\mathcal{L}' = \mathcal{O}_K$ to be the full ring of integers of a number field K, and $\mathcal{L} = \mathcal{O}$ to be another order of K; see below for further details.

We stress that the only "loss" associated with the reduction, which seems inherently necessary, is that when $\mathcal{L} \neq \mathcal{L}'$, the lattice $\mathcal{Q}\mathcal{L}^{\vee}$ by which the resulting noisy products $b \approx s \cdot a$ are reduced is typically "denser" than the lattice $\mathcal{Q}'(\mathcal{L}')^{\vee} \subsetneq \mathcal{Q}\mathcal{L}^{\vee}$ by which the original noisy products $b' \approx s' \cdot a'$ are reduced. One can alternatively see this as the (unchanging) error distribution being "wider" relative to the target lattice than to the original one. This can have consequences for applications, where we typically need the error from some (possibly combined) samples to be decodable modulo $\mathcal{Q}\mathcal{L}^{\vee}$. That is, we need to be able to efficiently recover e (or at least a large portion of it) from the coset $e + \mathcal{Q}\mathcal{L}^{\vee}$; to do this, standard decoding algorithms require

sufficiently short elements of the lattice $(\mathcal{Q}\mathcal{L}^\vee)^\vee = (\mathcal{L} : \mathcal{Q})$. (When $\mathcal{Q}$ is invertible as a fractional $\mathcal{O}$-ideal, this lattice is simply $\mathcal{Q}^{-1}\mathcal{L}$; in particular, when $\mathcal{Q} = q\mathcal{O}$ is principal, this lattice is $q^{-1}\mathcal{L}$.) So, the "sparser" we take $\mathcal{L} \subseteq \mathcal{L}'$ to be, the sparser we need $\mathcal{Q}$ to be in order to compensate. This weakens both the theoretical guarantees and concrete hardness of the original $\mathcal{L}'$-LWE problem, and is reason to prefer denser $\mathcal{L}$.

Discussion and comparison to prior (and subsequent) work. We now describe some of the immediate implications of the above reduction, and compare to related ones. Take $\mathcal{L}' = \mathcal{O}' = \mathcal{O}_K$ to be the full ring of integers of K, which corresponds to Ring-LWE, for which we have worst-case hardness theorems [LPR10, PRS17]. Then these same hardness guarantees are immediately inherited by Order-LWE (and in particular, Poly-LWE in its "dual" form) by taking $\mathcal{L} = \mathcal{O}$ to be an arbitrary order of K, as long as the conductor $\mathcal{O}$-ideal $(\mathcal{O} : \mathcal{O}')$ is coprime with $\mathcal{Q}$. These guarantees are qualitatively similar to the ones established in [RSW18, BBPS19, BBS21], but are obtained in a much simpler and more straightforward way; in particular, we do not need to replicate all the technical machinery of the worst-case to average-case reductions from [LPR10, PRS17] for arbitrary orders $\mathcal{O}$, as was done in [BBPS19, BBS21].[6]

Our reduction can also yield hardness for Poly-LWE—in which the secret s, multipliers a_i, and (error-free) products $s \cdot a_i$ all are over an order $\mathcal{O}$ itself, not its dual—via a different choice of $\mathcal{L}$ (see the next paragraph). However, it is instructive to see why it is preferable to use problems in which the products are over the *dual* $\mathcal{O}^\vee$ of the order. The main reason is that these admit quite natural reductions, both *from* Ring-LWE and *to* Middle-Product LWE and Module-LWE, whose effects on the error distribution are easy to understand and bound entirely in terms of certain known short elements of $\mathcal{O}$; see Sections 1.2.3 and 1.2.4 for details.

By contrast, the reduction and analysis for (primal) Poly-LWE over $\mathcal{O} = \mathbb{Z}[\alpha]$ (as in [RSW18]) is much more complex and cumbersome. Because $\mathcal{O}^\vee \not\subseteq \mathcal{O}_K$ (except in the trivial case $K = \mathbb{Q}$), we cannot simply take $\mathcal{L} = \mathcal{O}^\vee$. Instead, we need to apply a suitable "tweak" factor $t \in K$, so that $\mathcal{L} = t\mathcal{O}^\vee \subseteq \mathcal{O}_K$ and hence $\mathcal{L}^\vee = t^{-1}\mathcal{O}$. Reducing to $\mathcal{L}$-LWE preserves the error distribution, but to finally convert the samples to Poly-LWE samples we need to multiply by t, which distorts the error distribution. It can be shown that t must lie in the product of the *different ideal* of $\mathcal{O}_K$ and the *conductor ideal* of $\mathcal{O}$ (among other constraints), so the reduction requires non-uniform advice in the form of such a "short" t that does not distort the error too much. The existence proof for such a t from [RSW18] is quite involved, requiring several pages of sophisticated analysis. Finally, the decodability of the (distorted) error modulo $\mathcal{Q}\mathcal{O}$ is mainly determined by short nonzero vectors in $(\mathcal{Q}\mathcal{O})^\vee = (\mathcal{O}^\vee : \mathcal{Q})$, which also must be found and analyzed. (All these issues arise under slightly different guises in [RSW18]; in fact, there the error is distorted by t^2, yielding an even looser reduction.)

1.2.3 Reduction from Order-LWE to Middle-Product LWE

In Section 5 we give a simple reduction from $\mathcal{O}$-LWE, for a *wide class* of number fields K and orders $\mathcal{O}$ including polynomial rings of the form $\mathcal{O} = \mathbb{Z}[\alpha] \cong \mathbb{Z}[x]/f(x)$, to a *single* Middle-Product LWE problem. Together with the $\mathcal{L}$-LWE reduction described above, this yields a Ring-LWE to MP-LWE reduction like the

---

[6]Formally, the worst-case hardness theorems obtained from our results (with [LPR10, PRS17]) versus [BBPS19] are *incomparable* and complementary: ours stem from (worst-case) ideal lattices corresponding to $\mathcal{O}_K$-ideals, whereas those of [BBPS19] stem from invertible $\mathcal{O}$-ideals. Because a number-field lattice can be invertible as an ideal of at most one order, the two classes of ideals are disjoint (when $\mathcal{O} \neq \mathcal{O}_K$). Subsequent work [BBS21] gave a worst-case hardness theorem for modulus $\mathcal{Q} = q\mathcal{O}, q \in \mathbb{Z}$ stemming from *all* $\mathcal{O}$-ideals, under an assumption that is equivalent to $(\mathcal{O} : \mathcal{O}_K) + \mathcal{Q} = \mathcal{O}$. This class of ideals is a superset of both the $\mathcal{O}_K$-ideals and the invertible $\mathcal{O}$-ideals, so the result from [BBS21] generalizes both ours and the one from [BBPS19]. We mention that the condition $(\mathcal{O} : \mathcal{O}') + \mathcal{Q} = \mathcal{O}$ implies that *all* fractional $\mathcal{O}$-ideals are *invertible modulo* $\mathcal{Q}$ (see Section 2.4.2), which gives a simpler proof of the general "Cancellation Lemma" from [BBS21] and also enables our Theorems 4.5 and 4.7.

one obtained by the combination of [RSSS17, RSW18], which says that breaking the MP-LWE problem in question is at least as hard as breaking *all* of a wide class of Ring-LWE problems over unrelated number fields.

However, our result subsumes the prior ones by being simpler, more general, and tighter; see Figure 1 for a comparison. More specifically, it drops certain technical conditions on the order, and the overall distortion in the error distribution (starting from Ring-LWE) is given entirely by the spectral norm $\|\vec{p}\|$ of a certain known power basis $\vec{p}$ of $\mathcal{O}$. In particular, spherical Gaussian error over the canonical embedding of $\mathcal{O}$ translates to spherical Gaussian MP-LWE error (over the reals) that is just a $\|\vec{p}\|$ factor wider. These advantages arise from the error-preserving nature of our $\mathcal{L}$-LWE reduction (described above), and the judicious use of dual lattices in the defintion of $\mathcal{O}$-LWE.[7]

At heart, what makes our reduction work is the hypothesis that the order $\mathcal{O}$ has a power basis $\vec{p} = (x^i)$ for some $x \in \mathcal{O}$; clearly any monogenic order $\mathcal{O} = \mathbb{Z}[\alpha]$ has such a basis, with $x = \alpha$. Using our generalized LWE framework, we show that when using a power basis $\vec{p}$ and its dual $\vec{p}^\vee$ for $\mathcal{O}$ and $\mathcal{O}^\vee$ respectively, all the "slices" $T_{i\cdot\cdot}$ of the tensor T representing multiplication $\mathcal{O}^\vee \times \mathcal{O} \to \mathcal{O}^\vee$ are *Hankel* matrices. So, using the fact that the slices $M_{i\cdot\cdot}$ of the middle-product tensor M form the standard basis for the space of all Hankel matrices, we can transform $\mathcal{O}$-LWE samples to MP-LWE samples. The resulting MP-LWE error distribution is simply the original error distribution represented in the $\vec{p}^\vee$ basis, which is easily characterized using the geometry of $\vec{p}$.

The above perspective is helpful for revealing other reductions from wide classes of LWE problems to a single LWE problem. Essentially, it suffices that all the slices $T_{i\cdot\cdot}$ of all the source-problem tensors T over a ring $\mathcal{O}_Q$ lie in the $\mathcal{O}_Q$-span of the slices of the target-problem tensor. We use this observation in our final main reduction, described next.

1.2.4 Reduction Between Module-LWE Instantiations

Lastly, in Section 6 we give a reduction establishing the hardness of Module-LWE over an order $\mathcal{O}$ of a number field K, based on the hardness of Module-LWE (or Ring-LWE, as a special case) over *any one of a wide class* of orders $\mathcal{O}'$ of number field extensions K'/K. This is qualitatively analogous to what is known for Middle-Product LWE, but is potentially more beneficial because Module-LWE is easier to use in applications, and is much more widely used in theory and in practice.

A bit more precisely, we give a simple reduction from $\mathcal{O}'$-LWE$^{k'}$, for a *wide class* of orders $\mathcal{O}'$, to a *single* $\mathcal{O}$-LWE^k problem. The only needed condition is that $\mathcal{O}'$ is a rank-(k/k') free $\mathcal{O}$-module; therefore, the "total rank" (over $\mathbb{Z}$) is preserved.[8] This condition is easily achieved, for example, by defining $\mathcal{O}' = \mathcal{O}[\alpha] \cong \mathcal{O}[x]/f(x)$ for some root α of an arbitrary degree-(k/k') monic irreducible polynomial $f(x) \in \mathcal{O}[x]$. Once again, due to the use of duality in the definition of the problems, the reduction's effect on the error distribution is very easy to characterize: the output error is simply the trace (from K' to K) of the input error. In particular, the typical example of spherical Gaussian error in the canonical embedding of K' maps to spherical Gaussian error in the canonical embedding of K, because the trace just sums over a certain equi-partition of the coordinates.

We point out that our result is reminiscent of, but formally incomparable to, the kind of worst-case hardness theorem for $\mathcal{O}$-LWE^k (for certain $\mathcal{O}$) given in [LS15]: there the source problem involves arbitrary (worst-case) rank-k *module lattices* over $\mathcal{O}$, whereas here our source problem is an average-case rank-k' LWE

---

[7] We note that an independent and concurrent work [SSZ19] gave a reduction from (primal) Poly-LWE to MP-LWE that preserves the error distribution for a wide range of defining polynomials. This can be used to somewhat reduce the error expansion in the prior chain of reductions from Ring-LWE to MP-LWE [RSSS17, RSW18]. However, the earlier steps in the chain still incur a larger error expansion than ours does; see Figure 1 for details.

[8] We stress that K' can be any extension field of K. Proving that one specific Ring-/Module-LWE instantiation is at least as hard as many others without such a restriction remains a very interesting open problem.

problem over a rank-(k/k') $\mathcal{O}$-module. Our result is also somewhat complementary to a reduction of [AD17] (generalizing [BLP^{+}13]) that can *reduce* the rank of Module-LWE problems, with a corresponding increase in the modulus size and error rate (relative to the modulus). Specifically, the reduction is from Module-LWE of rank k (over any ring of integers in a number field) and modulus q to Module-LWE of rank k/ℓ (over the same ring) and modulus q^{ℓ}, where the error rate grows by roughly the square root of the total rank of the source problem.

1.3 Differences From Conference Version

This work is the full version of a preliminary conference paper [PP19]. Compared to the preliminary version, the present one has the following main additions and improvements:

- We give a formal definition of the *generalized LWE* problem (see Section 3.1), which encompasses all previously proposed algebraic LWE problems over polynomials or number fields, and in particular allows for an arbitrary modulus $\mathcal{O}$-ideal $\mathcal{Q}$, not just an integer q.

- Several results from the conference version involved hypotheses of the form $\gcd(|\mathcal{L}'/\mathcal{L}|, q) = 1$ for relevant lattices $\mathcal{L} \subseteq \mathcal{L}'$, whereas the present version asks that $(\mathcal{L} : \mathcal{L}') + \mathcal{Q} = \mathcal{O}$. It turns out that for $\mathcal{Q} = q\mathcal{O}$, the former hypothesis immediately implies the latter one (because $|\mathcal{L}'/\mathcal{L}| \in (\mathcal{L} : \mathcal{L}')$ by Lagrange's Theorem), so the present results are strictly more general, with no stronger hypotheses than the previous versions.[9]

- We give a more general and modular collection of reductions between instantiations of generalized LWE (Section 4), and have reworked the hardness theorems for Middle-Product LWE (Section 5) and Module-LWE (Section 6) as direct consequences of these reductions.

2 Preliminaries

In this work, by "ring" we always mean a commutative ring with identity. We let $[n] = \{0, \ldots, n-1\}$ for a non-negative integer n.

Computational issues. All of the computational operations needed by the reductions in this work—e.g., computing the sum, intersection, or product of lattices; evaluating a group homomorphism, and inverting it when it is injective; computing the kernel of a group homomorphism, etc.—can be performed in polynomial time (given suitable representations) using generic algorithms on finitely generated (and in particular, finite) abelian groups. These mainly rely on algorithms for computing the Smith and Hermite normal forms; see, e.g., [Coh99, Chapter 4] and [CDO01] for details. (More efficient algorithms are frequently available for important special cases, like cyclotomic fields or full rings of integers.) For simplicity of exposition, we omit explicit algorithmic descriptions in favor of abstract mathematical notation, and typically treat the inputs (groups, homomorphisms, etc.) as implicit by context, except where they may not be obvious.

2.1 Vectors, Matrices, and Tensors

In this work we frequently work with tensors, which generalize vectors and matrices to higher dimensions. Formally, a tensor T over a base set S has a finite index set I and a value $T_i \in S$ for each $i \in I$. If $I = I_1 \times \cdots \times I_r$ is seen as the Cartesian product of r components, we say that T has *order* r, and index

---

[9]In fact, the two hypotheses can be shown to be *equivalent* for $\mathcal{Q} = q\mathcal{O}$, by a version of the Jordan-Hölder theorem.

it as $T_{i_1 i_2 \cdots i_r}$. (A tensor's order may vary depending on how we choose to factor I.) Vectors are merely order-one tensors, which we denote by lower-case letters in bold, like $\mathbf{a}$, or with arrows, like $\vec{a}$, depending on the base set. Matrices are order-two tensors, which we denote by upper-case bold letters, like $\mathbf{A}$. We denote higher-order tensors by ordinary upper-case letters, like T.

For tensors A, B over a common set S supporting multiplication, and having respective index sets I, J, their *Kronecker product* (also known as *tensor product*) $A \otimes B$ is the tensor having index set $I \times J$ whose (i, j)th entry is $(A \otimes B)_{ij} = A_i B_j$. In general, then, the order of $A \otimes B$ is the sum of the orders of A and B. However, when A and B have the same order r with $I = I_1 \times \cdots \times I_r$ and $J = J_1 \times \cdots \times J_r$, we can also treat $A \otimes B$ as an order-r tensor as well, by reindexing it to have index set $K_1 \times \cdots \times K_r$ where $K_i = I_i \times J_i$, and $((i_1, j_1), \ldots, (i_r, j_r))$th entry $A_{i_1 \cdots i_r} B_{j_1 \cdots j_r}$. Moreover, a product index set $[n] \times [n']$ for non-negative integers n, n' can be reindexed as $[n \cdot n']$ (and vice versa), where (i, j) corresponds to $in' + j$.

2.2 Number Fields, Lattices, and Duality

For the remainder of this section we recall the key concepts from algebraic number theory that will be used in this work; for further details and proofs, see any standard text, e.g., [Lan94].

An (algebraic) *number field* K is a finite-dimensional field extension of the rationals $\mathbb{Q}$. More concretely, it can be written as $K = \mathbb{Q}(\zeta)$, by adjoining to $\mathbb{Q}$ some element ζ that satisfies the relation $f(\zeta) = 0$ for some monic irreducible polynomial $f(x) \in \mathbb{Q}[x]$. The polynomial f is called the *minimal polynomial* of ζ, and the degree of f is called the *degree* of K, which is denoted by n in what follows.

The (field) *trace* $\mathrm{Tr} = \mathrm{Tr}_{K/\mathbb{Q}} \colon K \to \mathbb{Q}$ is the trace of the $\mathbb{Q}$-linear transformation on K (viewed as a vector space over $\mathbb{Q}$) representing multiplication by x. More concretely, fixing any $\mathbb{Q}$-basis of K lets us uniquely represent every element of K as a vector in $\mathbb{Q}^n$, and multiplication by any $x \in K$ corresponds to multiplication by a matrix $M_x \in \mathbb{Q}^{n \times n}$; the trace of x is the trace of this matrix.

For the purposes of this work, a *lattice* $\mathcal{L}$ in K is a full-rank free additive subgroup of K, i.e., a free $\mathbb{Z}$-module of rank n. Equivalently, any lattice is generated as the integer linear combinations of some n $\mathbb{Q}$-linearly independent *basis* elements $\vec{b} = (b_1, \ldots, b_n) \in K^n$, as $\mathcal{L} = \{\sum_{i=1}^n \mathbb{Z} \cdot b_i\}$; the choice of basis is not unique.

For any two lattices $\mathcal{L}, \mathcal{L}' \subset K$, their product $\mathcal{L} \cdot \mathcal{L}'$ is the set of all integer linear combinations of terms $x \cdot x'$ for $x \in \mathcal{L}, x' \in \mathcal{L}'$. This set is itself a lattice, because it is torsion free and finitely generated by the tensor product of any basis of $\mathcal{L}$ and of $\mathcal{L}'$.

For a lattice $\mathcal{L}$, its *dual lattice* $\mathcal{L}^\vee$ (which is indeed a lattice) is defined as

$$\mathcal{L}^\vee := \{x \in K : \mathrm{Tr}(x\mathcal{L}) \subseteq \mathbb{Z}\}.$$

It is easy to see that if $\mathcal{L} \subseteq \mathcal{L}'$ are lattices in K, then $(\mathcal{L}')^\vee \subseteq \mathcal{L}^\vee$, and if $\vec{b}$ is a basis of $\mathcal{L}$, then its *dual basis* $\vec{b}^\vee = (b_1^\vee, \ldots, b_n^\vee)$ is a basis of $\mathcal{L}^\vee$, where $\vec{b}^\vee$ is defined so that $\mathrm{Tr}(b_i \cdot b_j^\vee)$ is 1 when $i = j$, and is 0 otherwise. Observe that by definition, $x = \vec{b}^t \cdot \mathrm{Tr}(\vec{b}^\vee \cdot x)$ for every $x \in K$.

2.3 Orders and Ideals

An *order* $\mathcal{O}$ of a number field K is a lattice in K that is also a subring, i.e., $1 \in \mathcal{O}$ and $\mathcal{O}$ is closed under multiplication. An element $\alpha \in K$ is an *algebraic integer* if there exists a monic integer polynomial $g(X) \in \mathbb{Z}[X]$ such that $g(\alpha) = 0$. The set of algebraic integers in K, denoted $\mathcal{O}_K$, is called the *ring of integers* of K, and is its maximal order, i.e., every order $\mathcal{O}$ of K is a subset of $\mathcal{O}_K$. For any order $\mathcal{O}$ of K, we have $\mathcal{O} \cdot \mathcal{O}^\vee = \mathcal{O}^\vee$ because $\mathcal{O}^\vee = 1 \cdot \mathcal{O}^\vee \subseteq \mathcal{O} \cdot \mathcal{O}^\vee$ and $\mathrm{Tr}((\mathcal{O} \cdot \mathcal{O}^\vee) \cdot \mathcal{O}) = \mathrm{Tr}(\mathcal{O}^\vee \cdot \mathcal{O}) \subseteq \mathbb{Z}$, since $\mathcal{O} \cdot \mathcal{O} = \mathcal{O}$.

9

An *ideal* of an order $\mathcal{O}$, also called an $\mathcal{O}$-ideal, is an additive subgroup $\mathcal{I} \subseteq \mathcal{O}$ that is closed under multiplication by $\mathcal{O}$, i.e., $\mathcal{O}\mathcal{I} \subseteq \mathcal{I}$; in fact, this is an equality, since $1 \in \mathcal{O}$. Throughout this work, we always implicitly restrict ideals to be *nontrivial* subgroups, in order to rule out the inconvenient "zero ideal" $\mathcal{I} = \{0\}$. With this restriction, every ideal $\mathcal{I}$ is a (full-rank) sublattice of $\mathcal{O}$.

A proper ideal $\mathfrak{p} \subsetneq \mathcal{O}$ is *maximal* if there does not exist any $\mathcal{O}$-ideal $\mathcal{I}$ strictly between $\mathfrak{p}$ and $\mathcal{O}$, i.e., $\mathfrak{p} \subsetneq \mathcal{I} \subsetneq \mathcal{O}$; it is *prime* if for every $a, b \in \mathcal{O}$ for which $ab \in \mathfrak{p}$, either $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$ (or both). It turns out that in any order of a number field, an ideal is prime if and only if it is maximal. Two $\mathcal{O}$-ideals $\mathcal{I}$ and $\mathcal{J}$ are *coprime* (also known as *comaximal*) if $\mathcal{I} + \mathcal{J} = \mathcal{O}$. Finally, a *fractional* ideal of an order $\mathcal{O}$ is a set $\mathcal{I} \subset K$ for which there exists a $d \in \mathcal{O}$ such that $d\mathcal{I}$ is an $\mathcal{O}$-ideal; in particular, $\mathcal{I}$ is a lattice.

A fractional $\mathcal{O}$-ideal $\mathcal{I}$ is *invertible* if there exists a fractional $\mathcal{O}$-ideal $\mathcal{I}^{-1}$ for which $\mathcal{I}\mathcal{I}^{-1} = \mathcal{O}$. Such an $\mathcal{I}^{-1}$, which is unique, is called the *inverse* of $\mathcal{I}$. Every fractional $\mathcal{O}_K$-ideal is invertible, but every non-maximal order $\mathcal{O} \subsetneq \mathcal{O}_K$ has some non-invertible ideal (see the next subsection for further details).

2.4 Lattice Quotients

Here we introduce a very useful notion called the *quotient* of two number-field lattices, and recall how several important objects can be obtained from it.[10]

Definition 2.1 (Lattice quotient). For lattices $\mathcal{L}, \mathcal{L}'$ in K, their *quotient* is $(\mathcal{L} : \mathcal{L}') := \{x \in K : x\mathcal{L}' \subseteq \mathcal{L}\}$.

The above can be seen as a kind of quotient because $\mathcal{I}\mathcal{L}' \subseteq \mathcal{L}$ if and only if $\mathcal{I} \subseteq (\mathcal{L} : \mathcal{L}')$. For any lattices $\mathcal{L}, \mathcal{L}', \mathcal{L}''$ in K, it follows immediately from the definition that

$$(\mathcal{L} : \mathcal{L}')(\mathcal{L}' : \mathcal{L}'') \subseteq (\mathcal{L} : \mathcal{L}''). \tag{2.1}$$

In addition, the lattice quotient generalizes the notion of invertibility for ideals.

Lemma 2.2. *Let $\mathcal{O}$ be an order of K, and let $\mathcal{I}, \mathcal{I}'$ be a fractional $\mathcal{O}$-ideals, with $\mathcal{I}'$ invertible. Then $(\mathcal{I} : \mathcal{I}') = \mathcal{I}(\mathcal{I}')^{-1}$. In particular, $(\mathcal{O} : \mathcal{I}') = (\mathcal{I}')^{-1}$.*

Proof. We have $\mathcal{I}(\mathcal{I}')^{-1} \cdot \mathcal{I}' = \mathcal{I}$, so $\mathcal{I}(\mathcal{I}')^{-1} \subseteq (\mathcal{I} : \mathcal{I}')$ by definition of lattice quotient. Now observe that

$$\mathcal{I} = \mathcal{I}(\mathcal{I}')^{-1}\mathcal{I}' \subseteq (\mathcal{I} : \mathcal{I}')\mathcal{I}' \subseteq \mathcal{I},$$

so the inclusions are equalities. Multiplying both sides by $(\mathcal{I}')^{-1}$ yields the claim. $\square$

As we shall see below, several sets of interest in this work can be defined as quotients of various lattices. The following gives an alternative characterization of the lattice quotient, which shows that it is also a lattice.

Lemma 2.3. *For any lattices $\mathcal{L}, \mathcal{L}'$ in K, we have $(\mathcal{L} : \mathcal{L}') = (\mathcal{L}^{\vee}\mathcal{L}')^{\vee}$.*

Proof. For any $x \in K$, we have

$$x \in (\mathcal{L}^{\vee}\mathcal{L}')^{\vee} \iff \mathrm{Tr}((x\mathcal{L}')\mathcal{L}^{\vee}) \subseteq \mathbb{Z} \iff x\mathcal{L}' \subseteq (\mathcal{L}^{\vee})^{\vee} = \mathcal{L}. \qquad \square$$

---

[10]We stress that this notion is entirely different from the quotient group $\mathcal{L}'/\mathcal{L}$ of lattices $\mathcal{L} \subseteq \mathcal{L}'$; in particular, no containment relationship is required.

2.4.1   Coefficient Rings and Conductor Ideals

Every lattice has an associated (maximal) order called its *coefficient ring*, which is simply the quotient of the lattice with itself.

Definition 2.4 (Coefficient ring). For a lattice $\mathcal{L}$ in K, its *coefficient ring* is

$$\mathcal{O}^{\mathcal{L}} := (\mathcal{L} : \mathcal{L}) = \{x \in K : x\mathcal{L} \subseteq \mathcal{L}\}.$$

Recall from Lemma 2.3 that $\mathcal{O}^{\mathcal{L}} = (\mathcal{L}\mathcal{L}^{\vee})^{\vee}$, so $\mathcal{O}^{\mathcal{L}} = \mathcal{O}^{\mathcal{L}^{\vee}}$. Moreover, if $\mathcal{L}$ is an order $\mathcal{O}$ or its dual $\mathcal{O}^{\vee}$, then $\mathcal{O}^{\mathcal{L}} = \mathcal{O}$. The following lemma explains the choice of the name "coefficient ring."

Lemma 2.5. *For any lattice $\mathcal{L}$ in K, its coefficient ring $\mathcal{O}^{\mathcal{L}}$ is the maximal order of K for which $\mathcal{L}$ is a fractional ideal.*

Proof. First we show that $\mathcal{O}^{\mathcal{L}}$ is an order of K. It is clear that $\mathcal{O}^{\mathcal{L}} = (\mathcal{L} \cdot \mathcal{L}^{\vee})^{\vee}$ is a lattice in K (because $\mathcal{L} \cdot \mathcal{L}^{\vee}$ is), thus we only need to show that it is a subring of K with unity. We clearly have $1 \in \mathcal{O}^{\mathcal{L}}$ by definition. Moreover, for any $x, y \in \mathcal{O}^{\mathcal{L}}$, we have $(xy)\mathcal{L} = x(y\mathcal{L}) \subseteq x\mathcal{L} \subseteq \mathcal{L}$, so $xy \in \mathcal{O}^{\mathcal{L}}$, as required.

By definition we have that $\mathcal{O}^{\mathcal{L}} \cdot \mathcal{L} \subseteq \mathcal{L}$, so $\mathcal{L}$ is a fractional $\mathcal{O}^{\mathcal{L}}$-ideal. And for any order $\mathcal{O}$ of K, if $\mathcal{O} \cdot \mathcal{L} \subseteq \mathcal{L}$ then $\mathcal{O} \subseteq (\mathcal{L} : \mathcal{L}) = \mathcal{O}^{\mathcal{L}}$, so $\mathcal{O}^{\mathcal{L}}$ is the maximal order of K for which $\mathcal{L}$ is a fractional ideal. $\square$

Lemma 2.6. *For any lattices $\mathcal{L}, \mathcal{L}'$ in K, both $\mathcal{O}^{\mathcal{L}}$ and $\mathcal{O}^{\mathcal{L}'}$ are contained in $\mathcal{O}^{(\mathcal{L}:\mathcal{L}')}$. In particular, $(\mathcal{L} : \mathcal{L}')$ is a fractional ideal of all three of these coefficient rings.*

Proof. We have $\mathcal{O}^{\mathcal{L}}(\mathcal{L} : \mathcal{L}') = (\mathcal{L} : \mathcal{L})(\mathcal{L} : \mathcal{L}') \subseteq (\mathcal{L} : \mathcal{L}')$, so $\mathcal{O}^{\mathcal{L}} \subseteq \mathcal{O}^{(\mathcal{L}:\mathcal{L}')}$ by definition of coefficient ring. The proof for $\mathcal{O}^{\mathcal{L}'}$ is essentially identical, but with the order of multiplication reversed. $\square$

Throughout the paper, let $\mathcal{L}_{\mathcal{Q}}$ denote the (finite) quotient $\mathcal{O}^{\mathcal{L}}$-module $\mathcal{L}/\mathcal{Q}\mathcal{L}$ for any lattice $\mathcal{L}$ of K and any subset $\mathcal{Q} \subseteq \mathcal{O}^{\mathcal{L}}$ having a nonzero element.[11] Without loss of generality, $\mathcal{Q}$ may be taken to be an $\mathcal{O}^{\mathcal{L}}$-ideal, i.e., we can replace $\mathcal{Q}$ with the $\mathcal{O}^{\mathcal{L}}$-ideal $\mathcal{Q}\mathcal{O}^{\mathcal{L}}$ because $\mathcal{Q}\mathcal{O}^{\mathcal{L}}\mathcal{L} = \mathcal{Q}\mathcal{L}$. When $\mathcal{Q} = \{q\}$ is a singleton set, or equivalently, when $\mathcal{Q} = q\mathcal{O}^{\mathcal{L}}$, we often just use q as the subscript, as in $\mathcal{L}_q$.

Definition 2.7 (Conductor ideal). For any lattices $\mathcal{L} \subseteq \mathcal{L}'$ in K, their *conductor ideal* is $(\mathcal{L} : \mathcal{L}')$.

We remark that the conductor is often defined only for *orders* $\mathcal{O} \subseteq \mathcal{O}'$, but we find it useful to define it more generally for (sub)lattices, since many of the conductor's properties still carry over to this setting. For example, the following lemma explains the choice of the name "conductor ideal."

Lemma 2.8. *For any lattices $\mathcal{L} \subseteq \mathcal{L}'$ in K, the conductor ideal $(\mathcal{L} : \mathcal{L}')$ is an (integral) ideal of both $\mathcal{O}^{\mathcal{L}}$ and $\mathcal{O}^{\mathcal{L}'}$, and hence also of the order $\mathcal{O}^{\mathcal{L}} \cap \mathcal{O}^{\mathcal{L}'}$.*

Proof. By Lemma 2.6, $(\mathcal{L} : \mathcal{L}')$ is a fractional ideal of both $\mathcal{O}^{\mathcal{L}}$ and $\mathcal{O}^{\mathcal{L}'}$, so it suffices to show that it is contained in both orders. Because $\mathcal{L} \subseteq \mathcal{L}'$ and by definition of lattice quotient, we have $1 \in (\mathcal{L}' : \mathcal{L})$. Therefore,

$$(\mathcal{L} : \mathcal{L}') \subseteq (\mathcal{L} : \mathcal{L}')(\mathcal{L}' : \mathcal{L}) \subseteq (\mathcal{L} : \mathcal{L}) = \mathcal{O}^{\mathcal{L}}.$$

The proof for $\mathcal{O}^{\mathcal{L}'}$ proceeds almost identically, but with the order of multiplication reversed. $\square$

---

[11]Note that $\mathcal{Q}\mathcal{L} = \mathcal{Q}\mathcal{O}^{\mathcal{L}}\mathcal{L} \subseteq \mathcal{O}^{\mathcal{L}}\mathcal{L} = \mathcal{L}$, so $\mathcal{L}/\mathcal{Q}\mathcal{L}$ is well defined, and is finite because $\mathcal{Q}\mathcal{L}$ is a lattice.

2.4.2 Alternative Notions of Invertibility

Recall from above that every non-maximal order of K has a non-invertible ideal. In particular, for any orders $\mathcal{O} \subsetneq \mathcal{O}'$, the conductor ideal $\mathcal{C} = (\mathcal{O} : \mathcal{O}')$ is not invertible as an $\mathcal{O}$-ideal: for if $\mathcal{CI} = \mathcal{O}$ for some fractional $\mathcal{O}$-ideal $\mathcal{I}$, then because $\mathcal{C}$ is an $\mathcal{O}'$-ideal, we have $\mathcal{O}' = \mathcal{O}'\mathcal{O} = \mathcal{O}'\mathcal{CI} = \mathcal{CI} = \mathcal{O}$, a contradiction. Yet despite the lack of ideal inverses in general, there is a proxy that turns out to be just as good for our purposes. Throughout the rest of this subsection let $\mathcal{O}$ be an arbitrary order of K.

Definition 2.9 (Pseudoinverse). The *pseudoinverse* of a fractional $\mathcal{O}$-ideal $\mathcal{I}$ is $(\mathcal{O} : \mathcal{I})$.

We stress that both the inverse and pseudoinverse are defined with respect to a particular order $\mathcal{O}$. Note that by Lemma 2.2, for an invertible ideal, the pseudoinverse is simply the inverse ideal. In any case, the pseudoinverse is a fractional $\mathcal{O}$-ideal, (i.e., $\mathcal{O}(\mathcal{O} : \mathcal{I}) \subseteq (\mathcal{O} : \mathcal{I})$), because $\mathcal{O}(\mathcal{O} : \mathcal{I})\mathcal{I} \subseteq \mathcal{O}\mathcal{O} = \mathcal{O}$. Using the pseudoinverse we generalize the notion of ideal invertibility, by defining it modulo another (integral) ideal.

Definition 2.10. For an $\mathcal{O}$-ideal $\mathcal{Q}$, we say that a fractional $\mathcal{O}$-ideal $\mathcal{I}$ is *invertible modulo* $\mathcal{Q}$ if the $\mathcal{O}$-ideals $(\mathcal{O} : \mathcal{I})\mathcal{I}$ and $\mathcal{Q}$ are coprime, i.e., $(\mathcal{O} : \mathcal{I})\mathcal{I} + \mathcal{Q} = \mathcal{O}$.

Observe that if $\mathcal{I}$ is invertible, then it is also invertible modulo any $\mathcal{Q}$, simply because $(\mathcal{O} : \mathcal{I})\mathcal{I} = \mathcal{O}$. Furthermore, if $\mathcal{Q}$ and the conductor ideal $(\mathcal{O} : \mathcal{O}_K)$ are coprime (in $\mathcal{O}$), then *any* fractional $\mathcal{O}$-ideal $\mathcal{I}$ is invertible modulo $\mathcal{Q}$, because $\mathcal{O} = (\mathcal{O} : \mathcal{O}_K) + \mathcal{Q} \subseteq (\mathcal{O} : \mathcal{I})\mathcal{I} + \mathcal{Q}$ by the following lemma.

Lemma 2.11. *For any fractional $\mathcal{O}$-ideal $\mathcal{I}$, we have $(\mathcal{O} : \mathcal{O}_K) \subseteq (\mathcal{O} : \mathcal{I})\mathcal{I} \subseteq (\mathcal{O} : \mathcal{O}^{\mathcal{I}})$. In particular, both inclusions are equalities when $\mathcal{I}$ is an $\mathcal{O}_K$-ideal.*

Proof. The second sentence follows immediately from the fact that $\mathcal{O}^{\mathcal{I}} = \mathcal{O}_K$. For the second inclusion, by (2.1) we have that

$$(\mathcal{O} : \mathcal{I})\mathcal{I}\mathcal{O}^{\mathcal{I}} = (\mathcal{O} : \mathcal{I})(\mathcal{I} : \mathcal{I})\mathcal{I} \subseteq (\mathcal{O} : \mathcal{I})\mathcal{I} \subseteq \mathcal{O},$$

as needed. For the first inclusion, to show that some $A \subseteq \mathcal{O}$ is also a subset of $(\mathcal{O} : \mathcal{I})\mathcal{I}$, it suffices to express $A = A'\mathcal{I}$ for some $A' \subseteq K$, because then $A' \subseteq (\mathcal{O} : \mathcal{I})$. Since $(\mathcal{O}_K : \mathcal{O}) \subseteq \mathcal{O}$ is an $\mathcal{O}_K$-ideal, and $\mathcal{O}_K\mathcal{I}$ is a fractional $\mathcal{O}_K$-ideal and hence invertible (with respect to $\mathcal{O}_K$), we have that

$$(\mathcal{O} : \mathcal{O}_K) = (\mathcal{O} : \mathcal{O}_K)\mathcal{O}_K = (\mathcal{O} : \mathcal{O}_K)(\mathcal{O}_K\mathcal{I})^{-1}\mathcal{O}_K\mathcal{I}. \qquad \square$$

2.5 Chinese Remainder Theorem

We now recall a general form of the Chinese Remainder Theorem (CRT) and its consequences for our work. The theorem is often stated for the special case of $\mathcal{M} = \mathcal{O}$, in which case it additionally yields a ring isomorphism; the more general form below immediately follows by tensoring the isomorphism with $\mathcal{M}$, which can be done because $\mathcal{M}$ is an $\mathcal{O}$-module.

Theorem 2.12 (Chinese Remainder Theorem). *Let $\mathcal{I}_1, \ldots, \mathcal{I}_r$ be any pairwise coprime (equivalently, co-maximal) $\mathcal{O}$-ideals, let $\mathcal{I} = \prod_{i=1}^{r} \mathcal{I}_i$, and let $\mathcal{M}$ be any fractional $\mathcal{O}$-ideal. Then the natural $\mathcal{O}$-module homomorphism*

$$\mathcal{M}/\mathcal{I}\mathcal{M} \to \bigoplus_{i=1}^{r} \mathcal{M}/\mathcal{I}_i\mathcal{M}$$

is an isomorphism (which is efficiently computable and invertible given the $\mathcal{I}_i$ and $\mathcal{M}$).

We next show several important consequences of CRT for our purposes. The following generalizes [LPR10, Lemmas 2.14 and 2.15] and [BBPS19, Lemma 2.35] to arbitrary orders $\mathcal{O}$ and possibly *non-invertible* (fractional) ideals $\mathcal{I}$.

Lemma 2.13. *Let $\mathcal{O}$ be an order of a number field K, $\mathcal{Q}$ be an $\mathcal{O}$-ideal, $\mathcal{I}$ and $\mathcal{I}'$ be fractional $\mathcal{O}$-ideals, and suppose that $(\mathcal{I}' : \mathcal{I})(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$. Then there exists $t \in (\mathcal{I}' : \mathcal{I})$ such that $t(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$. Moreover, such a t can be found in polynomial time given $\mathcal{O}, \mathcal{I}, \mathcal{I}'$, and all the prime (equivalently, maximal) $\mathcal{O}$-ideals that contain $\mathcal{Q}$.*

As a special case, a fractional $\mathcal{O}$-ideal $\mathcal{I}'$ is invertible modulo $\mathcal{Q}$ if and only if there exists $t \in \mathcal{I}'$ such that $t(\mathcal{O} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$.

Proof. Let $\mathfrak{p}_1, \dots, \mathfrak{p}_r$ be the distinct prime (equivalently, maximal) $\mathcal{O}$-ideals that contain $\mathcal{Q}$. First we show that any $t \in (\mathcal{I}' : \mathcal{I}) \setminus \bigcup_{i=1}^r \mathfrak{p}_i(\mathcal{I}' : \mathcal{I})$ suffices. We have that $t(\mathcal{I} : \mathcal{I}') + \mathcal{Q} \subseteq (\mathcal{I}' : \mathcal{I})(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$. Suppose for contradiction that the inclusion is not an equality. Then $t(\mathcal{I} : \mathcal{I}') + \mathcal{Q} \subseteq \mathfrak{p}$ for some maximal (and hence prime) ideal $\mathfrak{p} \subseteq \mathcal{O}$, which implies that $t(\mathcal{I} : \mathcal{I}') \subseteq \mathfrak{p}$ and $\mathcal{Q} \subseteq \mathfrak{p}$, so $\mathfrak{p} = \mathfrak{p}_i$ for some i. By these inclusions, the fact that $t \in (\mathcal{I}' : \mathcal{I})$, and the hypothesis that $(\mathcal{I}' : \mathcal{I})(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$, we have that

$$t \in t\mathcal{O} = t((\mathcal{I}' : \mathcal{I})(\mathcal{I} : \mathcal{I}') + \mathcal{Q}) \subseteq \mathfrak{p}_i(\mathcal{I}' : \mathcal{I}) + t\mathcal{Q} \subseteq \mathfrak{p}_i(\mathcal{I}' : \mathcal{I}) , \tag{2.2}$$

so $t \in \mathfrak{p}_i(\mathcal{I}' : \mathcal{I})$, which contradicts the choice of t.

Now, we show that such a t exists and can be computed efficiently. First, note that $(\mathcal{I}' : \mathcal{I}) \neq \mathfrak{p}_i(\mathcal{I}' : \mathcal{I})$ for all i.[12] So, for each i, choose some non-zero $t_i \in (\mathcal{I}' : \mathcal{I})/\mathfrak{p}_i(\mathcal{I}' : \mathcal{I})$, and let $t \in (\mathcal{I}' : \mathcal{I})$ be (an arbitrary representative of) the preimage of $(t_1, \dots, t_r) \in \bigoplus_{i=1}^r (\mathcal{I}' : \mathcal{I})/\mathfrak{p}_i(\mathcal{I}' : \mathcal{I})$ under the isomorphism given by Theorem 2.12 (which we can invoke here because the $\mathfrak{p}_i$ are distinct maximal ideals, and hence pairwise coprime). Clearly, $t \in (\mathcal{I}' : \mathcal{I}) \setminus \bigcup_{i=1}^r \mathfrak{p}_i(\mathcal{I}' : \mathcal{I})$, as desired.

For the special case, the 'only if' part follows immediately by letting $\mathcal{I} = \mathcal{O}$. For the 'if' part, if such a $t \in \mathcal{I}'$ exists then $\mathcal{O} = t(\mathcal{O} : \mathcal{I}') + \mathcal{Q} \subseteq \mathcal{I}'(\mathcal{O} : \mathcal{I}') + \mathcal{Q} \subseteq \mathcal{O}$, so the inclusions are equalities, hence $\mathcal{I}'$ is invertible modulo $\mathcal{Q}$. $\qquad\square$

Lemma 2.14. *Let $\mathcal{O}$ be an order of a number field K, $\mathcal{Q}$ be an $\mathcal{O}$-ideal, $\mathcal{I}, \mathcal{I}', \mathcal{J}$ be fractional $\mathcal{O}$-ideals, and $t \in (\mathcal{I}' : \mathcal{I})$ be such that $t(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$. Then the function $\theta_t \colon K \to K$ defined as $\theta_t(u) = t \cdot u$ induces an $\mathcal{O}$-module isomorphism from $(\mathcal{I}\mathcal{J})_\mathcal{Q}$ to $(\mathcal{I}'\mathcal{J})_\mathcal{Q}$.*

Proof. That θ_t induces an $\mathcal{O}$-module homomorphism follows immediately from the fact that it is multiplication by a fixed $t \in K$. Now consider the function from $\mathcal{I}\mathcal{J}$ to $(\mathcal{I}'\mathcal{J})_\mathcal{Q}$ that is induced by θ_t. Its kernel clearly contains $\mathcal{Q}\mathcal{I}\mathcal{J}$, because $t \in (\mathcal{I}' : \mathcal{I})$, and is in fact equal to $\mathcal{Q}\mathcal{I}\mathcal{J}$, which may be seen as follows. If $u \cdot t \in \mathcal{Q}\mathcal{I}'\mathcal{J}$ for some $u \in \mathcal{I}\mathcal{J}$, then $u \cdot t(\mathcal{I} : \mathcal{I}') \subseteq \mathcal{Q}\mathcal{I}'\mathcal{J}(\mathcal{I} : \mathcal{I}') \subseteq \mathcal{Q}\mathcal{I}\mathcal{J}$. Because $t(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$, we get that $u \in u\mathcal{O} = u(t(\mathcal{I} : \mathcal{I}') + \mathcal{Q}) \subseteq \mathcal{Q}\mathcal{I}\mathcal{J}$, as desired. So, the function from $(\mathcal{I}\mathcal{J})_\mathcal{Q}$ to $(\mathcal{I}'\mathcal{J})_\mathcal{Q}$ induced by θ_t is injective. It remains to show that it can be inverted, which also implies that it is an isomorphism.

Let $v \in \mathcal{I}'\mathcal{J}$ be arbitrary. By hypothesis, $t(\mathcal{I} : \mathcal{I}')$ and $\mathcal{Q}$ are coprime (in $\mathcal{O}$). Therefore, by Theorem 2.12 we can compute some $c \in t(\mathcal{I} : \mathcal{I}')$ such that $c = 1 \pmod{\mathcal{Q}}$. Then let $a = c \cdot v \in t(\mathcal{I} : \mathcal{I}')\mathcal{I}'\mathcal{J} \subseteq t\mathcal{I}\mathcal{J}$, and observe that $a - v = (c-1)v \in \mathcal{Q}\mathcal{I}'\mathcal{J}$. Let $w = a/t \in \mathcal{I}\mathcal{J}$; then $\theta_t(w) = a = v \pmod{\mathcal{Q}\mathcal{I}'\mathcal{J}}$, so $w \bmod \mathcal{Q}\mathcal{I}\mathcal{J}$ is the preimage of $v \bmod \mathcal{Q}\mathcal{I}'\mathcal{J}$. $\qquad\square$

---

[12]This is by Nakayama's lemma (see [Mat89, Theorem 2.2, page 8]), which says that if $\mathcal{I} = \mathfrak{p}_i\mathcal{I}$, then there is an $r \in \mathcal{O}$ such that $r = 1 \bmod \mathfrak{p}_i$ (hence $r \neq 0$) and $r\mathcal{I} = \{0\}$. Since K is an integral domain, this implies that $\mathcal{I} = \{0\}$ is the zero ideal, which we have ruled out for the entire paper.

Special case: bijective natural inclusions. For any lattices $\mathcal{L} \subseteq \mathcal{L}'$ in K and set $\mathcal{Q} \subseteq \mathcal{O}^{\mathcal{L}} \cap \mathcal{O}^{\mathcal{L}'}$ containing a nonzero element (so that both $\mathcal{L}_{\mathcal{Q}}$ and $\mathcal{L}'_{\mathcal{Q}}$ are well defined), the *natural inclusion map* $\mathcal{L}_{\mathcal{Q}} \to \mathcal{L}'_{\mathcal{Q}}$ sends the coset $x + \mathcal{Q}\mathcal{L}$ to the coset $x + \mathcal{Q}\mathcal{L}'$. (This can be seen as the composition of an inclusion map and a natural homomorphism.) The following lemmas give conditions under which maps of this kind are bijections.

Corollary 2.15. *Adopt the notation of Lemma 2.14, and also assume that $\mathcal{I} \subseteq \mathcal{I}'$ and $(\mathcal{I} : \mathcal{I}') + \mathcal{Q} = \mathcal{O}$. Then the natural inclusion map $h \colon (\mathcal{I}\mathcal{J})_{\mathcal{Q}} \to (\mathcal{I}'\mathcal{J})_{\mathcal{Q}}$ is a bijection. Moreover, because $(\mathcal{I}')^{\vee} \subseteq \mathcal{I}^{\vee}$ and $((\mathcal{I}')^{\vee} : \mathcal{I}^{\vee}) = (\mathcal{I} : \mathcal{I}')$, the same holds for the natural inclusion map $((\mathcal{I}')^{\vee}\mathcal{J})_{\mathcal{Q}} \to (\mathcal{I}^{\vee}\mathcal{J})_{\mathcal{Q}}$.*

Proof. This follows immediately from Lemma 2.14, by letting $t = 1 \in (\mathcal{I}' : \mathcal{I})$, which makes the $\mathcal{O}$-module isomorphism induced by θ_t simply the natural inclusion map. $\qquad\square$

The following shows that if the hypotheses of Corollary 2.15 hold, and we have a corresponding inclusion for the ideals' *coefficient rings*, then the coprimality condition also holds for the coefficient rings themselves (thus yielding another bijective natural inclusion).

Corollary 2.16. *Let $\mathcal{L} \subseteq \mathcal{L}'$ be lattices in a number field K, let $\mathcal{O} = \mathcal{O}^{\mathcal{L}}$ and $\mathcal{O}' = \mathcal{O}^{\mathcal{L}'}$ denote their respective coefficent rings with $\mathcal{O} \subseteq \mathcal{O}'$, and let $\mathcal{Q}$ be an $\mathcal{O}$-ideal such that $(\mathcal{L} : \mathcal{L}') + \mathcal{Q} = \mathcal{O}$. Then $(\mathcal{O} : \mathcal{O}') + \mathcal{Q} = \mathcal{O}$.*

Note that the converse does not hold: let $\mathcal{L} = q\mathcal{L}'$ and $\mathcal{Q} = q\mathcal{O}$ for some positive integer q; then $\mathcal{O} = \mathcal{O}' = (\mathcal{O} : \mathcal{O}')$, but $(\mathcal{L} : \mathcal{L}') = q\mathcal{O} = \mathcal{Q}$.

Proof. It suffices to show that $(\mathcal{L} : \mathcal{L}') \subseteq (\mathcal{O} : \mathcal{O}')$, because the latter is contained in $\mathcal{O}$ by Lemma 2.8. To see this, observe that

$$(\mathcal{L} : \mathcal{L}')\mathcal{O}' = (\mathcal{L} : \mathcal{L}')(\mathcal{L}' : \mathcal{L}') \subseteq (\mathcal{L} : \mathcal{L}') \subseteq (\mathcal{L} : \mathcal{L}')(\mathcal{L}' : \mathcal{L}) \subseteq (\mathcal{L} : \mathcal{L}) = \mathcal{O},$$

where for the second inclusion we use the fact that $1 \in (\mathcal{L}' : \mathcal{L})$ because $\mathcal{L} \subseteq \mathcal{L}'$. $\qquad\square$

2.6 Gaussians

To formally define Gaussian distributions over number fields, we need the field tensor product $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$, which is essentially the "real analogue" of $K/\mathbb{Q}$, obtained by generalizing rational scalars to real ones. In general this is not a field, but it is a ring; in fact, it is isomorphic to the ring product $\mathbb{R}^{s_1} \times \mathbb{C}^{s_2}$, where K has s_1 real embeddings and s_2 conjugate pairs of complex ring embeddings, and $n = s_1 + 2s_2$. Therefore, there is a "complex conjugation" involution $\tau \colon K_{\mathbb{R}} \to K_{\mathbb{R}}$, which corresponds to the identity map on each $\mathbb{R}$ component, and complex conjugation on each $\mathbb{C}$ component.

We extend the trace to $K_{\mathbb{R}}$ in the natural way, writing $\mathrm{Tr}_{K_{\mathbb{R}}/\mathbb{R}}$ for the resulting $\mathbb{R}$-linear transform. It turns out that under the ring isomorphism with $\mathbb{R}^{s_1} \times \mathbb{C}^{s_2}$, this trace corresponds to the sum of the real components plus twice the sum of the real parts of the complex components. From this it can be verified that $K_{\mathbb{R}}$ is an n-dimensional real inner-product space, with inner product $\langle x, y \rangle = \mathrm{Tr}_{K_{\mathbb{R}}/\mathbb{R}}(x \cdot \tau(y))$. In particular, $K_{\mathbb{R}}$ has some (non-unique) orthonormal basis $\vec{b}$, and hence $\vec{b}^{\vee} = \tau(\vec{b})$.

Now let H be any n-dimensional real inner-product space (e.g., $H = \mathbb{R}^n$ or $H = K_{\mathbb{R}}$) and fix an orthonormal basis, so that any element $x \in H$ may be uniquely represented as a real vector $\mathbf{x} \in \mathbb{R}^n$ relative to that basis.

Definition 2.17. For a positive semidefinite $\Sigma \in \mathbb{R}^{n \times n}$, which we call the *covariance* matrix, the *Gaussian function* $\rho_{\sqrt{\Sigma}} \colon H \to (0, 1]$ is defined as $\rho_{\sqrt{\Sigma}}(\mathbf{x}) := \exp(-\pi \mathbf{x}^t \cdot \Sigma^- \cdot \mathbf{x})$ for $\mathbf{x} \in \operatorname{span}(\Sigma) = \Sigma \cdot \mathbb{R}^n$ and $\rho_{\sqrt{\Sigma}}(\mathbf{x}) := 0$ otherwise, where Σ^- denotes the (Moore-Penrose) pseudoinverse. The *Gaussian distribution* $D_{\sqrt{\Sigma}}$ on H is the one whose probability density function (when restricted to $\operatorname{span}(\Sigma)$) is proportional to $\rho_{\sqrt{\Sigma}}$.[13]

When $\Sigma = r^2 \cdot \mathbf{I}$ for some $r \geq 0$ we often write ρ_r and D_r instead, and refer to these as *spherical* Gaussians with parameter r. (In this case, the choice of orthonormal basis for H is immaterial, i.e., any orthonormal basis yields the same $\Sigma = r^2 \cdot \mathbf{I}$.)

It is easy to verify that for any positive semidefinite Σ and any matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, the distribution $\mathbf{A} \cdot D_{\sqrt{\Sigma}} = D_{\sqrt{\Sigma'}}$, where $\Sigma' = \mathbf{A} \cdot \Sigma \cdot \mathbf{A}^t$. It is also well known that the sum of two independent Gaussians having covariances Σ_1, Σ_2 (respectively) is distributed as a Gaussian with covariance $\Sigma_1 + \Sigma_2$. Therefore, a Gaussian of covariance Σ can be transformed into one of any desired covariance $\Sigma' \succeq \Sigma$, i.e., one for which $\Sigma' - \Sigma$ is positive semidefinite, simply by adding an independent compensating Gaussian of covariance $\Sigma' - \Sigma$.

2.7 Extension Fields

For the material in Section 6 we need to generalize some of our definitions to number field extensions K'/K, where possibly $K \neq \mathbb{Q}$. The (field) *trace* $\operatorname{Tr} = \operatorname{Tr}_{K'/K} \colon K' \to K$ is the function mapping any element $x \in K'$ to the trace of the K-linear transformation on K', viewed as a vector space over K, representing multiplication by x. (In more detail, fixing an arbitrary K-basis of K', every element of K' can be represented as a vector with entries in K, and multiplication by x corresponds to multiplication by a square matrix with entries in K; the trace of x is the trace of this matrix.) We extend the trace to the real inner-product spaces $K'_{\mathbb{R}}$ and $K_{\mathbb{R}}$ in the natural way, writing $\operatorname{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}$ for the resulting $K_{\mathbb{R}}$-linear transform.

Let $\vec{b} = (b_1, \ldots, b_k)$ be a K-basis of K'. Its *dual basis* $\vec{b}^{\vee} = (b_1^{\vee}, \ldots, b_k^{\vee})$ is defined so that $\operatorname{Tr}_{K'/K}(b_i b_j^{\vee})$ is 1 when $i = j$, and is 0 otherwise. For a lattice $\mathcal{L}$ in K', its *dual lattice relative to an order $\mathcal{O}$ of K*, also called the $\mathcal{O}$-*dual of $\mathcal{L}$* for short, is defined as

$$\mathcal{L}^{\vee \mathcal{O}} := \{x \in K' : \operatorname{Tr}_{K'/K}(x\mathcal{L}) \subseteq \mathcal{O}\}.$$

Notice that this generalizes our prior definition of the dual lattice for $K = \mathbb{Q}$, whose only order is $\mathcal{O} = \mathbb{Z}$. Also, it is easy to verify that if $\vec{b}$ is an $\mathcal{O}$-basis of $\mathcal{L}$, then $\vec{b}^{\vee}$ is an $\mathcal{O}$-basis of $\mathcal{L}^{\vee \mathcal{O}}$.

For a tower $K''/K'/K$ of number fields extensions (i.e., K''/K' and K'/K are both number field extensions), it is easy to verify from the definitions that the trace decomposes as $\operatorname{Tr}_{K''/K} = \operatorname{Tr}_{K'/K} \circ \operatorname{Tr}_{K''/K'}$. Moreover, if $\vec{c}_2$ is a K'-basis of K'' and $\vec{c}_1$ is a K-basis of K', then by definition and K'-linearity of $\operatorname{Tr}_{K''/K'}$ we have $(\vec{c}_2 \otimes \vec{c}_1)^{\vee} = \vec{c}_2^{\vee} \otimes \vec{c}_1^{\vee}$.

Lemma 2.18. *Let $K''/K'/K$ be a tower of number field extensions and $\mathcal{O}''$, $\mathcal{O}'$, and $\mathcal{O}$ be orders of K'', K', and K, respectively. Then $(\mathcal{O}'')^{\vee \mathcal{O}'} \cdot (\mathcal{O}')^{\vee \mathcal{O}} \subseteq (\mathcal{O}'')^{\vee \mathcal{O}}$, with equality if $\mathcal{O}''$ is a free $\mathcal{O}'$-module and $\mathcal{O}'$ is a free $\mathcal{O}$-module.*

(Equality may hold more generally, but the above statement is all we will need in our application.)

---

[13]Note that the covariance of $D_{\sqrt{\Sigma}}$ is actually $\Sigma/(2\pi)$, due to the normalization factor in the definition of $\rho_{\sqrt{\Sigma}}$.

Proof. For the first claimed inclusion, observe that by the decomposition $\mathrm{Tr}_{K''/K} = \mathrm{Tr}_{K'/K} \circ \mathrm{Tr}_{K''/K'}$, K'-linearity of $\mathrm{Tr}_{K''/K'}$, and the definition of dual lattice, we have that

$$\mathrm{Tr}_{K''/K}\big(\mathcal{O}'' \cdot (\mathcal{O}'')^{\vee_{\mathcal{O}'}} \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}}\big) = \mathrm{Tr}_{K'/K}\big(\mathrm{Tr}_{K''/K'}\big(\mathcal{O}'' \cdot (\mathcal{O}'')^{\vee_{\mathcal{O}'}}\big) \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}}\big)$$
$$\subseteq \mathrm{Tr}_{K'/K}\big(\mathcal{O}' \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}}\big) \subseteq \mathcal{O},$$

as needed.

For the reverse inclusion (under the additional hypothesis), let $\vec{c}_2$ be a (finite) $\mathcal{O}'$-basis of $\mathcal{O}''$, and $\vec{c}_1$ be a (finite) $\mathcal{O}$-basis of $\mathcal{O}'$. Then $\vec{c}_2 \otimes \vec{c}_1$ is an $\mathcal{O}$-basis of $\mathcal{O}''$, thus $(\vec{c}_2 \otimes \vec{c}_1)^\vee = \vec{c}_2^\vee \otimes \vec{c}_1^\vee$ is an $\mathcal{O}$-basis of $(\mathcal{O}'')^{\vee_{\mathcal{O}}}$, and all of its entries are in the $\mathcal{O}$-ideal $(\mathcal{O}'')^{\vee_{\mathcal{O}'}} \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}}$. Therefore, $(\mathcal{O}'')^{\vee_{\mathcal{O}}} \subseteq (\mathcal{O}'')^{\vee_{\mathcal{O}'}} \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}}$, as needed. $\qquad\square$

3 Generalized (Algebraic) Learning With Errors

In this section we define a generalized form of LWE and relate it to the various prior LWE variants. First, in Section 3.1 we give a unified framework that encompasses all LWE variants (over commutative rings) that we are aware of. Then, in Section 3.2 we show in particular how to obtain all "algebraic" forms of LWE over number fields, including Ring-, Order-, and Poly-LWE, simply by parameterizing our generalized LWE by a lattice in the number field.

3.1 Generalized LWE

Here we describe a general framework that captures all variants of Learning With Errors (over commutative rings) of which we are aware, and is helpful in linking them together. Our starting point is the observation that in all such problems, the secret s, public multipliers a, and their (noiseless) products $s \cdot a$ each belong to a quotient $\mathcal{I}_{\mathcal{Q}} = \mathcal{I}/\mathcal{Q}\mathcal{I}$ (or its many-fold Cartesian product) for some respective *fractional ideals* $\mathcal{I}$ of some common order $\mathcal{O}$ of a number field and a fixed "modulus" $\mathcal{O}$-ideal $\mathcal{Q}$. Moreover, the products are given by some fixed $\mathcal{O}_{\mathcal{Q}}$-*bilinear map* on s and a. As a few examples:

- Ordinary LWE uses the $\mathbb{Z}_q$-bilinear inner-product map $\langle \cdot, \cdot \rangle \colon \mathbb{Z}_q^n \times \mathbb{Z}_q^n \to \mathbb{Z}_q$, where the secret, multipliers, and products all are associated with the ideal $\mathcal{I} = \mathbb{Z}$ of the unique order $\mathcal{O} = \mathbb{Z}$ in the rational number field $K = \mathbb{Q}$, and q is an integer modulus.

- Ring-LWE uses the R_q-bilinear multiplication map $R_q^\vee \times R_q \to R_q^\vee$, where the multipliers are associated with the maximal order $R = \mathcal{O}_K$ of a number field K, and the secret is associated with the "codifferent" fractional ideal $R^\vee$.

- Module-LWE interpolates between the above two cases, using the R_q-bilinear inner-product map $(R_q^\vee)^d \times R_q^d \to R_q^\vee$, where R and $R^\vee$ are as above.

A generalized LWE distribution is parameterized by:

1. an order $\mathcal{O}$ in a number field K and an $\mathcal{O}$-ideal $\mathcal{Q}$;

2. suitable fractional $\mathcal{O}$-ideals $\mathcal{S}, \mathcal{A}, \mathcal{B} = \mathcal{S}\mathcal{A}$;

3. dimensions k_s, k_a, k_b, and an order-three tensor $T \in \mathcal{O}_{\mathcal{Q}}^{k_s \times k_a \times k_b}$, which induces an $\mathcal{O}_{\mathcal{Q}}$-bilinear map $T \colon \mathcal{S}_{\mathcal{Q}}^{k_s} \times \mathcal{A}_{\mathcal{Q}}^{k_a} \to \mathcal{B}_{\mathcal{Q}}^{k_b}$ defined as $T(\vec{s}, \vec{a})_\ell = \sum_{i,j} T_{ij\ell} s_i a_j$; and

4. an error distribution ψ over $K_{\mathbb{R}}^{k_b}$.

Informally, the associated computational problems are concerned with "noisy products" ($\vec{a} \leftarrow \mathcal{A}_{\mathcal{Q}}^{k_a}, \vec{b} \approx T(\vec{s}, \vec{a})$) for some fixed $\vec{s} \in \mathcal{S}_{\mathcal{Q}}^{k_s}$. Clearly, different choices of the tensor T and/or error distribution ψ may yield different distributions of noisy products.

Definition 3.1 (LWE **distribution**). Adopt the above notation. For $\vec{s} \in \mathcal{S}_{\mathcal{Q}}^{k_s}$, a sample from the distribution $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ over $\mathcal{A}_{\mathcal{Q}}^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b}$ is generated by choosing $\vec{a} \leftarrow \mathcal{A}_{\mathcal{Q}}^{k_a}$ uniformly at random, choosing $\vec{e} \leftarrow \psi$, and outputting

$$(\vec{a}\ ,\ T(\vec{s}, \vec{a}) + \vec{e} \bmod (\mathcal{Q}\mathcal{B})^{k_b}).$$

For notational convenience, we also define the uniform distribution $U_{T,\mathcal{S},\mathcal{A}} = U(\mathcal{A}_{\mathcal{Q}}^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b})$.[14]

Definition 3.2 (LWE **problem, search**). The search-LWE$_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ problem is: given ℓ independent samples from $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ where $\vec{s} \leftarrow \mathcal{S}_{\mathcal{Q}}^{k_s}$, find $\vec{s}$.

Definition 3.3 (LWE **problem, decision**). The decision-LWE$_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ problem is to distinguish between ℓ independent samples from either $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ where $\vec{s} \leftarrow U(\mathcal{S}_{\mathcal{Q}}^{k_s})$, or $U_{T,\mathcal{S},\mathcal{A}}$.

Discussion. Requiring $\mathcal{S}$ and $\mathcal{A}$ to be fractional $\mathcal{O}$-ideals, and $\mathcal{Q}$ to be an $\mathcal{O}$-ideal, for some common order $\mathcal{O}$—rather than just some arbitrary lattices in K—is without loss of generality. Given lattices $\mathcal{S}$, $\mathcal{A}$, and $\mathcal{Q}$, in order for $\mathcal{S}_{\mathcal{Q}}$ and $\mathcal{A}_{\mathcal{Q}}$ to be well defined (see Section 2.4.1), without loss of generality $\mathcal{Q}$ is an ideal of both coefficient rings $\mathcal{O}^{\mathcal{S}}, \mathcal{O}^{\mathcal{A}}$, and hence of their intersection $\mathcal{O}$ as well. Then both $\mathcal{S}, \mathcal{A}$ are fractional $\mathcal{O}$-ideals, because they are fractional ideals of $\mathcal{O}^{\mathcal{S}}, \mathcal{O}^{\mathcal{A}}$ (respectively), and hence of any common suborder as well. (More generally, we may take $\mathcal{O}$ to be any common suborder of $\mathcal{O}^{\mathcal{S}}, \mathcal{O}^{\mathcal{A}}$ that contains $\mathcal{Q}$.)

Using a tensor over $\mathcal{O}_{\mathcal{Q}}$ to represent an $\mathcal{O}_{\mathcal{Q}}$-bilinear map $\mathcal{S}_{\mathcal{Q}}^{k_s} \times \mathcal{A}_{\mathcal{Q}}^{k_a} \to \mathcal{B}_{\mathcal{Q}}^{k_b}$ is also without loss of generality, under the mild assumption (which is also needed for some of our reductions) that the ideals $\mathcal{S}$ and $\mathcal{A}$—and hence $\mathcal{B} = \mathcal{S}\mathcal{A}$ as well—are invertible modulo $\mathcal{Q}$. That is, any bilinear map with such domain and range can be represented by an order-three tensor over $\mathcal{O}_{\mathcal{Q}}$. To see this, first note that by Lemma 2.14, for each $\mathcal{I} \in \{\mathcal{S}, \mathcal{A}, \mathcal{B}\}$ the $\mathcal{O}_{\mathcal{Q}}$-module $\mathcal{I}_{\mathcal{Q}}$ is isomorphic to $\mathcal{O}_{\mathcal{Q}}$. Because the latter module has a one-element $\mathcal{O}_{\mathcal{Q}}$-basis $\{1\}$, the former also has a one-element $\mathcal{O}_{\mathcal{Q}}$-basis $\{g_{\mathcal{I}}\}$ for some $g_{\mathcal{I}} \in \mathcal{I}_{\mathcal{Q}}$, where $g_{\mathcal{B}} = g_{\mathcal{S}} g_{\mathcal{A}}$. This naturally extends to the "standard basis" of $\mathcal{I}_{\mathcal{Q}}^{k}$, whose ith vector has $g_{\mathcal{I}}$ in its ith component and zeros elsewhere. Using these bases, any bilinear map T can be uniquely represented as an order-three tensor T over $\mathcal{O}_{\mathcal{Q}}$ by letting $T_{ij\ell} \in \mathcal{O}_{\mathcal{Q}}$ be the ℓth coefficient (with respect to the standard basis) of $T(\vec{e}_i, \vec{e}_j)$, where $\vec{e}_i, \vec{e}_j$ are respectively the ith and jth standard basis elements of their modules. By $\mathcal{O}_{\mathcal{Q}}$-bilinearity of the map and the fact that $g_{\mathcal{B}} = g_{\mathcal{S}} g_{\mathcal{A}}$, it follows that this tensor induces the bilinear map.

3.2 Parameterizing by a Single Lattice

We now define a special case of generalized LWE that still encompasses all prior algebraic LWE problems over number fields, including Ring-, Module-, Order-, and Poly-LWE. The key observation is that all of these problems can be obtained simply by taking the secret to be over (the dual of[15]) a certain *lattice* $\mathcal{L}$ in a given number field, taking the public multipliers to be over the lattice's *coefficient ring* $\mathcal{O} = \mathcal{O}^{\mathcal{L}}$, and using a tensor corresponding to an identity matrix. Indeed, the first two of these simplifications are without loss of generality under a mild assumption, by the reductions given in Theorems 4.1 and 4.5 (see Remark 4.6).

---

[14]Note that the *contents* of T are not used in the definition of $U_{T,\mathcal{S},\mathcal{A}}$, but its modulus $\mathcal{Q}$ and dimensions k_a, k_b are. So, for consistent notation between the two kinds of distributions, we include T as a subscript for both.

[15]We take the secret from the dual lattice for technical reasons that simplify the correspondence with prior problems like Ring- and Order-LWE, and the reductions between $\mathcal{L}$-LWE problems for different choices of the lattice $\mathcal{L}$.

Moreover, taking $\mathcal{L} = \mathcal{O}$ itself to be an order is especially advantageous for reductions from Ring-LWE and Order-LWE (as shown in Section 4.3), and to Middle-Product LWE (as shown in Section 5).

Definition 3.4 ($\mathcal{L}$-LWE problem). Let $\mathcal{L}$ be a lattice in a number field K, $\mathcal{O} = \mathcal{O}^{\mathcal{L}}$ be the coefficient ring of $\mathcal{L}$ (and of $\mathcal{L}^{\vee}$), ψ be a distribution over $K_{\mathbb{R}}$, $\mathcal{Q}$ be an $\mathcal{O}$-ideal, and k be a positive integer. Let $T \in \mathcal{O}_{\mathcal{Q}}^{k \times k \times 1}$ be the order-three tensor whose single $k \times k$ layer is the identity matrix. The (search or decision) $\mathcal{L}$-LWE$_{\mathcal{Q},\psi,\ell}^{k}$ problem is then simply the (search or decision, respectively) LWE$_{T,\mathcal{L}^{\vee},\mathcal{O},\psi,\ell}$ problem, i.e., with $\mathcal{A} = \mathcal{O}$ and $\mathcal{S} = \mathcal{B} = \mathcal{L}^{\vee}$.

We often omit k when $k = 1$; in this case, we have $s \in \mathcal{L}_{\mathcal{Q}}^{\vee}$, $a \in \mathcal{O}_{\mathcal{Q}}$, and a sample from the distribution $A_{T,\mathcal{L}^{\vee},\mathcal{O},\psi}(s)$ has the form $(a, b = s \cdot a + e \bmod \mathcal{Q}\mathcal{L}^{\vee})$.

Let us now see how the above definition strictly generalizes all prior algebraic LWE variants defined over number fields or polynomial rings. For simplicity, take $k = 1$ (taking $k > 1$ simply yields the "Module" analogues). Recall that if $\mathcal{L}$ is some order $\mathcal{O}$ of K or its dual $\mathcal{O}^{\vee}$, then $\mathcal{O}^{\mathcal{L}} = \mathcal{O}$. Therefore, by taking $\mathcal{L} = \mathcal{O}_K$ to be the full ring of integers (i.e., the maximal order in K), we get the Ring-LWE problem as originally defined in [LPR10], and by taking $\mathcal{L} = \mathcal{O}$ to be an arbitrary order we get Order-LWE [BBPS19].[16] Alternatively, by taking $\mathcal{L} = \mathcal{O}^{\vee}$ when $\mathcal{O} = \mathbb{Z}[\alpha]$ for some $\alpha \in \mathcal{O}_K$, we get the (primal) Poly-LWE problem [RSW18], which does not involve any dual lattices (hence the choice of $\mathcal{L}$). As we will see in Section 5, the "dual" formulations (i.e., $\mathcal{O}$-LWE for orders $\mathcal{O}$) have advantages over the "primal" formulations in terms of simplicity and tightness of reductions. Finally, by taking $\mathcal{L}$ to be neither an order nor the dual of an order, we get other problems that are not covered by any of the prior ones.

4 Generalized LWE Reductions

In this section we give a modular collection of tight, "minimal" reductions between various instantiations of generalized LWE. Each reduction alters a subset of the parameters, changing:

1. the ideals $\mathcal{S}, \mathcal{A}$ (Section 4.1);

2. the order $\mathcal{O}$ in the number field (Section 4.2), including the special case of changing the lattice $\mathcal{L}$ in $\mathcal{L}$-LWE (Section 4.3);

3. the tensor T (Section 4.4); or

4. the number field K over which all the other parameters are defined (Section 4.5),

with no loss in hardness of the associated LWE problems. Moreover, almost all of the reductions establish tight *equivalences* between problems, i.e., reductions in both directions. In later sections, we will obtain our main results for Order-LWE, Middle-Product LWE, etc., by suitably composing these individual reductions as building blocks.

---

[16]In "dual" Order-LWE as defined in [BBPS19], the domains of the secret s and the multipliers a_i are, respectively, $\mathcal{O}_{\mathcal{Q}}$ and $\mathcal{O}_{\mathcal{Q}}^{\vee}$, which are swapped relative to our definition of $\mathcal{O}$-LWE. However, as shown below in Corollary 4.3, the two problems are tightly equivalent (with the same error distribution), under a mild condition that is also needed for other reductions. This generalizes a similar observation made in [LPR10], that all Ring-LWE variants having the same product of the domains of s and a_i, and the same error distribution, are tightly equivalent.

4.1 Changing the Ideals

In this section we give reductions that map from one choice of the fractional ideals $\mathcal{S}, \mathcal{A}$ to another, while preserving the tensor T, error distribution ψ, and number of samples.

Our first theorem shows that without loss of generality, the entries of $\vec{a}$ may be chosen from $\mathcal{O}_\mathcal{Q}$ instead of $\mathcal{A}_\mathcal{Q}$ (with a corresponding change to the domain of the entries of $\vec{s}$) when $\mathcal{A}$ is invertible modulo $\mathcal{Q}$; recall from Section 2.4.2 that this is the case for *all* fractional ideals if $\mathcal{Q}$ is coprime to the conductor ideal $(\mathcal{O} : \mathcal{O}_K)$. This transformation is tight in all respects and reversible, so in fact it yields an *equivalence* between $\mathsf{LWE}_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ and $\mathsf{LWE}_{T,\mathcal{S}',\mathcal{A}',\psi,\ell}$ (in either their search or decision forms) whenever $\mathcal{A}, \mathcal{A}'$ are invertible modulo $\mathcal{Q}$, and $\mathcal{A}\mathcal{S} = \mathcal{A}'\mathcal{S}'$ (see Corollary 4.3).

Theorem 4.1. *Let $\mathcal{O}$ be an order in a number field K; $\mathcal{S}$ and $\mathcal{A}$ be fractional $\mathcal{O}$-ideals with $\mathcal{B} = \mathcal{S}\mathcal{A}$; $\mathcal{Q}$ be an $\mathcal{O}$-ideal; k_s, k_a, and k_b be positive integers; $T \in \mathcal{O}_\mathcal{Q}^{k_s \times k_a \times k_b}$ be an order-three tensor; and ψ be a distribution over $K_\mathbb{R}^{k_b}$. If $\mathcal{A}$ is invertible modulo $\mathcal{Q}$, then there is a polynomial-time computable and invertible deterministic transform that, given all the prime (equivalently, maximal) $\mathcal{O}$-ideals that contain $\mathcal{Q}$:*

1. *maps distribution $U_{T,\mathcal{S},\mathcal{A}}$ to distribution $U_{T,\mathcal{B},\mathcal{O}}$, and*

2. *maps distribution $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ to distribution $A_{T,\mathcal{B},\mathcal{O},\psi}(\vec{s}')$, where $\vec{s}' = h(\vec{s})$ for some $\mathcal{O}$-module isomorphism $h: \mathcal{S}_\mathcal{Q} \to \mathcal{B}_\mathcal{Q}$.*

In particular, (search or decision) $\mathsf{LWE}_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ is deterministic polynomial-time equivalent to (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{B},\mathcal{O},\psi,\ell}$, for any ℓ.

Remark 4.2. We stress that for efficiency, the reduction needs to be given all the prime $\mathcal{O}$-ideals that contain $\mathcal{Q}$. These are a subset of the prime ideals $\mathfrak{p}$ that "lie over" some prime integer divisor $p \in \mathbb{Z}$ of the index $Q = |\mathcal{O}/\mathcal{Q}| \in \mathcal{Q} \cap \mathbb{Z}$ (i.e., $\mathfrak{p} \cap \mathbb{Z} = p\mathbb{Z}$), which themselves may be easy to compute for certain choices of $\mathcal{Q}$ and $\mathcal{O}$. This is because for any prime $\mathcal{O}$-ideal $\mathfrak{p} \supseteq \mathcal{Q} \supseteq Q\mathcal{O}$, we have $\mathfrak{p} \cap \mathbb{Z} \supseteq Q\mathcal{O} \cap \mathbb{Z} = Q\mathbb{Z}$, so $\mathfrak{p}$ must lie over some prime divisor p of Q.

Proof. Since $\mathcal{A}$ is invertible modulo $\mathcal{Q}$, Lemma 2.13 says there exists a polynomial-time computable $t \in \mathcal{A}$ such that $t(\mathcal{O} : \mathcal{A}) + \mathcal{Q} = \mathcal{O}$. Then by Lemma 2.14, the function $\theta_t(u) = t \cdot u$ induces the $\mathcal{O}$-module isomorphisms $g: \mathcal{O}_\mathcal{Q} \to \mathcal{A}_\mathcal{Q}$ and $h: \mathcal{S}_\mathcal{Q} \to \mathcal{B}_\mathcal{Q}$, where in both cases we invoke the statement with $\mathcal{I} = \mathcal{O}$ and $\mathcal{I}' = \mathcal{A}$, and for g we let $\mathcal{J} = \mathcal{O}$ and for h we let $\mathcal{J} = \mathcal{S}$.

The claimed transform is as follows: for each input sample $(\vec{a}, \vec{b}) \in \mathcal{A}_\mathcal{Q}^{k_a} \times (K_\mathbb{R}/\mathcal{Q}\mathcal{B})^{k_b}$, we output

$$(\vec{a}' = g^{-1}(\vec{a}), \vec{b}' = \vec{b}) \in \mathcal{O}_\mathcal{Q}^{k_a} \times (K_\mathbb{R}/\mathcal{Q}\mathcal{B})^{k_b}$$

where g^{-1} is evaluated coordinate-wise on the vector $\vec{a}$. It is clear that this maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$, because g is a bijection. And obviously, the distribution of $\vec{b}'$ is identical to that of $\vec{b}$.

To complete the proof, it suffices to show that $T(\vec{s}, \vec{a}) = T(\vec{s}', \vec{a}')$, where $\vec{s}' = h(\vec{s})$. By linearity, it is enough to show that $s \cdot a = h(s)g^{-1}(a)$ for all $s \in \mathcal{S}_\mathcal{Q}$ and $a \in \mathcal{A}_\mathcal{Q}$. Note that $a = t \cdot g^{-1}(a) + \mathcal{Q}\mathcal{A}$ and $h(s) = t \cdot s + \mathcal{Q}\mathcal{B}$. Therefore,

$$\begin{aligned}
s \cdot a + \mathcal{Q}\mathcal{B} &= s \cdot (t \cdot g^{-1}(a) + \mathcal{Q}\mathcal{A}) + \mathcal{Q}\mathcal{B} \\
&= t \cdot s \cdot g^{-1}(a) + \mathcal{Q}\mathcal{B} \\
&= (t \cdot s + \mathcal{Q}\mathcal{B}) \cdot g^{-1}(a) + \mathcal{Q}\mathcal{B} \\
&= h(s) \cdot g^{-1}(a) + \mathcal{Q}\mathcal{B}.
\end{aligned}$$

For the claimed equivalences between LWE problems, simply apply the above transform or its inverse to each LWE sample. For the search problems, we may recover $\vec{s}$ from $\vec{s}'$, and vice versa, via h^{-1} or h, respectively. $\qquad\square$

Corollary 4.3. *Adopt the notation from Theorem 4.1, and let $\mathcal{S}', \mathcal{A}'$ be fractional $\mathcal{O}$-ideals with $\mathcal{S}'\mathcal{A}' = \mathcal{B} = \mathcal{S}\mathcal{A}$. If both $\mathcal{A}, \mathcal{A}'$ are invertible modulo $\mathcal{Q}$, and all the prime (equivalently, maximal) $\mathcal{O}$-ideals that contain $\mathcal{Q}$ are known, then (search or decision) $\mathsf{LWE}_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ is deterministic polynomial-time equivalent to (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{S}',\mathcal{A}',\psi,\ell}$, for any ℓ.*

Proof. By Theorem 4.1, both problems are deterministic polynomial-time equivalent to $\mathsf{LWE}_{T,\mathcal{B},\mathcal{O},\psi,\ell}$. $\qquad\square$

The next simple theorem shows that we can replace $\mathcal{S}$ and $\mathcal{B}$ with appropriately related super-ideals. Because this transformation may discard information by reducing modulo a "denser" lattice, it is typically not reversible, though under certain conditions it can be.

Theorem 4.4. *Let $\mathcal{O}$ be an order of a number field K; $\mathcal{S} \subseteq \mathcal{S}'$, $\mathcal{A}$, $\mathcal{B} = \mathcal{S}\mathcal{A} \subseteq \mathcal{B}' = \mathcal{S}'\mathcal{A}$ be fractional $\mathcal{O}$-ideals; $\mathcal{Q}$ be an $\mathcal{O}$-ideal; k_s, k_a, and k_b be positive integers; $T \in \mathcal{O}_{\mathcal{Q}}^{k_s \times k_a \times k_b}$ be an order-three tensor; and ψ be a distribution over $K_{\mathbb{R}}^{k_b}$. Then there is a polynomial-time deterministic transform that:*

1. *maps distribution $U_{T,\mathcal{S},\mathcal{A}}$ to distribution $U_{T,\mathcal{S}',\mathcal{A}}$, and*

2. *maps distribution $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ to distribution $A_{T,\mathcal{S}',\mathcal{A},\psi}(\vec{s}')$, where $\vec{s}' = \vec{s} \bmod \mathcal{Q}\mathcal{S}'$.*

In particular, there is a polynomial-time randomized reduction from decision-$\mathsf{LWE}_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ to decision-$\mathsf{LWE}_{T,\mathcal{S}',\mathcal{A},\psi,\ell}$, for any ℓ.

Moreover, if the natural inclusion map $\mathcal{S}_{\mathcal{Q}} \to \mathcal{S}'_{\mathcal{Q}}$ is a bijection (see, e.g., Corollary 2.15), then the reduction can be made deterministic*, and also extends to the* search *versions of the problems. In addition, if the natural inclusion map $\mathcal{B}_{\mathcal{Q}} \to \mathcal{B}'_{\mathcal{Q}}$ is a bijection and ψ is* discrete *over $\mathcal{B}^{k_b}$, then the above deterministic reductions are a polynomial-time* equivalences*.*

Proof. The claimed transform is as follows: for each input sample $(\vec{a}, \vec{b}) \in \mathcal{A}_{\mathcal{Q}}^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b}$, we output

$$(\vec{a}' = \vec{a} \, , \, \vec{b}' = \vec{b} \bmod (\mathcal{Q}\mathcal{B}')^{k_b}) \in \mathcal{A}_{\mathcal{Q}}^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B}')^{k_b}.$$

Trivially, this transform maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$. Also, since $\mathcal{Q}\mathcal{B} \subseteq \mathcal{Q}\mathcal{B}'$, the transform sends uniformly random $\vec{b}$ to uniformly random $\vec{b}'$.

To complete the proof, it suffices to show that $T(\vec{s}, \vec{a}) = T(\vec{s}', \vec{a}) \pmod{\mathcal{Q}\mathcal{B}'}$. By linearity, it suffices that $s \cdot a = s' \cdot a \pmod{\mathcal{Q}\mathcal{B}'}$ for all $s \in \mathcal{S}_{\mathcal{Q}}$ and $a \in \mathcal{A}_{\mathcal{Q}}$, where $s' = s \bmod \mathcal{Q}\mathcal{S}'$. This follows immediately from the fact that $a \in \mathcal{A}_{\mathcal{Q}}$ and $\mathcal{B}' = \mathcal{S}'\mathcal{A}$.

The claimed reductions follow immediately by applying the above transform to each LWE sample, and (if necessary for the decision-to-decision reduction) re-randomizing the secret $\vec{s}'$ in the standard way, by choosing a uniformly random $\vec{r}' \in (\mathcal{S}'_{\mathcal{Q}})^{k_s}$ and changing each sample $(\vec{a}', \vec{b}')$ to $(\vec{a}', \vec{b}' + T(\vec{r}', \vec{a}'))$. Moreover, when the natural inclusion map $\mathcal{S}_{\mathcal{Q}} \to \mathcal{S}'_{\mathcal{Q}}$ is a bijection, note that $\vec{s}'$ is uniform over $(\mathcal{S}'_{\mathcal{Q}})^{k_s}$ because $\vec{s}$ is uniform over $\mathcal{S}_{\mathcal{Q}}^{k_s}$, so the reduction is deterministic (re-randomization is not needed), and (for the search problems) we may recover $\vec{s}$ from $\vec{s}'$ simply by applying the inverse of the natural inclusion map. Finally, when additionally the natural inclusion map $\mathcal{B}_{\mathcal{Q}} \to \mathcal{B}'_{\mathcal{Q}}$ is bijection and ψ is discrete over $\mathcal{B}^{k_b}$, the $\vec{b}$ are over $\mathcal{B}_{\mathcal{Q}}$ so the above transform itself is efficiently invertible, hence the problems are equivalent. $\qquad\square$

4.2 Changing the Order

In this section we show a reduction (which in some cases is an equivalence) between generalized LWE problems defined over orders $\mathcal{O} \subseteq \mathcal{O}'$ of the same number field, for appropriately related ideals and tensors $T' = T \bmod \mathcal{Q}\mathcal{O}'$; the error distribution and the number of samples are preserved.

Theorem 4.5. *Let $\mathcal{O} \subseteq \mathcal{O}'$ be orders in a number field K; $\mathcal{S}$ and $\mathcal{A}$ be fractional $\mathcal{O}$-ideals with $\mathcal{B} = \mathcal{S}\mathcal{A}$; $\mathcal{S}' = \mathcal{O}'\mathcal{S}$, $\mathcal{A}' = \mathcal{O}'\mathcal{A}$, and $\mathcal{B}' = \mathcal{S}'\mathcal{A}' = \mathcal{O}'\mathcal{B}$; $\mathcal{Q}$ be an $\mathcal{O}$-ideal with $\mathcal{Q}' = \mathcal{O}'\mathcal{Q}$; $T \in \mathcal{O}_{\mathcal{Q}}^{k_s \times k_a \times k_b}$ be an order-three tensor for positive integers k_s, k_a, and k_b; and ψ be a distribution over $K_{\mathbb{R}}^{k_b}$. If $(\mathcal{O} : \mathcal{O}') + \mathcal{Q} = \mathcal{O}$, then there is a polynomial-time deterministic transform that:*

1. *maps distribution $U_{T,\mathcal{S},\mathcal{A}}$ to distribution $U_{T',\mathcal{S}',\mathcal{A}'}$, and*

2. *maps distribution $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ to distribution $A_{T',\mathcal{S}',\mathcal{A}',\psi}(\vec{s}')$ where $\vec{s}' = \vec{s} \bmod (\mathcal{Q}'\mathcal{S}')^{k_s}$,*

where $T' = T \bmod \mathcal{Q}'$, which is in $(\mathcal{O}'_{\mathcal{Q}'})^{k_s \times k_a \times k_b}$. Furthermore, if $\mathcal{B}' = \mathcal{B}$, or ψ is discrete over $\mathcal{B}^{k_b}$, then the transform is also polynomial-time invertible.

    In particular, there is a polynomial-time deterministic reduction from (search or decision) $\mathsf{LWE}_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ to (search or decision, respectively) $\mathsf{LWE}_{T',\mathcal{S}',\mathcal{A}',\psi,\ell}$ for any ℓ, and when the above transform is polynomial-time invertible, the problems are polynomial-time equivalent.

Proof. In all that follows, keep in mind that for any fractional $\mathcal{O}'$-ideal $\mathcal{I}'$, we have $\mathcal{Q}'\mathcal{I}' = \mathcal{Q}\mathcal{O}'\mathcal{I}' = \mathcal{Q}\mathcal{I}'$, and thus $\mathcal{I}'_{\mathcal{Q}'} = \mathcal{I}'_{\mathcal{Q}}$. By Corollary 2.15, the natural inclusion maps $\mathcal{O}_{\mathcal{Q}} \to \mathcal{O}'_{\mathcal{Q}'}$, $\mathcal{A}_{\mathcal{Q}} \to \mathcal{A}'_{\mathcal{Q}'}$, $\mathcal{S}_{\mathcal{Q}} \to \mathcal{S}'_{\mathcal{Q}'}$, $\mathcal{B}_{\mathcal{Q}} \to \mathcal{B}'_{\mathcal{Q}'}$ are bijections, where we invoke the statement with the order $\mathcal{O}$ and the fractional $\mathcal{O}$-ideals $\mathcal{I} = \mathcal{O}$ and $\mathcal{I}' = \mathcal{O}'$ in all cases, and with $\mathcal{J} = \mathcal{O}, \mathcal{A}, \mathcal{S}, \mathcal{B}$, respectively.

    The claimed transform is as follows: for each input sample $(\vec{a}, \vec{b}) \in \mathcal{A}_{\mathcal{Q}}^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b}$, we output

$$(\vec{a}' = \vec{a} \bmod (\mathcal{Q}'\mathcal{A}')^{k_a} , \vec{b}' = \vec{b} \bmod (\mathcal{Q}'\mathcal{B}')^{k_b}) \in (\mathcal{A}'_{\mathcal{Q}'})^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}'\mathcal{B}')^{k_b}.$$

It is clear that this transform maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$, since the natural inclusion map $\mathcal{A}_{\mathcal{Q}} \to \mathcal{A}'_{\mathcal{Q}'}$ is a bijection. Also, since $\mathcal{Q}\mathcal{B} \subseteq \mathcal{Q}'\mathcal{B}'$, the transform sends uniformly random $\vec{b}$ to uniformly random $\vec{b}'$.

    To complete the proof, it suffices to show that $T(\vec{s}, \vec{a}) = T'(\vec{s}', \vec{a}') \pmod{\mathcal{Q}'\mathcal{B}'}$. By linearity, it is enough to show that for all $s \in \mathcal{S}_{\mathcal{Q}}$, $r \in \mathcal{O}_{\mathcal{Q}}$ (representing an entry of T), and $a \in \mathcal{A}_{\mathcal{Q}}$, we have $s \cdot r \cdot a = s' \cdot r' \cdot a' \pmod{\mathcal{Q}'\mathcal{B}'}$ where $s' = s + \mathcal{Q}'\mathcal{S}'$, $r' = r + \mathcal{Q}'$, and $a' = a + \mathcal{Q}'\mathcal{A}'$. Indeed,

$$s \cdot r \cdot a + \mathcal{Q}'\mathcal{B}' = (s + \mathcal{Q}'\mathcal{S}') \cdot (r + \mathcal{Q}') \cdot (a + \mathcal{Q}'\mathcal{A}') + \mathcal{Q}'\mathcal{B}' = s' \cdot r' \cdot a' + \mathcal{Q}'\mathcal{B}'.$$

    To see that the transform is polynomial-time invertible (under one of the additional hypotheses), first recall that the natural inclusion map $\mathcal{A}_{\mathcal{Q}} \to \mathcal{A}'_{\mathcal{Q}'}$ is a bijection, and therefore efficiently invertible. Thus, it suffices to show that the modular reduction $\vec{b}' = \vec{b} \bmod (\mathcal{Q}'\mathcal{B}')^{k_b}$ is efficiently invertible. This is the case if $\mathcal{B}' = \mathcal{B}$, because $\mathcal{Q}'\mathcal{B}' = \mathcal{Q}\mathcal{O}'\mathcal{B}' = \mathcal{Q}\mathcal{B}' = \mathcal{Q}\mathcal{B}$ and hence $\vec{b}' = \vec{b}$. Alternatively, if ψ is discrete over $\mathcal{B}^{k_b}$, then each entry of $\vec{b}$ is mapped to the corresponding entry of $\vec{b}'$ by the natural inclusion map $\mathcal{B}_{\mathcal{Q}} \to \mathcal{B}'_{\mathcal{Q}'}$, which a bijection, as already noted.

    The claimed reductions between LWE problems follow immediately from the above transform and its inverse (when applicable), simply by applying them to each LWE sample. For the search problems, we can recover $\vec{s}$ from $\vec{s}'$ (or vice versa) using the inverse (or the forward direction, respectively) of the natural inclusion map $\mathcal{S}_{\mathcal{Q}} \to \mathcal{S}'_{\mathcal{Q}'}$. $\qquad\square$

Remark 4.6. Theorems 4.1 and 4.5 imply that the specific choices of order and ideals in the definition of $\mathcal{L}$-LWE (Definition 3.4) as an instantiation of generalized LWE are *canonical* ones, under mild conditions. Specifically, by Theorem 4.1, a generalized LWE instantiation over an order $\mathcal{O}$ is equivalent to one where $\mathcal{A} = \mathcal{O}$ and $\mathcal{S}$ is some fractional $\mathcal{O}$-ideal (assuming all the prime $\mathcal{O}$-ideals containing $\mathcal{Q}$ are known). Letting $\mathcal{O}' = \mathcal{O}^{\mathcal{S}} \supseteq \mathcal{O}$, by Theorem 4.5 this instantiation is equivalent to one over $\mathcal{O}'$ with $\mathcal{A}' = \mathcal{O}'\mathcal{A} = \mathcal{O}'$ and $\mathcal{S}' = \mathcal{O}'\mathcal{S} = \mathcal{S}$, which are exactly the choices made in $\mathcal{L}$-LWE with $\mathcal{L} = \mathcal{S}^{\vee}$. Note that a sufficient condition to satisfy the hypotheses of Theorems 4.1 and 4.5 is $(\mathcal{O} : \mathcal{O}_K) + \mathcal{Q} = \mathcal{O}$, because then any invertible $\mathcal{O}$-ideal is invertible modulo $\mathcal{Q}$ by Lemma 2.11, and because $(\mathcal{O} : \mathcal{O}_K) \subseteq (\mathcal{O} : \mathcal{O}')$.

4.3 Changing the Lattice in $\mathcal{L}$-LWE

As a straightforward corollary to Theorem 4.5 we get the following efficient, deterministic reduction from $\mathcal{L}'$-LWE to $\mathcal{L}$-LWE, under mild conditions on the lattices $\mathcal{L} \subseteq \mathcal{L}'$. Notice that, while the hypothesized containments of the lattices and their coefficient rings syntactically match those of the preceding results, the reduction actually goes in the *opposite* direction, from a "denser" order $\mathcal{O}' = \mathcal{O}^{\mathcal{L}'}$ to a "sparser" one $\mathcal{O} = \mathcal{O}^{\mathcal{L}} \subseteq \mathcal{O}'$. A main case of interest is when $\mathcal{L} = \mathcal{O}$ and $\mathcal{L}' = \mathcal{O}'$ are themselves orders; the theorem then says that hardness of Order-LWE over $\mathcal{O}'$—and in particular Ring-LWE, where $\mathcal{O}'$ is the full ring of integers—implies hardness over suborder $\mathcal{O}$.[17]

Theorem 4.7. *Let $\mathcal{L} \subseteq \mathcal{L}'$ be lattices in a number field K, $\mathcal{O} = \mathcal{O}^{\mathcal{L}}$ and $\mathcal{O}' = \mathcal{O}^{\mathcal{L}'}$ be their respective coefficient rings with $\mathcal{O} \subseteq \mathcal{O}'$, $\mathcal{Q}$ be an $\mathcal{O}$-ideal such that $(\mathcal{L} : \mathcal{L}') + \mathcal{Q} = \mathcal{O}$, $\mathcal{Q}' = \mathcal{O}'\mathcal{Q}$, ψ be a distribution over $K_{\mathbb{R}}$, and k be a positive integer. Then there is a polynomial-time deterministic reduction from (search or decision) $\mathcal{L}'\text{-LWE}^{k}_{\mathcal{Q}',\psi,\ell}$ to (search or decision, respectively) $\mathcal{L}\text{-LWE}^{k}_{\mathcal{Q},\psi,\ell}$, for any ℓ. Moreover, if ψ is discrete over $(\mathcal{L}')^{\vee}$, then the problems are polynomial-time equivalent.*

Proof. We proceed in two steps, by composing the "reverse" direction of Theorem 4.5 with Theorem 4.4.

First, $(\mathcal{O} : \mathcal{O}') + \mathcal{Q} = \mathcal{O}$ by Corollary 2.16. Invoke Theorem 4.5 with orders $\mathcal{O} \subseteq \mathcal{O}'$ and $\mathcal{A} = \mathcal{O}, \mathcal{S} = \mathcal{B} = (\mathcal{L}')^{\vee}$ (which is a fractional $\mathcal{O}$-ideal because it a fractional $\mathcal{O}'$-ideal), making $\mathcal{A}' = \mathcal{O}'\mathcal{A} = \mathcal{O}'$ and $\mathcal{S}' = \mathcal{O}'\mathcal{S} = \mathcal{S}$, and hence $\mathcal{B}' = \mathcal{S}' = \mathcal{S} = \mathcal{B}$. This yields a deterministic polynomial-time *equivalence* between (search or decision) $\mathcal{L}'\text{-LWE}^{k}_{\mathcal{Q}',\psi,\ell} = \text{LWE}_{T',(\mathcal{L}')^{\vee},\mathcal{O}',\psi,\ell}$ and (search or decision, respectively) $\text{LWE}^{k}_{T,(\mathcal{L}')^{\vee},\mathcal{O},\psi,\ell}$, where T' and T respectively correspond to k-by-k identity matrices over $\mathcal{O}'_{\mathcal{Q}'}$ and $\mathcal{O}_{\mathcal{Q}}$.

Next, because $(\mathcal{L} : \mathcal{L}') + \mathcal{Q} = \mathcal{O}$, by Corollary 2.15 the natural inclusion map $(\mathcal{L}')^{\vee}_{\mathcal{Q}} \to \mathcal{L}^{\vee}_{\mathcal{Q}}$ is a bijection. So, Theorem 4.4 gives a deterministic reduction from $\text{LWE}^{k}_{T,(\mathcal{L}')^{\vee},\mathcal{O},\psi,\ell}$ to (search or decision, respectively) $\text{LWE}_{T,\mathcal{L}^{\vee},\mathcal{O},\psi,\ell} = \mathcal{L}\text{-LWE}_{\mathcal{Q},\psi,\ell}$, as desired. Moreover, if ψ is discrete over $(\mathcal{L}')^{\vee}$, then again by Theorem 4.4 the problems are polynomial-time equivalent. □

4.4 Changing the Tensor

We now give a reduction from one generalized LWE problem to another, when their associated tensors are suitably related by multiplication with appropriately invertible matrices. In particular, the reduction is even error preserving when a certain one of those matrices is the identity.

Theorem 4.8. *Let*

---

[17]We remark that [RSW18, Theorem 4.2] proves a similar result for any order $\mathcal{O} \subseteq \mathcal{O}_K$, under the hypothesis $(\mathcal{O} : \mathcal{O}_K) + \mathcal{O}_K\mathcal{Q} = \mathcal{O}_K$ (among others). By [Con09, Theorem 3.8], this hypothesis is equivalent to ours for this choice of orders, so our result is strictly more general than the one of [RSW18].

- $\mathcal{O}$ be an order in a number field K, and $\mathcal{S}$ and $\mathcal{A}$ be fractional $\mathcal{O}$-ideals with $\mathcal{B} = \mathcal{S}\mathcal{A}$;

- $T \in \mathcal{O}_{\mathcal{Q}}^{k_s \times k_a \times k_b}$ and $T' \in \mathcal{O}_{\mathcal{Q}}^{k_s' \times k_a' \times k_b'}$ be tensors over $\mathcal{O}_{\mathcal{Q}}$ for positive integers $k_s, k_a, k_b, k_s', k_a', k_b'$ and $\mathcal{O}$-ideal $\mathcal{Q}$;

- $\mathbf{S} \in \mathcal{O}_{\mathcal{Q}}^{k_s' \times k_s}$, $\mathbf{A} \in \mathcal{O}_{\mathcal{Q}}^{k_a \times k_a'}$, and $\mathbf{B} \in \mathcal{O}^{k_b' \times k_b}$ be matrices where $\mathbf{A}$ and $\mathbf{B}$ are right invertible over $\mathcal{O}_{\mathcal{Q}}$ and K (respectively) and $\sum_{j\ell} T_{ij\ell}\, \mathbf{A}_{jj'}\, \mathbf{B}_{\ell'\ell} = \sum_{i'} T'_{i'j'\ell'}\, \mathbf{S}_{i'i}$ for all i, j', ℓ'; and

- ψ be a distribution over $K_{\mathbb{R}}^{k_b}$.

There is a polynomial-time randomized transform that:

1. *maps distribution $U_{T,\mathcal{S},\mathcal{A}}$ to distribution $U_{T',\mathcal{S},\mathcal{A}}$, and*

2. *maps the distribution $A_{T,\mathcal{S},\mathcal{A},\psi}(\vec{s})$ to $A_{T',\mathcal{S},\mathcal{A},\psi'}(\vec{s}')$, where $\psi' = \mathbf{B}\psi$ and $\vec{s}' = \mathbf{S}\vec{s}$.*

Furthermore, if $\mathbf{A}$ is square (and hence invertible), the transform is deterministic.

In particular, there is a polynomial-time randomized reduction from decision-$\mathsf{LWE}_{T,\mathcal{S},\mathcal{A},\psi,\ell}$ to decision-$\mathsf{LWE}_{T',\mathcal{S},\mathcal{A},\psi,\ell}$ for any ℓ, and similarly for the search problems if $\mathbf{S}$ is left invertible (over $\mathcal{O}_{\mathcal{Q}}$). Furthermore, when $\mathbf{A}$ is square, there is a deterministic decision-to-decision reduction if $\mathbf{S}$ is right invertible, and similarly for search-to-search if $\mathbf{S}$ is invertible.

Proof. First, let $\nu \colon \mathcal{A}_{\mathcal{Q}}^{k_a'} \to \mathcal{A}_{\mathcal{Q}}^{k_a}$ be defined by $\nu(\vec{a}') = \mathbf{A}\vec{a}'$. This map is surjective because $\mathbf{A}$ is right invertible, and a basis of its kernel (which is a finite group) can be computed in polynomial time using the techniques referenced at the start of [Section 2](). Therefore, we can efficiently sample uniformly from $\nu^{-1}(\vec{a})$, and such a sample is uniformly random over $\mathcal{A}_{\mathcal{Q}}^{k_a'}$ when $\vec{a}$ is uniformly random over $\mathcal{A}_{\mathcal{Q}}^{k_a}$. Furthermore, notice that when $\mathbf{A}$ is square, $\nu^{-1}(\vec{a})$ is unique, so we can deterministically sample from it.

The claimed transform is as follows: for each input sample $(\vec{a}, \vec{b}) \in \mathcal{A}_{\mathcal{Q}}^{k_a} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b}$, we output

$$(\vec{a}' \leftarrow \nu^{-1}(\vec{a})\, ,\ \vec{b}' = \mathbf{B}\vec{b}) \in \mathcal{A}_{\mathcal{Q}}^{k_a'} \times (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b'}.$$

As already observed, this maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$. It also maps uniformly random $\vec{b} \in (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b}$ to uniformly random $\vec{b}' \in (K_{\mathbb{R}}/\mathcal{Q}\mathcal{B})^{k_b'}$, because multiplication by $\mathbf{B}$ is a surjective map from $K_{\mathbb{R}}^{k_b}$ to $K_{\mathbb{R}}^{k_b'}$ (since $\mathbf{B}$ is right invertible over K), and $\mathbf{B}(\mathcal{Q}\mathcal{B})^{k_b} = (\mathcal{Q}\mathcal{B})^{k_b'}$.

It remains to show that if $\vec{b} = T(\vec{s}, \vec{a}) + \vec{e} \bmod (\mathcal{Q}\mathcal{B})^{k_b}$ for some $\vec{e} \leftarrow \psi$, then $\vec{b}' = T'(\vec{s}', \vec{a}') + \vec{e}' \bmod (\mathcal{Q}\mathcal{B})^{k_b'}$ where $\vec{e}' = \mathbf{B}\vec{e}$. To see this, observe that for any index ℓ' of $\vec{b}'$,

$$
\begin{aligned}
b'_{\ell'} = \sum_{\ell} \mathbf{B}_{\ell'\ell} b_\ell &= \sum_{i,j,\ell} \mathbf{B}_{\ell'\ell}(T_{ij\ell}\, s_i\, a_j + e_\ell) \\
&= \sum_{i,j,j',\ell} \mathbf{B}_{\ell'\ell}\, T_{ij\ell}\, s_i\, \mathbf{A}_{jj'}\, a'_{j'} + e'_{\ell'} \\
&= \sum_{i,i',j'} T'_{i'j'\ell'}\, \mathbf{S}_{i'i}\, s_i\, a'_{j'} + e'_{\ell'} \\
&= \sum_{i',j'} T'_{i'j'\ell'}\, s'_{i'} a'_{j'} + e'_{\ell'} \\
&= T'(\vec{s}', \vec{a}')_{\ell'} + e'_{\ell'},
\end{aligned}
$$

as desired.

For the claimed reductions, note that it may not suffice to simply apply the claimed transformation to each input sample: while $\vec{s}'$ is uniformly distributed when $\mathbf{S}$ is right invertible, it may not be otherwise. This is easily addressed by the standard technique of re-randomizing the secret, choosing a uniformly random $\vec{r}' \in \mathcal{S}_{\mathcal{Q}}^{k'_s}$ and transforming each sample $(\vec{a}', \vec{b}')$ to $(\vec{a}', \vec{b}' + T'(\vec{r}', \vec{a}'))$. This preserves the uniform distribution, and for LWE samples it maps any secret $\vec{s}'$ to a uniformly random secret $\vec{s}' + \vec{r}'$.

The above establishes the claimed reductions between the decision problems. For the claimed reductions between the search problems, apply the above transform, and given the secret $\vec{s}'$ for the resulting samples, simply compute the original secret as $\vec{s} = \mathbf{S}^+ \vec{s}'$, where $\mathbf{S}^+$ is a left inverse of $\mathbf{S}$. $\qquad\square$

Remark 4.9. Theorem 4.8 can be used to reshape the error distribution in a generalized LWE problem by a factor of any $t \in \mathcal{O}$ that is invertible modulo $\mathcal{Q}$, i.e., for which there exists some $u \in \mathcal{O}$ such that $t \cdot u = 1 \bmod \mathcal{Q}$. This is equivalent to the condition that $t\mathcal{O}$ and $\mathcal{Q}$ are coprime, i.e., $t\mathcal{O} + \mathcal{Q} = \mathcal{O}$. In this case, we can set $\mathbf{B} = t\mathbf{I}$, $\mathbf{A} = u\mathbf{I}$, and $\mathbf{S} = \mathbf{I}$ to obtain a reduction between generalized LWE problems in which the error distribution is multiplied by t, and all the other parameters are unchanged.

4.5 Changing the Number Field

In this section we show an equivalence between generalized LWE problems over number fields $K \subseteq K'$ and orders $\mathcal{O} \subset K, \mathcal{O}' \subset K'$ of different degrees, for suitably related ideals and tensors. Despite the rather technical nature of the theorem statement, the core idea is fairly straightforward. Let T' be an order-three LWE tensor over $\mathcal{O}'_{\mathcal{Q}'}$ where $\mathcal{Q}' = \mathcal{O}'\mathcal{Q}$ for some $\mathcal{O}$-ideal $\mathcal{Q}$, and recall that it represents some bilinear map on a secret and a multiplier. Furthermore, suppose that the input and output domains of T' (as a bilinear map) are free $\mathcal{O}$-modules with some known bases. Then T' is equivalent to a larger order-three LWE tensor T over $\mathcal{O}_{\mathcal{Q}}$, simply by representing the inputs and output in their respective $\mathcal{O}$-bases, and replacing each entry of T' with the block representing multiplication by that entry (also relative to those $\mathcal{O}$-bases). Formally, this expansion is done by letting T be the entry-wise trace of the order-three Kronecker product $T' \otimes C$, where C is the Kronecker product of the input bases and the dual of the output basis.

Theorem 4.10. *Let*

- *K'/K be a k-dimensional number field extension with $\mathrm{Tr} = \mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}$ (which, to recall, coincides with $\mathrm{Tr}_{K'/K}$ on K'), with $\mathcal{O}$ an order of K and $\mathcal{O}'$ an order of K';*

- *$\mathcal{Q}$ be an $\mathcal{O}$-ideal with $\mathcal{Q}' = \mathcal{O}'\mathcal{Q}$;*

- *$\mathcal{M}'_s, \mathcal{M}'_a, \mathcal{M}'_b = \mathcal{M}'_s\mathcal{M}'_a$ be fractional $\mathcal{O}'$-ideals that are also rank-k free $\mathcal{O}$-modules with respective known $\mathcal{O}$-bases $\vec{b}_s, \vec{b}_a, \vec{b}_b$ (which hence are K-bases of K'), each with index set $[k]$;*

- *$\mathcal{S}, \mathcal{A}, \mathcal{B} = \mathcal{S}\mathcal{A}$ be fractional $\mathcal{O}$-ideals and $\mathcal{S}' = \mathcal{S}\mathcal{M}'_s$, $\mathcal{A}' = \mathcal{A}\mathcal{M}'_a$, $\mathcal{B}' = \mathcal{B}\mathcal{M}'_b = \mathcal{S}'\mathcal{A}'$ be the corresponding fractional $\mathcal{O}'$-ideals;*

- *k'_s, k'_a, k'_b be positive integers, with $k_s = k'_s \cdot k$, $k_a = k'_a \cdot k$, and $k_b = k'_b \cdot k$;*

- *$T' \in (\mathcal{O}'_{\mathcal{Q}'})^{k'_s \times k'_a \times k'_b}$ be an order-three tensor over $\mathcal{O}'_{\mathcal{Q}'}$ and $T = \mathrm{Tr}(T' \otimes C) \in \mathcal{O}_{\mathcal{Q}}^{k_s \times k_a \times k_b}$ be the entry-wise trace of the order-three tensor $T' \otimes C$ having index set $[k'_s \cdot k] \times [k'_a \cdot k] \times [k'_b \cdot k] = [k_s] \times [k_a] \times [k_b]$ (see Section 2.1), where $C = \vec{b}_s \otimes \vec{b}_a \otimes \vec{b}_b^{\vee}$;[18] and*

- *ψ' be a distribution over $(K'_{\mathbb{R}})^{k'_b}$.*

---

[18]Note that $T \in \mathcal{O}_{\mathcal{Q}}^{k_s \times k_a \times k_b}$ because $\vec{b}_b^{\vee}$ is an $\mathcal{O}$-basis of $(\mathcal{M}'_b)^{\vee\mathcal{O}}$, and hence each element of $T = \mathrm{Tr}(T' \otimes C)$ is in $\mathrm{Tr}(\mathcal{O}'_{\mathcal{Q}'}\mathcal{M}'_s\mathcal{M}'_a(\mathcal{M}'_b)^{\vee\mathcal{O}}) \subseteq \mathrm{Tr}(\mathcal{J}_{\mathcal{Q}'}) \subseteq \mathcal{O}_{\mathcal{Q}}$, where $\mathcal{J} = \mathcal{M}'_b(\mathcal{M}'_b)^{\vee\mathcal{O}}$.

Then there is a polynomial-time computable and invertible transform that:

1. *maps distribution $U_{T',S',\mathcal{A}'}$ to $U_{T,S,\mathcal{A}}$, and*

2. *maps distribution $A_{T',S',\mathcal{A}',\psi'}(\vec{s}')$ to $A_{T,S,\mathcal{A},\psi}(\vec{s})$, for $\vec{s} = \mathrm{Tr}(\vec{s}' \otimes \vec{b}_s^\vee)$ and $\psi = \mathrm{Tr}(\psi' \otimes \vec{b}_b^\vee)$.*

In particular, (search or decision) $\mathsf{LWE}_{T',S',\mathcal{A}',\psi',\ell}$ is polynomial-time equivalent to (search or decision, respectively) $\mathsf{LWE}_{T,S,\mathcal{A},\psi,\ell}$, for any ℓ.

Proof. We begin by showing that the bases $\vec{b}_s, \vec{b}_a, \vec{b}_b$ yield a number of polynomial-time computable bijections, which we utilize throughout the proof. We show this in detail for $\vec{b}_s$; the reasoning for $\vec{b}_a$ and $\vec{b}_b$ is similar.

Note that because $\vec{b}_s$ is an $\mathcal{O}$-basis of a lattice in K', it is also a K-basis of K', and a $K_\mathbb{R}$-basis of $K_\mathbb{R}'$. In particular, the function $\varphi\colon K_\mathbb{R}' \to K_\mathbb{R}^k$ defined by $\varphi(x) = \mathrm{Tr}(x \cdot \vec{b}_s^\vee)$ is a bijection with inverse $\vec{x} \mapsto \langle \vec{x}, \vec{b}_s \rangle$, and these also induce bijections between K' and K^k. Furthermore, since $S' = S\mathcal{M}_s'$, it follows that φ maps surjectively (and hence bijectively) from S' to S^k: for if $x \in S'$, then

$$\mathrm{Tr}(x \cdot \vec{b}_s^\vee) \in S \cdot \mathrm{Tr}(\mathcal{M}_s'\vec{b}_s^\vee) = S \cdot \mathcal{O}^k = S^k,$$

and if $\vec{x} \in S^k$, then $\varphi^{-1}(\vec{x}) = \langle \vec{x}, \vec{b}_s \rangle \in S\mathcal{M}_s' = S'$. Similarly, φ maps bijectively from $Q'S' = QS'$ to $(QS)^k$, simply by K-linearity of φ.

The claimed transform is as follows: given a sample $(\vec{a}', \vec{b}') \in (\mathcal{A}_{Q'}')^{k_a'} \times (K_\mathbb{R}'/Q'\mathcal{B}')^{k_b'}$, we output

$$(\vec{a} = \mathrm{Tr}(\vec{a}' \otimes \vec{b}_a^\vee), \, \vec{b} = \mathrm{Tr}(\vec{b}' \otimes \vec{b}_b^\vee)) \in \mathcal{A}_Q^{k_a} \times (K_\mathbb{R}/Q\mathcal{B})^{k_b}.$$

By what we showed in the previous paragraph, $\mathrm{Tr}(\vec{a}' \otimes \vec{b}_a^\vee)$ simply extracts the coordinate vector (relative to $\vec{b}_a$) of each entry of $\vec{a}'$, and similarly for $\mathrm{Tr}(\vec{b}' \otimes \vec{b}_b^\vee)$. Therefore, the maps from $\vec{a}'$ to $\vec{a}$, and from $\vec{b}'$ to $\vec{b}$, are bijections between their respective domains, and hence preserve the corresponding uniform distributions.

To establish the second part of the claim, it suffices to show that $\mathrm{Tr}(T'(\vec{s}', \vec{a}') \otimes \vec{b}_b^\vee) = T(\vec{s}, \vec{a})$, where recall that $\vec{s} = \mathrm{Tr}(\vec{s}' \otimes \vec{b}_s^\vee)$. Also recall that $T = \mathrm{Tr}(T' \otimes C)$ where $C = \vec{b}_s \otimes \vec{b}_a \otimes \vec{b}_b^\vee$, and that we can index T by $((i', i), (j', j), (\ell', \ell))$ where (i', j', ℓ') is an index of T', and $i, j, \ell \in [k]$ are respectively indices of $\vec{b}_s, \vec{b}_a, \vec{b}_b^\vee$. By definition, we have

$$T_{(i',i)(j',j)(\ell',\ell)} = \mathrm{Tr}(T'_{i'j'\ell'} \cdot (\vec{b}_s)_i (\vec{b}_a)_j (\vec{b}_b^\vee)_\ell).$$

Also, $\vec{s}$ and $\vec{a}$ are indexed in a similar way, where

$$\vec{s}_{(i',i)} = \mathrm{Tr}(\vec{s}'_{i'} \cdot (\vec{b}_s^\vee)_i) \quad \text{and} \quad \vec{a}_{(j',j)} = \mathrm{Tr}(\vec{a}'_{j'} \cdot (\vec{b}_a^\vee)_j).$$

Therefore, by K-linearity of Tr and the definition of the dual basis, for any index (ℓ', ℓ) we have

$$
\begin{aligned}
T(\vec{s}, \vec{a})_{(\ell',\ell)} &= \sum_{i',i,j',j} T_{(i',i)(j',j)(\ell',\ell)} \cdot \vec{s}_{(i',i)} \cdot \vec{a}_{(j',j)} \\
&= \sum_{i',i,j',j} \mathrm{Tr}\Big(T'_{i'j'\ell'} \cdot (\vec{b}_s)_i (\vec{b}_a)_j (\vec{b}_b^\vee)_\ell \cdot \vec{s}_{(i',i)} \cdot \vec{a}_{(j',j)}\Big) \\
&= \sum_{i'j'} \mathrm{Tr}\Big(T'_{i'j'\ell'} \Big(\sum_i (\vec{b}_s)_i \, \mathrm{Tr}(\vec{s}'_{i'} \cdot (\vec{b}_s^\vee)_i)\Big) \Big(\sum_j (\vec{b}_a)_j \, \mathrm{Tr}(\vec{a}'_{j'} \cdot (\vec{b}_a^\vee)_j)\Big) (\vec{b}_b^\vee)_\ell\Big) \\
&= \sum_{i'j'} \mathrm{Tr}\big(T'_{i'j'\ell'} \cdot \vec{s}'_{i'} \cdot \vec{a}'_{j'} \cdot (\vec{b}_b^\vee)_\ell\big) \\
&= \mathrm{Tr}\big(T'(\vec{s}', \vec{a}')_{\ell'} \cdot (\vec{b}_b^\vee)_\ell\big) = \mathrm{Tr}\big(T'(\vec{s}', \vec{a}') \otimes \vec{b}_b^\vee\big)_{(\ell',\ell)},
\end{aligned}
$$

as desired. Finally, the transform is polynomial-time invertible because the maps applied to $\vec{a}'$ and $\vec{b}'$ are, by taking linear combinations with $\vec{b}_a$ and $\vec{b}_b$, respectively.

For the claimed equivalences between LWE problems, simply apply the above transform or its inverse to each sample. For the search problems, we may efficiently recover $\vec{s}$ from $\vec{s}'$, or vice versa, via the bijection φ or its inverse (from the start of the proof). $\qquad\square$

5 Hardness of Middle-Product LWE

Rosca *et al.* [RSSS17] introduced the Middle-Product LWE (MP-LWE) problem and gave a hardness theorem for it, by showing a reduction from a wide class of Poly-LWE instantiations—and by extension, Ring-LWE instantiations [RSW18]—over various polynomial rings of the form $\mathbb{Z}[\alpha] \cong \mathbb{Z}[x]/f(x)$ for $f(x)$ satisfying mild conditions.

In this section we give a reduction that, when combined with our reduction from Theorem 4.7, subsumes the prior Ring-/MP-LWE connection in the simplicity of its descriptions and analysis, and in its error expansion and distortion (see Figure 1). These advantages arise from our use of $\mathcal{O}$-LWE as an intermediate problem, and in particular its use of dual lattices, in contrast to the entirely "primal" definition of Poly-LWE.

5.1 Middle-Product LWE

Middle-Product LWE can be seen as an instance of generalized LWE, as follows. The d-middle-product operation takes two polynomials of certain degree bounds, multiplies them together, and outputs only the "middle" d coefficients of the product. More specifically, the product of two polynomials respectively having degrees $< n + d - 1$ and $< n$ has degree $< 2n + d - 2$; the middle-product discards the lowest and highest $n - 1$ coefficients, and outputs the remaining d coefficients. Middle-Product LWE is concerned with random noisy middle products with a secret polynomial over $\mathbb{Z}_q$, for an integer modulus q.

To see this as an instantiation of generalized LWE, take the trivial number field $K = \mathbb{Q}$ with its unique order $\mathcal{O} = \mathbb{Z}$, and take ideals $\mathcal{S} = \mathcal{A} = \mathcal{B} = \mathbb{Z}$. Let $k_s = n + d - 1$ and $k_a = n$, and respectively identify $\mathcal{S}_q^{k_s} = \mathbb{Z}_q^{n+d-1}$ and $\mathcal{A}_q^n = \mathbb{Z}_q^n$ with $\mathbb{Z}_q^{<n+d-1}[x]$ and $\mathbb{Z}_q^{<n}[x]$ (the $\mathbb{Z}_q$-modules of polynomials of degrees $< n + d - 1$ and $< n$, respectively), via the bases $\vec{s} = (1, x, \ldots, x^{n+d-2})$ and $\vec{a} = (x^{n-1}, x^{n-2}, \ldots, 1)$. (Basis $\vec{a}$ is in decreasing order by degree for reasons that will become clear shortly.) Finally, let $k_b = d$ and identify $\mathcal{B}_q^{k_b} = \mathbb{Z}_q^d$ with $x^{n-1} \cdot \mathbb{Z}_q^{<d}[x]$ via the basis $\vec{b} = (x^{n-1}, x^n, \ldots, x^{n+d-2})$.

The middle product is a $\mathbb{Z}_q$-bilinear form $M \colon \mathbb{Z}_q^{k_s} \times \mathbb{Z}_q^{k_a} \to \mathbb{Z}_q^{k_b}$ that is represented by the order-three tensor M (which is indexed from zero in all dimensions) defined by

$$M_{ij\ell} = \begin{cases} 1 & \text{if } i = j + \ell \\ 0 & \text{otherwise.} \end{cases} \tag{5.1}$$

This is because $s_i \cdot a_j = x^i \cdot x^{n-1-j} = x^{(n-1)+(i-j)}$, which equals b_{i-j} if $0 \le i - j < d$, and vanishes under the middle product otherwise. Therefore, the "slice" matrix $M_{i\cdot\cdot}$ (obtained by fixing the i coordinate) is the $n \times d$ rectangular Hankel matrix defined by the standard basis vector $\mathbf{e}_i \in \mathbb{Z}^{n+d-1}$, which is 1 in the ith coordinate and zero elsewhere (again indexing from zero).[19] Importantly, these $M_{i\cdot\cdot}$ slices form the standard basis of all $n \times d$ Hankel matrices, so we refer to M as the "Hankel tensor." With these observations, Middle-Product LWE is simply the following instantiation of generalized LWE.

---

[19] Recall that a matrix H is Hankel if each entry $H_{j\ell}$ is determined by $j + \ell$ (equivalently, it is an "upside down" Toeplitz matrix). So, an $n \times d$ Hankel matrix is defined by an $(n + d - 1)$-dimensional vector whose ith entry defines the entries $H_{j\ell}$ for $i = j + \ell$.

Definition 5.1 (MP-LWE **problem**). Let n, d, q be positive integers and ψ be a distribution over $\mathbb{R}^d$. The (search or decision) MP-LWE$_{n,d,q,\psi,\ell}$ problem is simply the (search or decision, respectively) LWE$_{M,\mathbb{Z},\mathbb{Z},\psi,\ell}$ problem, where M is the order-three tensor from Equation (5.1).

The above definition is specialized to the (unique) order $\mathbb{Z}$ of the rationals $\mathbb{Q}$, to match the definition given in [RSSS17]. However, it and Theorem 5.3 below generalize straightforwardly to other orders $\mathcal{O}$ of number fields K; see Remark 5.4.

We point out that MP-LWE becomes no easier as d decreases (and the corresponding final coordinate(s) of the error distribution are truncated), because the degree-$(n + d - 2)$ monomial of the secret can affect only the monomial of the same degree in the middle product. Therefore, dropping the latter just has the effect of dropping the former. In the tensor M, this corresponds to removing the "slices" $M_{(n+d-2)..}$ and $M_{..(d-1)}$, which yields the tensor for parameters n and $d - 1$.

5.2 Reduction

We start by recalling the notion of a power basis.

Definition 5.2. For an order $\mathcal{O}$, a *power* basis is a basis of the form $\vec{p} = (x^j)_{i \in [d]} = (1, x, x^2, \ldots, x^{d-1})$ for some $x \in \mathcal{O}$.

Theorem 5.3. *Let $d \leq n$ be positive integers, $\mathcal{O}$ be an order of a degree-d number field K having a known power $\mathbb{Z}$-basis $\vec{p} = (x^j)_{j \in [d]}$, ψ be a distribution over $K_{\mathbb{R}}$, and q be a positive integer modulus. There is a polynomial-time randomized reduction from (search or decision) $\mathcal{O}$-LWE$_{q,\psi,\ell}$ to (search or decision, respectively) MP-LWE$_{n,d,q,\psi',\ell}$, where $\psi' = \mathrm{Tr}_{K_{\mathbb{R}}/\mathbb{R}}(\psi \cdot \vec{p})$.*

Proof. The reader may wish to focus first on the special case $d = n$, in which case the matrix $\mathbf{A}$ constructed below is the identity matrix, and can be ignored.

First, because $\vec{p}$ is a $\mathbb{Z}$-basis of $\mathcal{O}$ and $\vec{p}^{\vee}$ is a $\mathbb{Z}$-basis of $\mathcal{O}^{\vee} = \mathcal{O} \cdot \mathcal{O}^{\vee}$, by Theorem 4.10 we have a polynomial-time deterministic reduction from (search or decision) $\mathcal{O}$-LWE$_{q,\psi,\ell}$ = LWE$_{1,\mathcal{O}^{\vee},\mathcal{O},q,\psi,\ell}$ to (search or decision, respectively) LWE$_{T,\mathbb{Z},\mathbb{Z},\psi',\ell}$, where $T_{ij\ell} = \mathrm{Tr}(p_i^{\vee} p_j p_{\ell}) \in \mathbb{Z}_q^{d \times d \times d}$ for $\mathrm{Tr} = \mathrm{Tr}_{K/\mathbb{Q}}$. Observe that each "slice" $T_{i..}$ is a $d \times d$ Hankel matrix, because $p_j p_{\ell} = x^{j+\ell}$ depends only on $j + \ell$. As we show next, this is the key property allowing us to relate T to M.

To complete the reduction to MP-LWE, we use Theorem 4.8 by exhibiting a suitable relationship between the $d \times d \times d$ tensor T and the $(n + d - 1) \times n \times d$ middle-product tensor M. Specifically, we will show that for suitable matrices $\mathbf{A} \in \mathbb{Z}_q^{d \times n}, \mathbf{S} \in \mathbb{Z}_q^{(n+d-1) \times d}$ (and $\mathbf{B} \in \mathbb{Z}^{d \times d}$ being the identity matrix),

$$\sum_j T_{ij\ell'} \mathbf{A}_{jj'} = \sum_{i'} M_{i'j'\ell'} \mathbf{S}_{i'i} . \tag{5.2}$$

First, extend the power $\mathbb{Z}$-basis $\vec{p} = (x^j)_{j \in [d]}$ of $\mathcal{O}$ to $\vec{p}' = (x^{j'})_{j' \in [n+d-1]}$, by including more powers of x if necessary. Define the matrix $\mathbf{A}_{jj'} = \mathrm{Tr}(p_j^{\vee} \cdot p'_{j'}) \in \mathbb{Z}_q^{d \times n}$, which is right-invertible because its left-most d

columns form the $d \times d$ identity matrix. Now, the left-hand side of Equation (5.2) is

$$
\begin{aligned}
T'_{ij'\ell'} &:= \sum_j T_{ij\ell'} \mathbf{A}_{jj'} \\
&= \sum_j \mathrm{Tr}(p_i^\vee p_j p_{\ell'}) \cdot \mathrm{Tr}(p_j^\vee \cdot p'_{j'}) && \text{(definition of } T \text{ and } \mathbf{A}) \\
&= \mathrm{Tr}\Big(p_i^\vee \cdot \big(\sum_j p_j \, \mathrm{Tr}(p_j^\vee \cdot p'_{j'})\big) \cdot p_{\ell'}\Big) && (\mathbb{Q}\text{-linearity of } \mathrm{Tr}\colon K \to \mathbb{Q}) \\
&= \mathrm{Tr}(p_i^\vee \cdot p'_{j'} \cdot p_{\ell'}) = \mathrm{Tr}(p_i^\vee \cdot x^{j'+\ell'}). && \text{(duality, definition of } \vec{p} \text{ and } \vec{p}')
\end{aligned}
$$

Observe that each "slice" $T'_{i..}$ is an $n \times d$ Hankel matrix. So, we can factor T' as the middle-product tensor M times a suitable matrix $\mathbf{S}$: each slice $T'_{i..}$ can be written as an (efficiently computable) $\mathbb{Z}_q$-linear combination of the slices $M_{i..}$, because these latter slices form the standard basis for the $n \times d$ Hankel matrices over $\mathbb{Z}_q$. More formally, defining $\mathbf{S} \in \mathbb{Z}_q^{(n+d-1)\times d}$ by $\mathbf{S}_{i'i} = \mathrm{Tr}(p'_{i'} \cdot p_i^\vee)$, by definition of M we have $T'_{ij'\ell'} = \sum_{i'} M_{i'j'\ell'} \mathbf{S}_{i'i}$ for all i. Furthermore, $\mathbf{S}$ is left-invertible, since its first d rows form the $d \times d$ identity matrix.

To conclude, we have satisfied Equation (5.2) with suitable $\mathbf{A}$ and $\mathbf{S}$, and hence Theorem 4.8 yields a polynomial-time randomized reduction from (search or decision) $\mathsf{LWE}_{T,\mathbb{Z},\mathbb{Z},\psi',\ell}$ to (search or decision, respectively) $\mathsf{LWE}_{M,\mathbb{Z},\mathbb{Z},\psi',\ell} = \mathsf{MP\text{-}LWE}_{n,d,q,\psi',\ell}$, as claimed. $\qquad\square$

Remark 5.4. Our definition of MP-LWE as an instantiation of generalized LWE, along with Theorem 5.3, naturally generalize to orders $\mathcal{O} \neq \mathbb{Z}$ of number fields $K \neq \mathbb{Q}$. Specifically, we define $\mathsf{MP\text{-}LWE}_{\mathcal{O},n,d,\mathcal{Q},\psi,\ell}$ as $\mathsf{LWE}_{M,\mathcal{O}^\vee,\mathcal{O},\psi,\ell}$, where $M \in \mathcal{O}_{\mathcal{Q}}^{(n+d-1)\times n \times d}$ is the order-three Hankel tensor whose entries are given by Equation (5.1).

Letting K'/K be a degree-d number field extension and $\mathcal{O}'$ be an order of K' that has a power $\mathcal{O}$-basis $\vec{p}$, notice that $\mathcal{O}'$ is a free $\mathcal{O}$-module of rank d with basis $\vec{p}$, and $(\mathcal{O}')^\vee$ is a free $\mathcal{O}$-module of rank d with basis $\vec{p}^\vee$. So, $\mathcal{O}' = \mathcal{O}\mathcal{O}'$ and $(\mathcal{O}')^\vee = (\mathcal{O}')^{\vee_\mathcal{O}}\mathcal{O}^\vee$ by Lemma 2.18 (for the tower $K'/K/\mathbb{Q}$ and orders $\mathcal{O}', \mathcal{O}, \mathbb{Z}$). Therefore, the statement and proof of Theorem 5.3 straightforwardly generalize: first, Theorem 4.10 with $\mathcal{M}'_a = \mathcal{O}'$ and $\mathcal{M}'_s = (\mathcal{O}')^{\vee_\mathcal{O}}$ gives a polynomial-time deterministic reduction from $\mathcal{O}'\text{-}\mathsf{LWE}_{\mathcal{Q},\psi,\ell} = \mathsf{LWE}_{1,(\mathcal{O}')^\vee,\mathcal{O}',\mathcal{Q},\psi',\ell}$ to $\mathsf{LWE}_{T,\mathcal{O}^\vee,\mathcal{O},\psi,\ell}$ where $T_{ij\ell} = \mathrm{Tr}_{K'/K}(p_i^\vee p_j p_\ell) \in \mathcal{O}_{\mathcal{Q}}^{d\times d\times d}$ and $\psi = \mathrm{Tr}_{K'_\mathbb{R}/K_\mathbb{R}}(\psi' \cdot \vec{p})$. Then, Theorem 4.8 gives a polynomial-time randomized reduction to $\mathsf{MP\text{-}LWE}_{\mathcal{O},n,d,\mathcal{Q},\psi,\ell}$.

5.3 Managing the Error Distribution

The reduction described in Theorem 5.3 reduces $\mathcal{O}$-LWE with error distribution ψ to MP-LWE with error distribution $\psi' = \mathrm{Tr}_{K_\mathbb{R}/\mathbb{R}}(\psi \cdot \vec{p})$ over $\mathbb{R}^d$, where $\vec{p}$ is some power $\mathbb{Z}$-basis of $\mathcal{O}$. However, we ultimately want a reduction from *many* $\mathcal{O}$-LWE problems to a *single* MP-LWE problem, so we need to further control the resulting error distribution. To this end, we consider the usual case where ψ is a Gaussian distribution over $K_\mathbb{R}$, in which case it turns out that ψ' is a Gaussian whose covariance is related to the Gram matrix of $\vec{p}$. Moreover, by a standard technique we can add some independent Gaussian error having a compensating covariance to arrive at any desired target covariance that is sufficiently large.

Throughout this section, we use the following notation. Let $\mathrm{Tr} = \mathrm{Tr}_{K_\mathbb{R}/\mathbb{R}}$, and given a basis $\vec{p}$ of $\mathcal{O}$, let $\mathbf{P} = \mathrm{Tr}(\vec{p} \cdot \tau(\vec{p})^t)$ denote the (positive definite) Gram matrix of $\vec{p}$, whose (i,j)th entry is $\langle p_i, p_j \rangle = \mathrm{Tr}(p_i \cdot \tau(p_j))$. Fix some orthonormal $\mathbb{R}$-basis $\vec{b} = \tau(\vec{b}^\vee)$ of $K_\mathbb{R}$, and let $\mathbf{P}_b = \mathrm{Tr}(\vec{b} \cdot \vec{p}^t)$. Then by $\mathbb{R}$-linearity of τ and trace, we have

$$
\mathbf{P} = \mathrm{Tr}(\vec{p} \cdot \tau(\vec{p})^t) = \mathrm{Tr}\Big(\vec{p} \cdot \tau\big((\vec{b}^\vee)^t \cdot \mathrm{Tr}(\vec{b} \cdot \vec{p}^t)\big)\Big) = \mathrm{Tr}(\vec{p} \cdot \vec{b}^t) \cdot \mathrm{Tr}(\vec{b} \cdot \vec{p}^t) = \mathbf{P}_b^t \cdot \mathbf{P}_b .
$$

For a real matrix $\mathbf{A}$, let

$$\|\mathbf{A}\| = \max_{\|\mathbf{u}\|_2=1} \|\mathbf{A}\mathbf{u}\|_2$$

denote the spectral (or operator) norm of $\mathbf{A}$; observe that by the above, we have $\|\mathbf{P}\| = \|\mathbf{P}_b\|^2$.

Corollary 5.5. *Let $d \leq n$ be positive integers, $\mathcal{O}$ be an order of a degree-d number field K with a power $\mathbb{Z}$-basis $\vec{p}$, $\Sigma \in \mathbb{R}^{d \times d}$ be a positive semidefinite matrix, and q be a positive integer. For any $\Sigma' \succeq \mathbf{P}_b^t \cdot \Sigma \cdot \mathbf{P}_b$, there is a polynomial-time randomized reduction from (search or decision) $\mathcal{O}\text{-LWE}_{q,D_{\sqrt{\Sigma}},\ell}$ to (search or decision, respectively) $\mathsf{MP\text{-}LWE}_{n,d,q,D_{\sqrt{\Sigma'}},\ell}$.*

In particular, for any $r' \geq r\sqrt{\|\mathbf{P}\|}$, there is a polynomial-time randomized reduction from (search or decision) $\mathcal{O}\text{-LWE}_{q,D_r,\ell}$ to (search or decision, respectively) $\mathsf{MP\text{-}LWE}_{n,d,q,D_{r'},\ell}$.

Proof. By applying Theorem 5.3 we obtain a polynomial-time randomized reduction from $\mathcal{O}\text{-LWE}_{q,D_{\sqrt{\Sigma}},\ell}$ to $\mathsf{MP\text{-}LWE}_{n,d,q,\psi',\ell}$, where $\psi' = \mathrm{Tr}_{K_{\mathbb{R}}/\mathbb{R}}(\vec{p} \cdot D_{\sqrt{\Sigma}})$ is a distribution over $\mathbb{R}^d$ and is analyzed as follows. Because $\vec{b}$ is an orthonormal basis of $K_{\mathbb{R}}$, the original error distribution $D_{\sqrt{\Sigma}}$ over $K_{\mathbb{R}}$ has the form $\vec{b}^t \cdot D_{\sqrt{\Sigma}}$ where $D_{\sqrt{\Sigma}}$ is a Gaussian over $\mathbb{R}^d$. Then by $\mathbb{R}$-linearity of the trace,

$$\psi' = \mathrm{Tr}_{K_{\mathbb{R}}/\mathbb{R}}(\vec{p} \cdot \vec{b}^t \cdot D_{\sqrt{\Sigma}}) = \mathrm{Tr}_{K_{\mathbb{R}}/\mathbb{R}}(\vec{p} \cdot \vec{b}^t) \cdot D_{\sqrt{\Sigma}} = \mathbf{P}_b^t \cdot D_{\sqrt{\Sigma}} = D_{\sqrt{\Sigma_1}},$$

where $\Sigma_1 = \mathbf{P}_b^t \cdot \Sigma \cdot \mathbf{P}_b$.

Since $\Sigma' \succeq \Sigma_1$ by assumption, we may transform the error distribution $D_{\sqrt{\Sigma_1}}$ to $D_{\sqrt{\Sigma'}}$ by adding (to the $\mathbf{b}$-part of each MP-LWE sample) a fresh error term from the compensating Gaussian distribution of covariance $\Sigma' - \Sigma_1 \succeq 0$. This yields the desired error distribution and completes the proof of the first claim.

For the second claim, notice that if $\Sigma = r^2 \cdot \mathbf{I}$, then $\Sigma' = (r')^2 \cdot \mathbf{I} \succeq \mathbf{P}_b^t \cdot \Sigma \cdot \mathbf{P}_b = r^2 \cdot \mathbf{P}$, because $(r')^2 \mathbf{I} - r^2 \mathbf{P}$ is positive semidefinite, since $\mathbf{x}^t \mathbf{P} \mathbf{x} \leq \|\mathbf{P}\| \cdot \|\mathbf{x}\|_2^2$ for any $\mathbf{x}$. $\qquad\square$

5.4 Example Instantiations

Corollary 5.5 bounds the expansion of the error distribution by the square root of the spectral norm of the Gram matrix $\mathbf{P}$ of a power $\mathbb{Z}$-basis $\vec{p}$ of $\mathcal{O}$. Here we show that there are large families of orders with well-behaved power bases.

Let α be an algebraic integer with minimal polynomial $f(x) \in \mathbb{Z}[x]$ of degree d, and consider the order $\mathcal{O} = \mathbb{Z}[\alpha] \subset K = \mathbb{Q}(\alpha)$, which has power $\mathbb{Z}$-basis $\vec{p} = (1, \alpha, \ldots, \alpha^{d-1})$. Consider the Vandermonde matrix

$$\mathbf{V} = \begin{pmatrix} 1 & \alpha_1 & \alpha_1^2 & & \alpha_1^{d-1} \\ 1 & \alpha_2 & \alpha_2^2 & \cdots & \alpha_2^{d-1} \\ 1 & \alpha_3 & \alpha_3^2 & & \alpha_3^{d-1} \\ & \vdots & & \ddots & \vdots \\ 1 & \alpha_d & \alpha_d^2 & \cdots & \alpha_d^{d-1} \end{pmatrix}$$

where the α_i are the d distinct roots of f, i.e., the conjugates of α. This $\mathbf{V}$ represents the linear transform σ that maps coefficient vectors with respect to $\vec{p}$ to the canonical (or Minkowski) embedding.

It is easy to see that the Gram matrix of $\vec{p}$ is $\mathbf{P} = \mathbf{V}^* \mathbf{V}$, where $\mathbf{V}^*$ denotes the conjugate transpose of $\mathbf{V}$, so $\sqrt{\|\mathbf{P}\|} = \|\mathbf{V}\|$. Therefore, we immediately have the bound $\sqrt{\|\mathbf{P}\|} \leq \|\mathbf{V}\|_2 \leq \sqrt{d} \cdot \max_i \|\sigma(\alpha^i)\|$, where the maximum is taken over $i \in [d]$. That is, the Frobenius and Euclidean norms of the power-basis elements (in the canonical embedding) yield bounds on the error expansion. The following lemma gives an alternative bound directly in terms of the minimal polynomial $f(x)$.

Lemma 5.6. *Adopt the above notation, and assume that the minimal polynomial $f(x) = x^d - g(x) \in \mathbb{Z}[x]$, where $g(x) = a_k x^k + \cdots + a_1 x + a_0$ has degree at most $k < d$. Then $\sqrt{\|\mathbf{P}\|} \leq d \cdot A^{d/(d-k)}$ where $A = \sum_{i=0}^{k} |a_i|$. In particular, if $k = (1-c)d$ for some $c \in (0,1)$, then $\sqrt{\|\mathbf{P}\|} \leq d \cdot A^{1/c}$.*

For example, if all the $|a_i| = \mathrm{poly}(d)$ and $c < 1$ is any positive constant, then $\sqrt{\|\mathbf{P}\|} = \mathrm{poly}(d)$. This enlarges the set of moduli $f(x)$ yielding polynomial error expansion from those considered in [RSSS17].

Proof. We bound $\|\mathbf{V}\|$ as follows. Let $\alpha_* = \max_i |\alpha_i| \geq 1$ be the maximum magnitude of any root of f. Then $\|\mathbf{V}\| \leq d \max |\mathbf{V}_{i,j}| \leq d \cdot \alpha_*^d$. Now, because the α_i satisfy $\alpha_i^d = g(\alpha_i)$, by the triangle inequality we have $\alpha_*^d \leq \alpha_*^k \cdot A$ and hence $\alpha_*^{d-k} \leq A$. The claim follows by raising to the $d/(d-k)$ power. $\qquad\square$

6 Hardness of Module-LWE

In this section we obtain a simple reduction from $\mathcal{O}'$-LWE, for a *wide class* of orders $\mathcal{O}'$, to a *single* $\mathcal{O}$-LWE problem of higher rank (i.e., Module-LWE) over a suborder $\mathcal{O} \subset \mathcal{O}'$ of lower rank. The reduction preserves the "total rank" of the problems, i.e., the product of the ranks of the LWE problem and the order over which it is defined.

Theorem 6.1. *Let K'/K be a degree-r number field extension, $\mathcal{O}$ be an order of K, $\mathcal{O}'$ be an order of K' that is a (rank-r) free $\mathcal{O}$-module with known basis $\vec{b}$, ψ' be a distribution over $K'_{\mathbb{R}}$, and $\mathcal{Q}$ be an $\mathcal{O}$-ideal with $\mathcal{Q}' = \mathcal{O}'\mathcal{Q}$. Then for any positive integer k', there is a polynomial-time deterministic reduction from (search or decision) $\mathcal{O}'$-LWE$_{\mathcal{Q}',\psi',\ell}^{k'}$ to (search or decision, respectively) $\mathcal{O}$-LWE$_{\mathcal{Q},\psi,\ell}^{k}$, where $k = k'r$ and $\psi = \mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}(\psi')$.*

Proof. First note that by Lemma 2.18 (for tower $K'/K/\mathbb{Q}$ and orders $\mathcal{O}', \mathcal{O}, \mathbb{Z}$), we have $(\mathcal{O}')^\vee = (\mathcal{O}')^{\vee_\mathcal{O}} \mathcal{O}^\vee$. Furthermore, $\vec{b}^\vee$ is an $\mathcal{O}$-basis of $(\mathcal{O}')^{\vee_\mathcal{O}}$, because $\vec{b}$ is an $\mathcal{O}$-basis of $\mathcal{O}'$. We invoke Theorem 4.10 with fractional $\mathcal{O}'$-ideals (and free $\mathcal{O}$-modules) $\mathcal{M}'_a = \mathcal{O}', \mathcal{M}'_s = \mathcal{M}'_b = (\mathcal{O}')^{\vee_\mathcal{O}}$ having respective bases $\vec{b}_a = \vec{b}, \vec{b}_s = \vec{b}_b = \vec{b}^\vee$, and ideals $\mathcal{A} = \mathcal{O}, \mathcal{S} = \mathcal{B} = \mathcal{O}^\vee$, which yield $\mathcal{A}' = \mathcal{O}', \mathcal{S}' = \mathcal{B}' = (\mathcal{O}')^\vee$. This gives a polynomial-time deterministic reduction from (search or decision) $\mathcal{O}'$-LWE$_{\mathcal{Q}',\psi',\ell}^{k'} = \mathsf{LWE}_{T',(\mathcal{O}')^\vee,\mathcal{O}',\mathcal{Q}',\psi',\ell}$, where $T'_{i'j'1} = \delta_{i'j'}$ is the $k' \times k' \times 1$ identity-matrix tensor over $\mathcal{O}'_{\mathcal{Q}'}$, to (search or decision, respectively) $\mathsf{LWE}_{\tilde{T},\mathcal{O}^\vee,\mathcal{O},\tilde{\psi},\ell}$, where $\tilde{T}_{(i',i)(j',j)\ell} = \delta_{i'j'} \cdot \mathrm{Tr}_{K'/K}(b_i^\vee b_j b_\ell) \in \mathcal{O}_\mathcal{Q}^{k \times k \times r}$ and $\tilde{\psi} = \mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}(\vec{b} \cdot \psi')$.[20]

To complete the reduction to $\mathcal{O}$-LWE$_{\mathcal{Q},\psi,\ell}^{k}$, we take an appropriate linear combination of the r "layers" of $\tilde{T}$ to obtain the $k \times k \times 1$ identity-matrix tensor, and invoke Theorem 4.8. Let $\mathrm{Tr} = \mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}$ and define $\mathbf{B} = \mathrm{Tr}(\vec{b}^\vee)^t \in \mathcal{O}^{1 \times r}$, where the membership holds because $\mathbf{B}$ is the coefficient vector of $1 \in \mathcal{O}'$ relative to the $\mathcal{O}$-basis $\vec{b}$. Notice that for any $x \in K'$, by K-linearity of Tr and the definition of dual basis we have that

$$\mathbf{B} \cdot \mathrm{Tr}(\vec{b} \cdot x) = \mathrm{Tr}(\vec{b}^\vee)^t \cdot \mathrm{Tr}(\vec{b} \cdot x) = \mathrm{Tr}\big((\vec{b}^\vee)^t \cdot \mathrm{Tr}(\vec{b} \cdot x)\big) = \mathrm{Tr}(x).$$

Therefore, $\mathrm{Tr}(\vec{b} \cdot r^{-1}) \in K^{r \times 1}$ is a right inverse of $\mathbf{B}$, and similarly, $\mathbf{B}\tilde{\psi} = \mathrm{Tr}(\psi') = \psi$. Next, define

---

[20]For convenience, for the first two indices of $\tilde{T}$ we identify $[k] = [k'r]$ with $[k'] \times [r]$, as described in Section 2.1.

$T \in \mathcal{O}_{\mathcal{Q}}^{k \times k \times 1}$ as

$$
\begin{aligned}
T_{(i',i)(j',j)1} &= \sum_\ell \tilde{T}_{(i',i)(j',j)\ell} \mathbf{B}_{1\ell} \\
&= \delta_{i'j'} \cdot \operatorname{Tr}_{K'/K}\Big(b_i^\vee b_j \sum_\ell b_\ell \operatorname{Tr}_{K'/K}(b_\ell^\vee) \Big) \\
&= \delta_{i'j'} \cdot \operatorname{Tr}_{K'/K}(b_i^\vee b_j) \\
&= \delta_{i'j'} \cdot \delta_{ij} = \delta_{(i',i)(j',j)},
\end{aligned}
$$

so T is the identity-matrix tensor, as desired. Therefore, taking $\mathbf{A} = \mathbf{S}$ to be the $k \times k$ identity matrix and invoking Theorem 4.8, we get a polynomial-time deterministic reduction from (search or decision) $\mathsf{LWE}_{\tilde{T}, \mathcal{O}^\vee, \mathcal{O}, \tilde{\psi}, \ell}$ to (search or decision, respectively) $\mathsf{LWE}_{T, \mathcal{O}^\vee, \mathcal{O}, \psi, \ell} = \mathcal{O}\text{-}\mathsf{LWE}_{\mathcal{Q}, \psi, \ell}^k$. $\qquad\square$

6.1 Managing the Error Distribution

Similarly to our reduction from $\mathcal{O}$-LWE to MP-LWE in Section 5, we want a reduction from many $\mathcal{O}'$-LWE problems to a single $\mathcal{O}$-LWE^k problem. To control the resulting error distribution, we consider the usual case where the original error distribution ψ' is a Gaussian, in which case it turns out that the resulting error distribution ψ is also a Gaussian. As in Section 5.3, we can add some independent Gaussian error with a compensating covariance to obtain any large enough desired target covariance. Alternatively, when ψ' is a *spherical* Gaussian, then ψ is one as well, with a covariance that is an r factor larger, so no compensating error is needed. (Also note that $(\mathcal{O}')^\vee$ is typically denser than $\mathcal{O}^\vee$ in their respective canonical embeddings—or seen another way, $\mathcal{O}$ can have shorter vectors than $\mathcal{O}'$—so the increase in covariance does not necessarily represent an actual increase in the relative error.)

In what follows, let K'/K be a degree-r number field extension, fix some orthonormal $\mathbb{R}$-bases $\vec{c}' = \tau((\vec{c}')^\vee)$ and $\vec{c} = \tau(\vec{c}^\vee)$ of $K'_{\mathbb{R}}$ and $K_{\mathbb{R}}$ (respectively) for defining Gaussian distributions, and let $\mathbf{A} = \operatorname{Tr}_{K'_{\mathbb{R}}/\mathbb{R}}(\vec{c}' \cdot \tau(\vec{c})^t)$ be the real matrix whose (i,j)th entry is $\langle c_i', c_j \rangle$.

Corollary 6.2. *Adopt the notation and hypotheses of Theorem 6.1, with $\psi' = D_{\sqrt{\Sigma'}}$ over $K'_{\mathbb{R}}$ for some positive semidefinite matrix Σ'. For any $\Sigma \succeq \mathbf{A}^t \cdot \Sigma' \cdot \mathbf{A}$ with $\psi = D_{\sqrt{\Sigma}}$ over $K_{\mathbb{R}}$, there is a polynomial-time randomized reduction from (search or decision) $\mathcal{O}'\text{-}\mathsf{LWE}_{\mathcal{Q}', \psi', \ell}^{k'}$ to (search or decision, respectively) $\mathcal{O}\text{-}\mathsf{LWE}_{\mathcal{Q}, \psi, \ell}^k$.*

Moreover, for $s = s'\sqrt{r}$, there is a polynomial-time deterministic reduction from (search or decision) $\mathcal{O}'\text{-}\mathsf{LWE}_{\mathcal{Q}', D_{s'}, \ell}^{k'}$ to (search or decision, respectively) $\mathcal{O}\text{-}\mathsf{LWE}_{\mathcal{Q}, D_s, \ell}^k$.

Proof. By Theorem 6.1, there exists a polynomial-time deterministic reduction from $\mathcal{O}'\text{-}\mathsf{LWE}_{\mathcal{Q}', \psi', \ell}^{k'}$ to $\mathcal{O}\text{-}\mathsf{LWE}_{\mathcal{Q}, \psi_1, \ell}^k$ where $\psi_1 = \operatorname{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}(\psi')$ is analyzed as follows. The original error distribution ψ' over $K'_{\mathbb{R}}$ has the form $\psi' = \vec{c}'^t \cdot D_{\sqrt{\Sigma'}}$ where here $D_{\sqrt{\Sigma'}}$ is a Gaussian over $\mathbb{R}^{\deg(K'/\mathbb{Q})}$. Then by linearity,

$$
\psi_1 = \operatorname{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}(\psi') = \vec{c}^t \cdot \operatorname{Tr}_{K'_{\mathbb{R}}/\mathbb{R}}(\tau(\vec{c}) \cdot \vec{c}'^t \cdot D_{\sqrt{\Sigma'}}) = \vec{c}^t \cdot \mathbf{A}^t \cdot D_{\sqrt{\Sigma'}} = \vec{c}^t \cdot D_{\sqrt{\Sigma_1}}
$$

is the distribution $D_{\sqrt{\Sigma_1}}$ over $K_{\mathbb{R}}$, where $\Sigma_1 = \mathbf{A}^t \cdot \Sigma' \cdot \mathbf{A}$. Since $\Sigma \succeq \Sigma_1$ by assumption, we can transform the error distribution $D_{\sqrt{\Sigma_1}}$ to $D_{\sqrt{\Sigma}}$ by adding (to the b-part of each $\mathcal{O}$-LWE sample) a fresh error term from the compensating Gaussian distribution of covariance $\Sigma - \Sigma_1 \succeq 0$. This yields the desired error distribution and completes the proof of the first claim.

For the second claim, observe that because $\vec{c}'$ and $\vec{c}$ are orthonormal,

$$\mathbf{A}^t \cdot \mathbf{A} = \mathrm{Tr}_{K'_\mathbb{R}/\mathbb{R}}(\vec{c} \cdot \tau(\vec{c})^t) = \mathrm{Tr}_{K_\mathbb{R}/\mathbb{R}}(\mathrm{Tr}_{K'_\mathbb{R}/K_\mathbb{R}}(1) \cdot \vec{c} \cdot \tau(\vec{c})^t) = \mathrm{Tr}_{K_\mathbb{R}/\mathbb{R}}(r \cdot \vec{c} \cdot \tau(\vec{c})^t) = r \cdot \mathbf{I}.$$

Therefore, if $\Sigma' = (s')^2 \cdot \mathbf{I}$ and $\Sigma = s^2 \cdot \mathbf{I}$, then $\Sigma_1 = \mathbf{A}^t \cdot \Sigma' \cdot \mathbf{A} = r(s')^2 \cdot \mathbf{I} = s^2 \cdot \mathbf{I} = \Sigma$, so no compensating error is needed, yielding a deterministic reduction. $\qquad\square$

6.2 Instantiations

It is straightforward to instantiate Corollary 6.2 to get reductions from a huge class of $\mathcal{O}'$-LWE problems over various orders $\mathcal{O}'$ to a single $\mathcal{O}$-LWE problem. Let $\mathcal{O}$ be an arbitrary order of a number field K, and let α denote some root of an arbitrary monic irreducible degree-r polynomial $f(X) \in \mathcal{O}[X]$. Then we can satisfy the hypotheses of Theorem 6.1 by letting $K' = K(\alpha)$ and $\mathcal{O}' = \mathcal{O}[\alpha]$, so that $(1, \alpha, \dots, \alpha^{r-1})$ is an $\mathcal{O}$-basis of $\mathcal{O}'$. (We emphasize that there are no restrictions on the choice of the algebraic integer α, other than its degree over $\mathcal{O}$.) Letting, e.g., $\psi' = D_{s'}$ be a spherical Gaussian over $K'_\mathbb{R}$ and $\psi = D_{s'\sqrt{r}}$ be the corresponding spherical Gaussian over $K_\mathbb{R}$, we obtain a polynomial-time deterministic reduction from $\mathcal{O}'\text{-LWE}^{k'}_{\mathcal{Q}',\psi',\ell}$ to $\mathcal{O}\text{-LWE}^k_{\mathcal{Q},\psi,\ell}$, where $k = k'r$.

References

[AD17]   M. R. Albrecht and A. Deo. Large modulus Ring-LWE $\geq$ Module-LWE. In *ASIACRYPT*, pages 267–296. 2017. Pages 2 and 8.

[BBPS19] M. Bolboceanu, Z. Brakerski, R. Perlman, and D. Sharma. Order-LWE and the hardness of Ring-LWE with entropic secrets. In *ASIACRYPT*, pages 91–120. 2019. Full version at https://eprint.iacr.org/2018/494. Pages 2, 4, 6, 13, and 18.

[BBS21]   M. Bolboceanu, Z. Brakerski, and D. Sharma. On algebraic embedding for unstructured lattices. Cryptology ePrint Archive, Paper 2021/053, 2021. https://eprint.iacr.org/2021/053. Page 6.

[BGV12]   Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. *TOCT*, 6(3):13, 2014. Preliminary version in ITCS 2012. Page 2.

[BLP+13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. 2013. Pages 2 and 8.

[CDO01]   H. Cohen, F. Diaz y Diaz, and M. Olivier. Algorithmic methods for finitely generated abelian groups. *J. Symb. Comput.*, 31(1/2):133–147, 2001. Page 8.

[CIV16]   W. Castryck, I. Iliashenko, and F. Vercauteren. Provably weak instances of Ring-LWE revisited. In *EUROCRYPT*, pages 147–167. 2016. Page 2.

[Coh99]   H. Cohen. *Advanced Topics in Computational Number Theory*. Graduate Texts in Mathematics. Springer, November 1999. Page 8.

[Con09]   K. Conrad. The conductor ideal, 2009. Available at http://www.math.uconn.edu/~kconrad/blurbs/, last accessed 14 May 2019. Page 22.

[DD12]   L. Ducas and A. Durmus. Ring-LWE in polynomial rings. In *Public Key Cryptography*, pages 34–51. 2012. Page 5.

[HPS98]   J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288. 1998. Page 2.

[Lan94]   S. Lang. *Algebraic number theory*, volume 110 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1994. Page 9.

[LM06]   V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155. 2006. Page 2.

[LPR10]   V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in Eurocrypt 2010. Pages 2, 4, 5, 6, 13, and 18.

[LS15]   A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015. Pages 2 and 7.

[Lyu16]   V. Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In *ASIACRYPT*, pages 196–214. 2016. Page 2.

[Mat89]   H. Matsumura. *Commutative ring theory*, volume 8 (2nd ed.) of *Cambridge Studied in Advanced Mathematics*. Cambridge University Press, 1989. Page 13.

[Mic02]   D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002. Page 2.

[Pei09]   C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009. Page 2.

[Pei16a] C. Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016. Page 2.

[Pei16b] C. Peikert. How (not) to instantiate Ring-LWE. In *SCN*, pages 411–430. 2016. Page 2.

[PP19]   C. Peikert and Z. Pepin. Algebraically structured LWE, revisited. In *TCC*, pages 1–23. 2019. Page 8.

[PR06]   C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166. 2006. Page 2.

[PRS17]   C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of Ring-LWE for any ring and modulus. In *STOC*, pages 461–473. 2017. Pages 2, 4, 5, and 6.

[Reg05]   O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005. Page 2.

[Reg10]   O. Regev. The learning with errors problem (invited survey). In *IEEE Conference on Computational Complexity*, pages 191–204. 2010. Page 2.

[RSSS17] M. Rosca, A. Sakzad, D. Stehlé, and R. Steinfeld. Middle-product learning with errors. In *CRYPTO*, pages 283–297. 2017. Pages 2, 3, 5, 7, 26, 27, and 30.

[RSW18] M. Rosca, D. Stehlé, and A. Wallet. On the Ring-LWE and Polynomial-LWE problems. In *EUROCRYPT*, pages 146–173. 2018. Pages 2, 4, 5, 6, 7, 18, 22, and 26.

[SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635. 2009. Page 2.

[SSZ19] R. Steinfeld, A. Sakzad, and R. K. Zhao. Practical MP-LWE-based encryption balancing security-risk versus efficiency. *Des. Codes Cryptogr.*, 87(12):2847–2884, 2019. Pages 5 and 7.