Algebraically Structured LWE, Revisited

Chris Peikert[*] Zachary Pepin[†]

October 26, 2020

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 some 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.

[*]Computer Science and Engineering, University of Michigan. Email: cpeikert@umich.edu. This material is based upon work supported by the National Science Foundation under Award CNS-1606362 and by DARPA under Agreement No. HR00112020025. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Government, the National Science Foundation, or DARPA.

[†]Computer Science and Engineering, University of Michigan. Email: zapepin@umich.edu.

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] and Micciancio's initial worst-case to average-case reductions for "algebraically structured" lattices over polynomial rings [Mic02], Lyubashevsky, Peikert, and Regev [LPR10] introduced *Ring-LWE* to improve the asymptotic and practical efficiency of LWE (see also [SSTX09]). 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 large polynomials. 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 precisely which problems and parameters are supported by reductions and proofs. The importance of having a clear, precise view of the landscape 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] for an overview. This work aims to provide such a view.

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 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 (noiseless) products $s \cdot a$ each belong to—or, more generally, are vectors over—a quotient $\mathcal{I}/q\mathcal{I}$ for some respective *fractional ideals* $\mathcal{I} = \mathcal{I}_s, \mathcal{I}_a, \mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$ of some common order $\mathcal{O}$ of a number field K. Moreover, the products are given by some fixed $\mathcal{O}_q$*-bilinear map* on s and a (where $\mathcal{O}_q = \mathcal{O}/q\mathcal{O}$); by fixing appropriate $\mathcal{O}_q$-bases, this bilinear map can be represented as an *order-three tensor* (i.e., a three-dimensional array) over $\mathcal{O}_q$.

A generalized LWE problem is defined by some fixed choices of the above parameters (order, ideals, and tensor), along with an error distribution. For example, plain LWE uses $\mathcal{O} = \mathcal{I}_s = \mathcal{I}_a = \mathbb{Z}$, with the ordinary n-dimensional 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 $\mathcal{I}_a = \mathcal{O}, \mathcal{I}_s = \mathcal{O}^\vee$ being the fractional "codifferent" ideal, and the bilinear map being ordinary multiplication in K, 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 matrix "slices" $M_{i\cdot\cdot}$ of the middle-product tensor M are seen to form the standard basis for the space of all *Hankel* matrices. (In a Hankel matrix, the (j, k)th entry is determined by $j + k$.) This formulation is central to our improved reduction from Ring-LWE over a wide class of number fields to Middle-Product LWE, described in Section 1.2.3 below.

Parameterizing by a single lattice. Next, in Section 3.2 we define a specialization of generalized LWE that 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* (i.e., discrete additive subgroup) $\mathcal{L}$ of a number field K. Define

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

to be the set of field elements by which $\mathcal{L}$ is closed under multiplication; this set is known as the *coefficient ring* of $\mathcal{L}$. 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, i.e., a subring with unity that is also a lattice. Note that if $\mathcal{L}$ itself is an order $\mathcal{O}$ of K or its dual $\mathcal{O}^\vee$, then $\mathcal{O}^{\mathcal{L}} = \mathcal{O}$, but in general $\mathcal{L}$ can be any lattice, and $\mathcal{O}^{\mathcal{L}}$ is just the largest order of K by which $\mathcal{L}$ is closed under multiplication.[1]

In what follows, let $\mathcal{L}_q$ denote the quotient $\mathcal{L}/q\mathcal{L}$ for any lattice $\mathcal{L}$ of K and positive integer q. In $\mathcal{L}$-LWE, there is a secret $s \in \mathcal{L}_q^\vee$, and we are given independent noisy random products

$$(a_i \leftarrow \mathcal{O}_q^{\mathcal{L}}\,,\, b = s \cdot a_i + e_i \bmod q\mathcal{L}^\vee),$$

where each a_i is uniformly random and each e_i is an error term that is drawn from a specified distribution.[2] We show that under mild conditions, taking the multipliers from the lattice's coefficient ring (modulo q) is

[1]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 inclusion 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 $c\mathcal{L} \subsetneq \mathcal{L}' \subsetneq \mathcal{L}$.

[2]Observe that the reduction modulo $q\mathcal{L}^\vee$ is well defined because the (noiseless) product $s \cdot a_i \in \mathcal{L}_q^\vee$, since $\mathcal{L}^\vee \cdot \mathcal{O}^{\mathcal{L}} \subseteq \mathcal{L}^\vee$ due to $\mathrm{Tr}(\mathcal{L}^\vee \cdot \mathcal{O}^{\mathcal{L}} \cdot \mathcal{L}) \subseteq \mathrm{Tr}(\mathcal{L}^\vee \cdot \mathcal{L}) \subseteq \mathbb{Z}$.

actually without loss of generality, which justifies using this specific definition of $\mathcal{L}$-LWE. More generally, s and a can be k-dimensional vectors over $\mathcal{L}_q^\vee$ and $\mathcal{O}_q^{\mathcal{L}}$ (respectively), with $s \cdot a \in \mathcal{L}_q^\vee$ denoting their inner product; we call this variant $\mathcal{L}$-LWE^k.

We now explain how $\mathcal{L}$-LWE generalizes prior algebraic LWE problems. 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 (and for general $k \geq 1$ we get "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] when $\mathcal{L} = \mathcal{O}$ or $\mathcal{L} = \mathcal{O}^\vee$ for some arbitrary order $\mathcal{O}$ of K.

Notice that in the latter two cases, $\mathcal{L}$ may be the *dual* of its coefficient ring $\mathcal{O} = \mathcal{O}^{\mathcal{L}}$, so the secret s and product $s \cdot a$ belong to $\mathcal{L}^\vee = \mathcal{O}$ itself (modulo q). But as we shall see, for reductions it turns out to be more natural and advantageous to let $\mathcal{L}$ itself be an order, not its dual. 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, 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 recover 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 $\mathcal{I}_s$ and $\mathcal{I}_a$, the order $\mathcal{O}$ (or the lattice $\mathcal{L}$ defining it, in the case of $\mathcal{L}$-LWE), the tensor T defining the bilinear map, or the number field over which all the other parameters are defined. All of the reductions preserve the number of samples, most of them are even tight polynomial-time *equivalences* (i.e., reductions in both directions), and most also preserve the error distribution. When the latter is not the case, the error distribution is changed by an easy-to-analyze linear transformation.

All of the main results of this work are then obtained by invoking and/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 from $\mathcal{L}$-LWE to $\mathcal{L}'$-LWE

As a first main result (see Theorem 4.6), we obtain a reduction from $\mathcal{L}$-LWE to $\mathcal{L}'$-LWE for any lattices $\mathcal{L}' \subseteq \mathcal{L}$ of K for which $\mathcal{O}^{\mathcal{L}'} \subseteq \mathcal{O}^{\mathcal{L}}$ and the index $|\mathcal{L}/\mathcal{L}'|$ is coprime with the modulus 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 $\mathcal{L}'$ to be any other order of K (whose index in $\mathcal{O}_K$ is coprime with q).

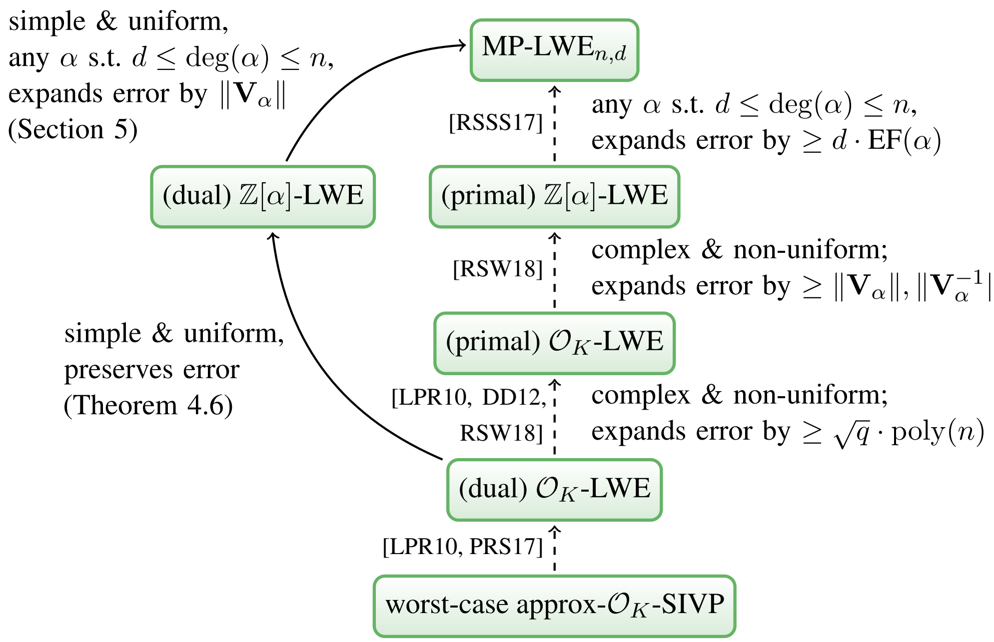

Figure 1: Summary of (some) known reductions among algebraic lattice and LWE problems. Dashed arrows represent prior reductions, and solid arrows represent (some of) the reductions given in this work.

We stress that the only "loss" associated with the reduction, which seems inherently necessary, is that when $\mathcal{L} \neq \mathcal{L}'$, the lattice $q(\mathcal{L}')^\vee$ by which the resulting noisy products $b' \approx s' \cdot a'$ are reduced is "denser" than the lattice $q\mathcal{L}^\vee \subsetneq q(\mathcal{L}')^\vee$ 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 accumulated error from some combined samples to be decodable modulo $q(\mathcal{L}')^\vee$. That is, we need to be able to efficiently recover e' (or at least a large portion of it) from the coset $e' + q(\mathcal{L}')^\vee$; to do this, standard decoding algorithms require sufficiently short elements of $q^{-1}\mathcal{L}'$. So, the "sparser" we take $\mathcal{L}' \subseteq \mathcal{L}$ to be, the denser $(\mathcal{L}')^\vee$ is, and the larger we need q to be 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 work. We now describe some of the immediate implications of the above reduction, and compare to prior related ones. Take $\mathcal{L} = \mathcal{O}_K$ 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 $\mathcal{L}'$ to be an arbitrary order $\mathcal{O}$ of K, as long as $|\mathcal{L}'/\mathcal{L}|$ is coprime with q. These guarantees are qualitatively similar to the ones established in [RSW18, BBPS19], 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].

Our reduction can also yield hardness for the "primal" form of Poly-LWE and Order-LWE via a different choice of $\mathcal{L}'$ (see the next paragraph); however, it is instructive to see why it is preferable to reduce to the "dual" form of these problems. The main reason is that the dual form admits 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 Section 1.2.3 and Section 1.2.4 below for further details.)

By contrast, the reduction and analysis for "primal" Order-LWE over an order $\mathcal{O}$—including 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 primal Order-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 $q\mathcal{O}$ is mainly determined by short nonzero vectors in $\mathcal{O}^\vee$, 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 lossier reduction.)

1.2.3 Reduction from $\mathcal{O}$-LWE to MP-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/MP-LWE connection similar to the one obtained in [RSSS17, RSW18], which implies 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 one by being simpler, more general, and tighter: 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.

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 from $\mathcal{O}'$-LWE$^{d'}$ to $\mathcal{O}$-LWE^d

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 indeed much more widely used in theory and in practice.

A bit more precisely, we give a simple reduction from $\mathcal{O}'$-$\mathsf{LWE}^{d'}$, for a wide class of orders $\mathcal{O}'$, to a *single* $\mathcal{O}$-LWE^d problem. The only condition we require is that $\mathcal{O}'$ should be a rank-(d/d') free $\mathcal{O}$-module. For example, this is easily achieved by defining $\mathcal{O} = \mathcal{O}[\alpha] \cong \mathcal{O}[x]/f(x)$ for some root α of an arbitrary degree-(d/d') 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^d (for certain $\mathcal{O}$) given in [LS15]: there the source problem involves arbitrary (worst-case) rank-d *module lattices* over $\mathcal{O}$, whereas here our source problem is an average-case rank-d' LWE problem over a rank-(d/d') $\mathcal{O}$-module.

2 Preliminaries

In this work, by "ring" we always mean a commutative ring with identity.

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 it as $T_{i_1 i_2 \cdots i_r}$. (The 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 $\mathbf{A}$.

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 often 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 $A_{i_1 \cdots i_r} B_{j_1 \cdots j_r}$.

2.2 Number Fields, Lattices, and Duality

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 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 discrete additive subgroup of K for which $\operatorname{span}_{\mathbb{Q}}(\mathcal{L}) = K$, i.e., every lattice has full rank. Any lattice is generated as the integer linear combinations of n *basis* elements $\vec{b} = (b_1, \ldots, b_n) \in K^n$, as $\mathcal{L} = \{\sum_{i=1}^n \mathbb{Z} \cdot b_i\}$; in other words, $\mathcal{L}$ is a free $\mathbb{Z}$-module of rank n. For convenience, we let $\mathcal{L}_q$ denote the quotient group $\mathcal{L}/q\mathcal{L}$ for any positive integer q.

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, and given bases for $\mathcal{L}, \mathcal{L}'$ we can efficiently compute a basis for $\mathcal{L} \cdot \mathcal{L}'$ via the Hermite normal form.

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

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

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 $\operatorname{Tr}(b_i \cdot b_j^\vee)$ is 1 when $i = j$, and is 0 otherwise. Observe that by definition, $x = \vec{b}^t \cdot \operatorname{Tr}(\vec{b}^\vee \cdot x)$ for every $x \in K$.

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 $\mathcal{I} \subseteq (\mathcal{L} : \mathcal{L}')$. As we shall see, several sets of interest for this work can be defined as quotients of various lattices. The following gives an alternative characterization of the lattice quotient, and yields a way to efficiently compute a basis for $(\mathcal{L} : \mathcal{L}')$ given bases for $\mathcal{L}$ and $\mathcal{L}'$.

Lemma 2.2. *For any lattices $\mathcal{L}, \mathcal{L}'$ in K, we have $(\mathcal{L} : \mathcal{L}') = (\mathcal{L}'\mathcal{L}^\vee)^\vee$.*

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

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

2.3 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 $\operatorname{Tr}_{K_{\mathbb{R}}/\mathbb{R}}$ 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 $\langle x, y \rangle = \operatorname{Tr}_{K_{\mathbb{R}}/\mathbb{R}}(x \cdot \tau(y))$. 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 an 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 2.3. 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 $\mathbf{x} \in \mathrm{span}(\Sigma) = \Sigma \cdot \mathbb{R}^n$ 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 $\mathrm{span}(\Sigma)$) is proportional to $\rho_{\sqrt{\Sigma}}$.[3]

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 $\Sigma' - \Sigma$.

2.4 Orders and Ideals

We recall some basic notions relating to orders and their ideals; see [Con09] for more details and missing proofs. Throughout this subsection let K be an arbitrary number field.

An *order* $\mathcal{O}$ of 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 f such that $f(\alpha) = 0$. 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 $\mathcal{O} \cdot \mathcal{O} = \mathcal{O}$.

An *ideal* of an order $\mathcal{O}$, also called an $\mathcal{O}$-ideal, is a nontrivial additive subgroup $\mathcal{I} \subseteq \mathcal{O}$ that is closed under multiplication by $\mathcal{O}$, i.e., $\mathcal{O} \cdot \mathcal{I} \subseteq \mathcal{I}$; in fact this is an equality, since $1 \in \mathcal{O}$.[4] 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 (hence, $\mathcal{I}$ is a lattice).

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 Definition 2.1 that $\mathcal{O}^{\mathcal{L}} = (\mathcal{L}\mathcal{L}^\vee)^\vee$, so $\mathcal{O}^{\mathcal{L}} = \mathcal{O}^{\mathcal{L}^\vee}$, and if $\mathcal{L}$ is itself 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 an order of K.*

Proof. 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. By definition of $\mathcal{O}^{\mathcal{L}}$, we clearly have $1 \in \mathcal{O}^{\mathcal{L}}$. 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 desired. ∎

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

[4] In this work we 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 $\mathcal{O}$.

Definition 2.6 (Conductor ideal). For any orders $\mathcal{O} \subseteq \mathcal{O}'$ in K, their *(relative) conductor ideal* is

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

We often omit the superscript $\mathcal{O}'$ when it is $\mathcal{O}_K$.

Observe that $\mathcal{C} = \mathcal{C}_{\mathcal{O}}^{\mathcal{O}'} \subseteq \mathcal{O}$ because $1 \in \mathcal{O}'$, and it is immediate from the definition that $\mathcal{C}$ is an ideal of both $\mathcal{O}$ and $\mathcal{O}'$.

Throughout the rest of this subsection let $\mathcal{O}$ be an arbitrary order of K. 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. In particular, any conductor ideal $\mathcal{C} = \mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$ for $\mathcal{O} \subsetneq \mathcal{O}'$ is not invertible as an $\mathcal{O}$-ideal: for if $\mathcal{C}\mathcal{I} = \mathcal{O}$ for some fractional $\mathcal{O}$-ideal $\mathcal{I}$, then $\mathcal{O}' = \mathcal{O}'\mathcal{O} = \mathcal{O}'\mathcal{C}\mathcal{I} \subseteq \mathcal{C}\mathcal{I} = \mathcal{O}$, a contradiction.

Despite the lack of ideal inverses in general, there is a proxy that turns out to be just as good for our purposes.

Definition 2.7 (Pseudoinverse). The *pseudoinverse* of a fractional $\mathcal{O}$-ideal $\mathcal{I}$ is the fractional $\mathcal{O}$-ideal

$$\widetilde{\mathcal{I}} := (\mathcal{O} : \mathcal{I}) = \{x \in K : x\mathcal{I} \subseteq \mathcal{O}\}.$$

We stress that both the inverse and pseudoinverse are defined with respect to the particular order $\mathcal{O}$. Furthermore, it is easy to prove that if $\mathcal{I}$ is invertible, then $\mathcal{I}^{-1} = \widetilde{\mathcal{I}}$.

Lemma 2.8. *For any fractional $\mathcal{O}$-ideal $\mathcal{I}$, we have $\mathcal{C}_{\mathcal{O}} \subseteq \mathcal{I}\widetilde{\mathcal{I}} \subseteq \mathcal{C}_{\mathcal{O}}^{\mathcal{O}^{\mathcal{I}}} \subseteq \mathcal{O}$.*

Proof. For $\mathcal{I}\widetilde{\mathcal{I}} \subseteq \mathcal{C}_{\mathcal{O}}^{\mathcal{O}^{\mathcal{I}}}$, by the definitions of coefficient ring and pseudoinverse we have $\mathcal{O}^{\mathcal{I}}\mathcal{I}\widetilde{\mathcal{I}} \subseteq \mathcal{I}\widetilde{\mathcal{I}} \subseteq \mathcal{O}$, as needed. For $\mathcal{C}_{\mathcal{O}} \subseteq \mathcal{I}\widetilde{\mathcal{I}}$, because $(\mathcal{O}_K : \mathcal{I}\mathcal{O}_K)$ is the inverse of $\mathcal{I}\mathcal{O}_K$ as a fractional $\mathcal{O}_K$-ideal, we have

$$\mathcal{C}_{\mathcal{O}} = \mathcal{O}_K\mathcal{C}_{\mathcal{O}} = \mathcal{I}\mathcal{O}_K(\mathcal{O}_K : \mathcal{I}\mathcal{O}_K)\mathcal{C}_{\mathcal{O}} \subseteq \mathcal{I}(\mathcal{C}_{\mathcal{O}} : \mathcal{I}) \subseteq \mathcal{I}(\mathcal{O} : \mathcal{I}) = \mathcal{I}\widetilde{\mathcal{I}},$$

where the first inclusion follows from the previous equalities and the definition of the lattice quotient. □

Using the pseudoinverse we generalize the notion of ideal invertibility, by defining it modulo another ideal.

Definition 2.9. For an $\mathcal{O}$-ideal $\mathcal{J}$, we say that a fractional $\mathcal{O}$-ideal $\mathcal{I}$ is *invertible modulo* $\mathcal{J}$ if $\mathcal{I}\widetilde{\mathcal{I}}$ and $\mathcal{J}$ are coprime, i.e., $\mathcal{I}\widetilde{\mathcal{I}} + \mathcal{J} = \mathcal{O}$.

Observe that if $\mathcal{I}$ is invertible, then it is also invertible modulo any $\mathcal{J}$, because $\mathcal{I}\widetilde{\mathcal{I}} = \mathcal{O}$. Furthermore, if $\mathcal{J}$ and the conductor ideal $\mathcal{C}_{\mathcal{O}}$ are coprime, then *any* fractional ideal $\mathcal{I}$ is invertible modulo $\mathcal{J}$, because $\mathcal{O} = \mathcal{C}_{\mathcal{O}} + \mathcal{J} \subseteq \mathcal{I}\widetilde{\mathcal{I}} + \mathcal{J}$ by Lemma 2.8.

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 $\mathcal{M}$ as an $\mathcal{O}$-module.

Theorem 2.10 (Chinese Remainder Theorem). *Let $\mathcal{I}_1, \ldots, \mathcal{I}_r$ be any pairwise coprime $\mathcal{O}$-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. Moreover, it is efficiently computable and invertible given (bases of) $\mathcal{O}$, the $\mathcal{I}_i$, and $\mathcal{M}$.

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 $\mathcal{I}$, by requiring that $\mathcal{I}$ is invertible modulo $\mathcal{J}$ (using $\widetilde{I}$).

Lemma 2.11. *Let $\mathcal{J}$ be an $\mathcal{O}$-ideal. Then a fractional $\mathcal{O}$-ideal $\mathcal{I}$ is invertible modulo $\mathcal{J}$ if and only if there exists $t \in \mathcal{I}$ such that $t\widetilde{\mathcal{I}} + \mathcal{J} = \mathcal{O}$. Moreover, in this case such t can be found efficiently given $\mathcal{O}, \mathcal{I}$ and all the prime $\mathcal{O}$-ideals that contain $\mathcal{J}$.*
In particular, if $\mathcal{I} = \mathcal{O}' \supseteq \mathcal{O}$ is an order, then $\mathcal{O}'$ is invertible modulo $\mathcal{J}$ if and only if $\mathcal{C}_{\mathcal{O}}^{\mathcal{O}'} + \mathcal{J} = \mathcal{O}$.

Proof. First, if such t exists then $\mathcal{O} = t\widetilde{\mathcal{I}} + \mathcal{J} \subseteq \mathcal{I}\widetilde{\mathcal{I}} + \mathcal{J} \subseteq \mathcal{O}$, so the inclusions are equalities, and $\mathcal{I}$ is invertible modulo $\mathcal{J}$.

For the other direction, let $\mathfrak{p}_1, \ldots, \mathfrak{p}_r$ be the distinct prime $\mathcal{O}$-ideals that contain $\mathcal{J}$. First we show that any $t \in \mathcal{I} \setminus \bigcup_{i=1}^r \mathfrak{p}_i\mathcal{I}$ satisfies the desired condition. Indeed, because $t\widetilde{\mathcal{I}} + \mathcal{J} \subseteq \mathcal{O}$, if this is not an equality then $t\widetilde{\mathcal{I}} + \mathcal{J} \subseteq \mathfrak{p}$ for some maximal (and hence prime) ideal $\mathfrak{p} \subseteq \mathcal{O}$, which implies that $t\widetilde{\mathcal{I}} \subseteq \mathfrak{p}$ and $\mathcal{J} \subseteq \mathfrak{p}$, so $\mathfrak{p} = \mathfrak{p}_i$ for some i. By these inclusions, the fact that $t \in \mathcal{I}$, and the hypothesis that $\mathcal{I}\widetilde{\mathcal{I}} + \mathcal{J} = \mathcal{O}$, we have

$$t \in t\mathcal{O} = t(\mathcal{I}\widetilde{\mathcal{I}} + \mathcal{J}) \subseteq \mathfrak{p}_i\mathcal{I} + t\mathcal{J} \subseteq \mathfrak{p}_i\mathcal{I} \ ,$$

so $t \in \mathfrak{p}_i\mathcal{I}$, which contradicts the choice of t.

Now, we show that such a t exists and can be computed efficiently. First, note that $\mathcal{I} \neq \mathfrak{p}_i\mathcal{I}$ for all i.[5] So, for each i, choose some non-zero $t_i \in \mathcal{I}/\mathfrak{p}_i\mathcal{I}$, and let $t \in \mathcal{I}$ be (an arbitrary representative of) the preimage of $(t_1, \ldots, t_r) \in \bigoplus_{i=1}^r \mathcal{I}/\mathfrak{p}_i\mathcal{I}$ under the isomorphism given by Theorem 2.10 (which we can invoke here because the $\mathfrak{p}_i$ are distinct maximal ideals, and hence pairwise coprime). Clearly, $t \in \mathcal{I} \setminus \bigcup_{i=1}^r \mathfrak{p}_i\mathcal{I}$, as desired.

For the case where $\mathcal{I} = \mathcal{O}' \supseteq \mathcal{O}$ is an order, observe that $\widetilde{\mathcal{O}'} = (\mathcal{O} : \mathcal{O}') = \mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$ by definition. Therefore, if $\mathcal{O}'$ is invertible modulo $\mathcal{J}$ then by the above we can take $t = 1 \in \mathcal{O}' \setminus \bigcup_{i=1}^r \mathfrak{p}_i\mathcal{O}'$ to get $t\mathcal{C}_{\mathcal{O}}^{\mathcal{O}'} + \mathcal{J} = \mathcal{O}$, hence $\mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$ is coprime with $\mathcal{J}$. ☐

Lemma 2.12. *Let $\mathcal{J}$ be an $\mathcal{O}$-ideal, $\mathcal{I}$ be a fractional $\mathcal{O}$-ideal that is invertible modulo $\mathcal{J}$, and $t \in \mathcal{I}$ be such that $t\widetilde{\mathcal{I}} + \mathcal{J} = \mathcal{O}$ (as guaranteed by Lemma 2.11). Then for any fractional $\mathcal{O}$-ideal $\mathcal{M}$, the function $\theta_t \colon K \to K$ defined as $\theta_t(u) = t \cdot u$ induces an $\mathcal{O}$-module isomorphism from $\mathcal{M}/\mathcal{J}\mathcal{M}$ to $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$. Moreover, this isomorphism is efficiently invertible given $\mathcal{O}, \mathcal{I}, \mathcal{J}, \mathcal{M}$, and t.*

In particular, if $\mathcal{I} = \mathcal{O}' \supseteq \mathcal{O}$ is an order, then by Lemma 2.11 we can take $t = 1$, making the induced $\mathcal{O}$-module isomorphism the natural inclusion map.

Proof. That θ_t induces an $\mathcal{O}$-module homomorphism follows immediately from the fact that it is multiplication by a fixed $t \in \mathcal{O}$. Now consider the function from $\mathcal{M}$ to $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$ that is induced by θ_t. Its kernel

[5]By Nakayama's lemma, if $\mathcal{I} = \mathfrak{p}_i\mathcal{I}$, then there must be 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.

clearly contains $\mathcal{J}\mathcal{M}$, and is in fact equal to $\mathcal{J}\mathcal{M}$, which may be seen as follows. If $u \cdot t \in \mathcal{I}\mathcal{J}\mathcal{M}$ for some $u \in \mathcal{M}$, then $u \cdot t\widetilde{\mathcal{I}} \subseteq \mathcal{I}\mathcal{J}\mathcal{M}\widetilde{\mathcal{I}} \subseteq \mathcal{J}\mathcal{M}$. Because $t\widetilde{\mathcal{I}} + \mathcal{J} = \mathcal{O}$, we get that $u \in u\mathcal{O} = u(t\widetilde{\mathcal{I}} + \mathcal{J}) \subseteq \mathcal{J}\mathcal{M}$, as desired. So, the function from $\mathcal{M}/\mathcal{J}\mathcal{M}$ to $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$ induced by θ_t is injective. It remains to show that it can be efficiently inverted, which also implies that it is an isomorphism.

Let $v \in \mathcal{I}\mathcal{M}$ be arbitrary. By hypothesis, $t\widetilde{\mathcal{I}}$ and $\mathcal{J}$ are coprime (in $\mathcal{O}$). Therefore, we can use the algorithm from [LPR10, Lemma 2.13] (which works for arbitrary orders $\mathcal{O}$) to compute some $c \in t\widetilde{\mathcal{I}}$ such that $c = 1 \pmod{\mathcal{J}}$. Then let $a = c \cdot v \in t\widetilde{\mathcal{I}}\mathcal{I}\mathcal{M} \subseteq t\mathcal{M}$, and observe that $a - v = v \cdot (c-1) \in \mathcal{I}\mathcal{J}\mathcal{M}$. Let $w = a/t \in \mathcal{M}$; then $\theta_t(w) = t \cdot (a/t) = v \pmod{\mathcal{I}\mathcal{J}\mathcal{M}}$, so $w \bmod \mathcal{J}\mathcal{M}$ is the preimage of $v \bmod \mathcal{I}\mathcal{J}\mathcal{M}$. □

2.6 Bijective Natural Inclusions

For any lattices $\mathcal{L}' \subseteq \mathcal{L}$ in K, the *natural inclusion map* $\mathcal{L}'_q \to \mathcal{L}_q$ sends $x + q\mathcal{L}'$ to $x + q\mathcal{L}$. (This can be seen as the composition of a natural homomorphism and an inclusion map.) The following lemmas give conditions under which maps of this kind are bijections.

Lemma 2.13. *Let $\mathcal{L}' \subseteq \mathcal{L}$ be lattices in a number field K and let q be a positive integer. Then the natural inclusion map $h\colon \mathcal{L}'_q \to \mathcal{L}_q$ is a bijection if and only if q is coprime with the index $|\mathcal{L}/\mathcal{L}'|$; in this case, h is efficiently computable and invertible given an arbitrary basis of $\mathcal{L}'$ relative to a basis of $\mathcal{L}$.*

Because $|\mathcal{L}/\mathcal{L}'| = |(\mathcal{L}')^\vee/\mathcal{L}^\vee|$, the same conclusions hold for the natural inclusion map $\mathcal{L}^\vee_q \to (\mathcal{L}')^\vee_q$.

Proof. Let $\vec{b}, \vec{b}'$ respectively be some $\mathbb{Z}$-bases of $\mathcal{L}, \mathcal{L}'$ (and hence $\mathbb{Z}_q$-bases of $\mathcal{L}_q, \mathcal{L}'_q$). Then $\vec{b}' = \mathbf{T} \cdot \vec{b}$ for some square matrix $\mathbf{T}$. This $\mathbf{T}$ is integral because $\mathcal{L}' \subseteq \mathcal{L}$, and we have $|\det(\mathbf{T})| = |\mathcal{L}/\mathcal{L}'|$. Letting $\mathbf{x}'$ be the coefficient vector (over $\mathbb{Z}_q$) of some arbitrary $x' = \langle \vec{b}', \mathbf{x} \rangle \in \mathcal{L}'_q$, we have $x' = \langle \mathbf{T} \cdot \vec{b}, \mathbf{x}' \rangle = \langle \vec{b}, \mathbf{T}^t \cdot \mathbf{x}' \rangle$, so $\mathbf{x} = \mathbf{T}^t \cdot \mathbf{x}'$ is the coefficient vector (over $\mathbb{Z}_q$) of $h(x') \in \mathcal{L}_q$ relative to $\vec{b}$. Moreover, $\mathbf{x}$ and $\mathbf{x}'$ are in bijective correspondence if and only if $\mathbf{T}$ is invertible modulo q, i.e., if $|\det(\mathbf{T})| = |\mathcal{L}/\mathcal{L}'|$ is coprime with q, and we can efficiently evaluate and invert this bijection given $\mathbf{T}$. □

Lemma 2.14. *Let $\mathcal{O}' \subseteq \mathcal{O}$ be orders in a number field K and q be a positive integer. Then the following statements are equivalent:*

1. *The natural inclusion map $\mathcal{O}'_q \to \mathcal{O}_q$ is a bijection.*
2. *The index $|\mathcal{O}/\mathcal{O}'|$ is coprime with q.*
3. *The conductor $\mathcal{C}^{\mathcal{O}}_{\mathcal{O}'}$ is coprime with $q\mathcal{O}'$, i.e., $\mathcal{C}^{\mathcal{O}}_{\mathcal{O}'} + q\mathcal{O}' = \mathcal{O}'$.*

Proof. First, Item 1 and Item 2 are equivalent by Lemma 2.13. Also, Item 1 follows from Item 3 by Lemma 2.12. Finally, to see that Item 3 follows from Item 2, let $m = |\mathcal{O}/\mathcal{O}'|$ be the order of the group $\mathcal{O}/\mathcal{O}'$, which is finite. The order of every element of $\mathcal{O}/\mathcal{O}'$ divides m, so $m \in (\mathcal{O}' : \mathcal{O}) = \mathcal{C}^{\mathcal{O}}_{\mathcal{O}'}$, which is an $\mathcal{O}'$-ideal. It follows that $m\mathcal{O}' \subseteq \mathcal{C}^{\mathcal{O}}_{\mathcal{O}'}$. Finally, since m is coprime with q, we have $\mathcal{O}' = m\mathcal{O}' + q\mathcal{O}' \subseteq \mathcal{C}^{\mathcal{O}}_{\mathcal{O}'} + q\mathcal{O}' \subseteq \mathcal{O}'$, so the inclusions are in fact equalities, as needed. □

Lemma 2.15. *Let $\mathcal{L}' \subseteq \mathcal{L}$ be lattices in a number field K, and let q be a positive integer that is coprime with the index $|\mathcal{L}/\mathcal{L}'|$. If $\mathcal{O}^{\mathcal{L}'} \subseteq \mathcal{O}^{\mathcal{L}}$, then the natural inclusion map $g\colon \mathcal{O}^{\mathcal{L}'}_q \to \mathcal{O}^{\mathcal{L}}_q$ is a bijection.*

Note that the converse does not hold, because for $\mathcal{L}' = q\mathcal{L}$ we have $\mathcal{O}^{\mathcal{L}'} = \mathcal{O}^{\mathcal{L}}$, so the natural inclusion map g is a bijection, but q is not coprime with $|\mathcal{L}/\mathcal{L}'| = q^{\deg(K/\mathbb{Q})}$.

Proof. Let $h\colon \mathcal{L}'_q \to \mathcal{L}_q$ be the natural inclusion map, which by Lemma 2.13 is a bijection. First, notice that for any $a \in \mathcal{O}_q^{\mathcal{L}'}$ and $x \in \mathcal{L}'_q$, we have $h(a \cdot x) = g(a) \cdot h(x)$. This is because

$$g(a) \cdot h(x) = (a + q\mathcal{O}^{\mathcal{L}}) \cdot (x + q\mathcal{L}) = a \cdot x + q(\mathcal{O}^{\mathcal{L}} \cdot x + a \cdot \mathcal{L} + \mathcal{O}^{\mathcal{L}} \cdot \mathcal{L}) = a \cdot x + q\mathcal{L} = h(a \cdot x).$$

Now, let $a, b \in \mathcal{O}_q^{\mathcal{L}'}$ satisfy $g(a) = g(b)$. Then for all $x \in \mathcal{L}'$, we have

$$h(a \cdot x) = g(a) \cdot h(x) = g(b) \cdot h(x) = h(b \cdot x).$$

Since h is a bijection, it follows that $a \cdot x = b \cdot x \pmod{q\mathcal{L}'}$ for all $x \in \mathcal{L}'$. Therefore,

$$(a - b) \cdot \mathcal{L}' \subseteq q\mathcal{L}' \Rightarrow a - b \in q\mathcal{O}^{\mathcal{L}'} \Rightarrow a = b \pmod{q\mathcal{O}^{\mathcal{L}'}}.$$

Thus, g is injective. Since the sets $\mathcal{O}_q^{\mathcal{L}'}$ and $\mathcal{O}_q^{\mathcal{L}}$ have the same cardinality $q^{\deg(K/\mathbb{Q})}$, g must bijective. □

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* $\mathrm{Tr} = \mathrm{Tr}_{K'/K}\colon K' \to K$ is the trace of the K-linear transformation on K' (viewed as a vector space over K) representing multiplication by x. We extend the trace to the real inner-product spaces $K'_{\mathbb{R}}$ and $K_{\mathbb{R}}$ in the natural way, writing $\mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}$ 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 $\mathrm{Tr}_{K'/K}(b_i b_j^\vee)$ 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* is defined as

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

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 see that if $\mathcal{L}$ has an $\mathcal{O}$-basis $\vec{b}$, then $\vec{b}^\vee$ is an $\mathcal{O}$-basis of $\mathcal{L}^{\vee_{\mathcal{O}}}$.

Lemma 2.16. *Let K'/K be a number field extension with K-basis $\vec{b}$, and let $x = \langle \vec{b}^\vee, \vec{x}\rangle, y = \langle \vec{b}, \vec{y}\rangle$ for some $\vec{x}, \vec{y}$ over K. Then $\mathrm{Tr}_{K'/K}(x \cdot y) = \langle \vec{x}, \vec{y}\rangle$.*

Proof. Letting $\mathrm{Tr} = \mathrm{Tr}_{K'/K}$, by K-linearity of Tr we have

$$\mathrm{Tr}(x \cdot y) = \mathrm{Tr}(\langle \vec{b}^\vee, \vec{x}\rangle \cdot \langle \vec{b}, \vec{y}\rangle) = \mathrm{Tr}(\vec{x}^t \cdot (\vec{b}^\vee \cdot \vec{b}^t) \cdot \vec{y}) = \vec{x}^t \cdot \mathrm{Tr}(\vec{b}^\vee \cdot \vec{b}^t) \cdot \vec{y} = \vec{x}^t \cdot I \cdot \vec{y} = \langle \vec{x}, \vec{y}\rangle. \quad \square$$

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 is transitive, decomposing as $\mathrm{Tr}_{K''/K} = \mathrm{Tr}_{K'/K} \circ \mathrm{Tr}_{K''/K'}$. Moreover, if $\vec{c}$ is a K'-basis of K'' and $\vec{b}$ is a K-basis of K', then by definition and K'-linearity of $\mathrm{Tr}_{K''/K'}$ we have $(\vec{c} \otimes \vec{b})^\vee = \vec{c}^\vee \otimes \vec{b}^\vee$.

Lemma 2.17. *Let $K''/K'/K$ be a tower of number field extensions and $\mathcal{O}''$, $\mathcal{O}'$, and $\mathcal{O}$ respectively be orders of K'', K', and K, where $\mathcal{O}''$ has an $\mathcal{O}'$-basis $\vec{c}$ and $\mathcal{O}'$ has an $\mathcal{O}$-basis $\vec{b}$. Then $(\mathcal{O}'')^{\vee_{\mathcal{O}}} = (\mathcal{O}'')^{\vee_{\mathcal{O}'}}(\mathcal{O}')^{\vee_{\mathcal{O}}}$.*

Proof. We prove this by demonstrating a common $\mathcal{O}$-basis for both sets. By hypothesis, $\vec{c} \otimes \vec{b}$ is an $\mathcal{O}$-basis of $\mathcal{O}''$, and thus $(\vec{c} \otimes \vec{b})^\vee = \vec{c}^\vee \otimes \vec{b}^\vee$ is an $\mathcal{O}$-basis of $(\mathcal{O}'')^{\vee_{\mathcal{O}}}$. We next show that this is also an $\mathcal{O}$-basis of $(\mathcal{O}'')^{\vee_{\mathcal{O}'}}(\mathcal{O}')^{\vee_{\mathcal{O}}}$. It is immediate from the definition that $(\mathcal{O}')^{\vee_{\mathcal{O}}}$ is a fractional $\mathcal{O}'$-ideal, and $\vec{c}^\vee$ is an $\mathcal{O}'$-basis of $(\mathcal{O}'')^{\vee_{\mathcal{O}'}}$, so

$$(\mathcal{O}'')^{\vee_{\mathcal{O}'}}(\mathcal{O}')^{\vee_{\mathcal{O}}} = \sum_i c_i^\vee \cdot \mathcal{O}' \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}} = \sum_i c_i^\vee \cdot (\mathcal{O}')^{\vee_{\mathcal{O}}}.$$

Since $\vec{b}^\vee$ is an $\mathcal{O}$-basis of $(\mathcal{O}')^{\vee_{\mathcal{O}}}$, we conclude that $\vec{c}^\vee \otimes \vec{b}^\vee$ is an $\mathcal{O}$-basis of $(\mathcal{O}'')^{\vee_{\mathcal{O}'}}(\mathcal{O}')^{\vee_{\mathcal{O}}}$, as desired. □

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 will be 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 $\mathcal{I}/q\mathcal{I}$ (or its many-fold Cartesian product) for some respective *fractional ideals* $\mathcal{I}$ of some common order $\mathcal{O}$ of a number field. Moreover, the products are given by some fixed $\mathcal{O}_q$*-bilinear map* on s and a (where recall that $\mathcal{O}_q = \mathcal{O}/q\mathcal{O}$). 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 $K = \mathbb{Q}$.
- 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;
2. suitable fractional $\mathcal{O}$-ideals $\mathcal{I}_s, \mathcal{I}_a, \mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$, respective dimensions k_s, k_a, k_b, and a modulus q;
3. an order-three tensor $T \in \mathcal{O}_q^{k_s \times k_a \times k_b}$, which induces an $\mathcal{O}_q$-bilinear map $T \colon (\mathcal{I}_s/q\mathcal{I}_s)^{k_s} \times (\mathcal{I}_a/q\mathcal{I}_a)^{k_a} \to (\mathcal{I}_b/q\mathcal{I}_b)^{k_b}$ defined as $T(\vec{s}, \vec{a})_k = \sum_{i,j} T_{ijk} s_i a_j$; and
4. an error distribution ψ over $K_\mathbb{R}^{k_b}$.

Informally, the associated computational problems are concerned with "noisy products" $(\vec{a} \leftarrow (\mathcal{I}_a/q\mathcal{I}_a)^{k_a}, \vec{b} \approx T(\vec{s}, \vec{a}))$ for some fixed $\vec{s} \in (\mathcal{I}_s/q\mathcal{I}_s)^{k_s}$. 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 $\vec{s} \in (\mathcal{I}_s/q\mathcal{I}_s)^{k_s}$, a sample from the distribution $A_{T,\mathcal{I}_s,\mathcal{I}_a,\psi}(\vec{s})$ over $(\mathcal{I}_a/q\mathcal{I}_a)^{k_a} \times (K_\mathbb{R}/q\mathcal{I}_b)^{k_b}$ is generated by choosing $\vec{a} \leftarrow (\mathcal{I}_a/q\mathcal{I}_a)^{k_a}$ uniformly at random, choosing $\vec{e} \leftarrow \psi$, and outputting

$$(\vec{a}\ ,\ T(\vec{s}, \vec{a}) + \vec{e} \bmod (q\mathcal{I}_b)^{k_b}).$$

For notational convenience, we also define the uniform distribution $U_{T,\mathcal{I}_s,\mathcal{I}_a} = U((\mathcal{I}_a/q\mathcal{I}_a)^{k_a} \times (K_\mathbb{R}/q\mathcal{I}_b)^{k_b})$.

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

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

Discussion. Once the order $\mathcal{O}$ is fixed, typical choices of $\mathcal{I}_s$ and $\mathcal{I}_a$ are $\mathcal{I}_a = \mathcal{O}$, which is actually without loss of generality under mild conditions (as shown in Section 4.1), and $\mathcal{I}_s = \mathcal{O}^\vee$, which is especially advantageous for reductions from Ring-LWE (as shown in Section 4.2), and to Middle-Product LWE (as shown in Section 5).

As noted above, the tensor defines an $\mathcal{O}_q$-bilinear map $(\mathcal{I}_s/q\mathcal{I}_s)^{k_s} \times (\mathcal{I}_a/q\mathcal{I}_a)^{k_a} \to (\mathcal{I}_b/q\mathcal{I}_b)^{k_b}$. Under the mild assumption (which is also needed for many of our reductions) that the ideals $\mathcal{I}_s$ and $\mathcal{I}_a$—and hence $\mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$ as well—are invertible modulo $q\mathcal{O}$, the converse holds as well: any bilinear map with such domain and range can be represented by an order-three tensor over $\mathcal{O}_q$. To see this, first note that by Lemma 2.12, for each $z \in \{s, a, b\}$ the $\mathcal{O}_q$-module $\mathcal{I}_z/q\mathcal{I}_z$ is isomorphic to $\mathcal{O}_q$. Because the latter module has a one-element $\mathcal{O}_q$-basis $\{1\}$, the former also has a one-element $\mathcal{O}_q$-basis $\{g_z\}$ for some $g_z \in \mathcal{I}_z/q\mathcal{I}_z$, where $g_b = g_s g_a$. This naturally extends to the "standard basis" of $(\mathcal{I}_z/q\mathcal{I}_z)^{k_z}$, whose ith vector has g_z in its ith component and zeros elsewhere.

Using the above bases, any bilinear map T can be uniquely represented as an order-three tensor T over $\mathcal{O}_q$ by letting $T_{ijk} \in \mathcal{O}_q$ be the kth 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 $\mathcal{O}_q$-bilinearity of the map and the fact that $g_b = g_s g_a$, 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 prior algebraic LWE problems, including Ring-, Module-, Order-, and Poly-LWE. The key observation is that all of these problems can be obtained simply by parameterizing by a single *lattice* in a given number field, then taking the public multipliers to be over the lattice's *coefficient ring* modulo q, and using a tensor corresponding to an identity matrix. Indeed, the first two of these simplifications are without loss of generality among a broad class of generalized LWE parameterizations, by the reductions we give in Theorem 4.1 and Theorem 4.4 (see Remark 4.5).

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 $\mathcal{L}$, ψ be a distribution over $K_{\mathbb{R}}$, and q and k be positive integers. Let $T \in \mathcal{O}_q^{k\times k\times 1}$ be the order-three tensor whose single $k \times k$ layer is the identity matrix. The (search or decision) $\mathcal{L}$-$\mathsf{LWE}^k_{q,\psi,\ell}$ problem is then simply the (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{L}^\vee,\mathcal{O},\psi,\ell}$ problem.

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

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 "Module" analogues). Recall that if $\mathcal{L}$ is an 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, we get the Ring-LWE problem as originally defined in [LPR10], and by taking $\mathcal{L} = \mathcal{O}$ to be some order we get the "dual" form of Order-LWE [BBPS19]. Alternatively, by taking $\mathcal{L} = \mathcal{O}^\vee$ we get the "primal" form of Order-LWE, which corresponds to the Poly-LWE problem [RSW18] when

$\mathcal{O} = \mathbb{Z}[\alpha]$ for some $\alpha \in \mathcal{O}_K$. As we will see, the "dual" formulations have advantages 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 $\mathcal{I}_s, \mathcal{I}_a$ (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.

4.1 Changing the Ideals

In this section we give reductions that map from one choice of the ideals $\mathcal{I}_s, \mathcal{I}_a$ 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 $\mathcal{O}_q := \mathcal{O}/q\mathcal{O}$ instead of $\mathcal{I}_a/q\mathcal{I}_a$ (with a corresponding change to the domain of the entries of $\vec{s}$) when $\mathcal{I}_a$ is invertible modulo $q\mathcal{O}$; recall that this is the case for *all* fractional ideals when $q\mathcal{O}$ is coprime to the conductor ideal $\mathcal{C}_\mathcal{O}$. This transformation is tight in all respects and reversible, so in fact it yields an equivalence between (search or decision) $\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$ and $\mathsf{LWE}_{T,\mathcal{I}'_s,\mathcal{I}'_a,\psi,\ell}$ whenever $\mathcal{I}_a, \mathcal{I}'_a$ are invertible modulo $q\mathcal{O}$, and $\mathcal{I}_a\mathcal{I}_s = \mathcal{I}'_a\mathcal{I}'_s$.

Theorem 4.1. *Let $\mathcal{O}$ be an order in a number field K; $\mathcal{I}_s$ and $\mathcal{I}_a$ be fractional $\mathcal{O}$-ideals with $\mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$; k_s, k_a, k_b, and q be positive integers; $T \in \mathcal{O}_q^{k_s \times k_a \times k_b}$ be an order-three tensor; and ψ be a distribution over $K_\mathbb{R}^{k_b}$. If $\mathcal{I}_a$ is invertible modulo $q\mathcal{O}$, then there is an efficiently computable and invertible deterministic transform which:*

1. *maps distribution $U_{T,\mathcal{I}_s,\mathcal{I}_a}$ to distribution $U_{T,\mathcal{I}_b,\mathcal{O}}$, and*
2. *maps distribution $A_{T,\mathcal{I}_s,\mathcal{I}_a,\psi}(\vec{s})$ to distribution $A_{T,\mathcal{I}_b,\mathcal{O},\psi}(\vec{s}')$, where $\vec{s}' = h(\vec{s})$ for some efficiently computable and invertible $h\colon \mathcal{I}_s/q\mathcal{I}_s \to \mathcal{I}_b/q\mathcal{I}_b$.*

In particular, (search or decision) $\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$ is polynomial-time equivalent to (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{I}_b,\mathcal{O},\psi,\ell}$.

Proof. Since $\mathcal{I}_a$ is invertible modulo $q\mathcal{O}$, Lemma 2.11 says there exists a $t \in \mathcal{I}_a$ such that $t\widetilde{\mathcal{I}_a} + q\mathcal{O} = \mathcal{O}$. Then by Lemma 2.12, the function $\theta_t(u) = t \cdot u$ induces efficiently computable and invertible $\mathcal{O}$-module isomorphisms $g\colon \mathcal{O}/q\mathcal{O} \to \mathcal{I}_a/q\mathcal{I}_a$ and $h\colon \mathcal{I}_s/q\mathcal{I}_s \to \mathcal{I}_b/q\mathcal{I}_b$.

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

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

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 $s \in \mathcal{I}_s/q\mathcal{I}_s$ and $a \in \mathcal{I}_a/q\mathcal{I}_a$. Note that $a = t \cdot g^{-1}(a) + q\mathcal{I}_a$ and $h(s) = t \cdot s + q\mathcal{I}_b$. Therefore,

$$\begin{aligned} s \cdot a + q\mathcal{I}_b &= s \cdot (t \cdot g^{-1}(a) + q\mathcal{I}_a) + q\mathcal{I}_b \\ &= t \cdot s \cdot g^{-1}(a) + q\mathcal{I}_b \\ &= (t \cdot s + q\mathcal{I}_b) \cdot g^{-1}(a) + q\mathcal{I}_b \\ &= h(s) \cdot g^{-1}(a) + q\mathcal{I}_b. \end{aligned}$$

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. ☐

Corollary 4.2. *Adopt the notation from Theorem 4.1, and let $\mathcal{I}'_s, \mathcal{I}'_a$ be fractional $\mathcal{O}$-ideals with $\mathcal{I}'_s\mathcal{I}'_a = \mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$. If both $\mathcal{I}_a, \mathcal{I}'_a$ are invertible modulo $q\mathcal{O}$, then (search or decision) $\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$ is polynomial-time equivalent to (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{I}'_s,\mathcal{I}'_a,\psi,\ell}$.*

Proof. By Theorem 4.1, both of the problems in question are polynomial-time equivalent to $\mathsf{LWE}_{T,\mathcal{I}_b,\mathcal{O},\psi,\ell}$. ☐

The next simple theorem shows that we can replace $\mathcal{I}_s$ and $\mathcal{I}_b$ with appropriately related super-ideals. However, because this transformation may discard information, it is typically not reversible.

Theorem 4.3. *Let $\mathcal{O}$ be an order of a number field K; $\mathcal{I}_s \subseteq \mathcal{I}'_s$, $\mathcal{I}_a$, $\mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a \subseteq \mathcal{I}'_b = \mathcal{I}'_s\mathcal{I}_a$ be fractional $\mathcal{O}$-ideals; k_s, k_a, k_b, and q be positive integers; $T \in \mathcal{O}_q^{k_s \times k_a \times k_b}$ be an order-three tensor; and ψ be a distribution over $K_\mathbb{R}^{k_b}$. Then there is an efficiently computable deterministic transform which:*

1. *maps distribution $U_{T,\mathcal{I}_s,\mathcal{I}_a}$ to distribution $U_{T,\mathcal{I}'_s,\mathcal{I}_a}$, and*
2. *maps distribution $A_{T,\mathcal{I}_s,\mathcal{I}_a,\psi}(\vec{s})$ to distribution $A_{T,\mathcal{I}'_s,\mathcal{I}_a,\psi}(\vec{s}')$, where $\vec{s}' = \vec{s} \bmod q\mathcal{I}'_s$.*

In particular, there is an efficient randomized reduction from decision-$\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$ to decision-$\mathsf{LWE}_{T,\mathcal{I}'_s,\mathcal{I}_a,\psi,\ell}$. Moreover, if the natural inclusion map $\mathcal{I}_s/q\mathcal{I}_s \to \mathcal{I}'_s/q\mathcal{I}'_s$ is a bijection, then there is an efficient deterministic *reduction from (search or decision of) the former problem to (search or decision, respectively of) the latter problem.*

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

$$(\vec{a}' = \vec{a}\,, \vec{b}' = \vec{b} \bmod (q\mathcal{I}'_b)^{k_b}) \in (\mathcal{I}_a/q\mathcal{I}_a)^{k_a} \times (K_\mathbb{R}/q\mathcal{I}'_b)^{k_b}.$$

It is clear that this transform maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$. Also, since $q\mathcal{I}_b \subseteq q\mathcal{I}'_b$, the transform sends uniformly random $\vec{b}$ to uniformly random $\vec{b}'$.

To complete the proof, it suffices to show that $T(\vec{s}, \vec{a}) = T(\vec{s}', \vec{a}) \pmod{q\mathcal{I}'_b}$. By linearity, it is enough to show that for all $s \in \mathcal{I}_s/q\mathcal{I}_s$ and $a \in \mathcal{I}_a/q\mathcal{I}_a$, we have $s \cdot a = s' \cdot a \pmod{q\mathcal{I}'_b}$, where $s' = s + q\mathcal{I}'_s$. Indeed,

$$s \cdot a + q\mathcal{I}'_b = (s + q\mathcal{I}'_s) \cdot a + q\mathcal{I}'_b = s' \cdot a + q\mathcal{I}'_b.$$

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 $\vec{r}' \in (\mathcal{I}'_s/q\mathcal{I}'_s)^{k_s}$ and changing each sample $(\vec{a}', \vec{b}')$ to $(\vec{a}', \vec{b}' + T(\vec{r}', \vec{a}'))$. Moreover, when the natural inclusion map $\mathcal{I}_s/q\mathcal{I}_s \to \mathcal{I}'_s/q\mathcal{I}'_s$ is a bijection, note that $\vec{s}'$ is uniform over $(\mathcal{I}'_s/q\mathcal{I}'_s)^{k_s}$ because $\vec{s}$ is uniform over $(\mathcal{I}_s/q\mathcal{I}_s)^{k_s}$, so re-randomization is not needed and the reduction is deterministic. Finally, for the search problems, we may recover $\vec{s}$ from $\vec{s}'$ simply by applying the inverse of the natural inclusion map. □

4.2 Changing the Order

In this section we show 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 $T' = T \bmod q\mathcal{O}'$; the error distribution and the number of samples are preserved. We note that in particular, the theorem gives an equivalence between $\mathcal{L}$-LWE and $\mathcal{L}'$-LWE for lattices $\mathcal{L}' = (\mathcal{L}^\vee \mathcal{O}')^\vee = (\mathcal{L} : \mathcal{O}') \subseteq \mathcal{L}$ in K, when $\mathcal{O} := \mathcal{O}^{\mathcal{L}} \subseteq \mathcal{O}' := \mathcal{O}^{\mathcal{L}'}$. When $\mathcal{L} = \mathcal{O}$ is itself an order, $\mathcal{L}' = (\mathcal{O} : \mathcal{O}') = \mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$ is the conductor ideal of the two orders.

Theorem 4.4. *Let $\mathcal{O} \subseteq \mathcal{O}'$ be orders in a number field K; $\mathcal{I}_s$ and $\mathcal{I}_a$ be fractional $\mathcal{O}$-ideals with $\mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$; $\mathcal{I}'_s = \mathcal{I}_s\mathcal{O}'$, $\mathcal{I}'_a = \mathcal{I}_a\mathcal{O}'$, and $\mathcal{I}'_b = \mathcal{I}'_s\mathcal{I}'_a = \mathcal{I}_b\mathcal{O}'$; $T \in \mathcal{O}_q^{k_s \times k_a \times k_b}$ be an order-three tensor for positive integers k_s, k_a, k_b, and q; and ψ be a distribution over $K_{\mathbb{R}}^{k_b}$. If $\mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$ and $q\mathcal{O}$ are coprime as $\mathcal{O}$-ideals, then there is an efficiently computable deterministic transform which:*

1. *maps distribution $U_{T,\mathcal{I}_s,\mathcal{I}_a}$ to distribution $U_{T',\mathcal{I}'_s,\mathcal{I}'_a}$, and*
2. *maps distribution $A_{T,\mathcal{I}_s,\mathcal{I}_a,\psi}(\vec{s})$ to distribution $A_{T',\mathcal{I}'_s,\mathcal{I}'_a,\psi}(\vec{s}')$ where $\vec{s}' = \vec{s} \bmod (q\mathcal{I}'_s)^{k_s}$,*

where $T' = T \bmod q\mathcal{O}'$ (which is in $(\mathcal{O}'_q)^{k_s \times k_a \times k_b}$). Furthermore, if $\mathcal{I}'_b = \mathcal{I}_b$, or ψ is over $\mathcal{I}_b^{k_b} \subset K_{\mathbb{R}}^{k_b}$, then the transform is also efficiently invertible.

In particular, there is an efficient deterministic reduction from (search or decision) $\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$ to (search or decision, respectively) $\mathsf{LWE}_{T',\mathcal{I}'_s,\mathcal{I}'_a,\psi,\ell}$, and when the above transform is efficiently invertible, the problems are polynomial-time equivalent.

Proof. First, observe that $\mathcal{O}'$ is a fractional $\mathcal{O}$-ideal. By hypothesis, $t = 1 \in \mathcal{O}'$ satisfies $t\widetilde{\mathcal{O}'} + q\mathcal{O} = \mathcal{O}$. So by taking $\mathcal{I} = \mathcal{O}'$ in Lemma 2.12, the natural inclusion maps $\mathcal{O}/q\mathcal{O} \to \mathcal{O}'/q\mathcal{O}'$, $\mathcal{I}_a/q\mathcal{I}_a \to \mathcal{I}'_a/q\mathcal{I}'_a$, and $\mathcal{I}_s/q\mathcal{I}_s \to \mathcal{I}'_s/q\mathcal{I}'_s$ are efficiently invertible bijections.

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

$$(\vec{a}' = \vec{a} \bmod (q\mathcal{I}'_a)^{k_a}, \vec{b}' = \vec{b} \bmod (q\mathcal{I}'_b)^{k_b}) \in (\mathcal{I}'_a/q\mathcal{I}'_a)^{k_a} \times (K_{\mathbb{R}}/q\mathcal{I}'_b)^{k_b}.$$

It is clear that this transform maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$, since the natural inclusion map $\mathcal{I}_a/q\mathcal{I}_a \to \mathcal{I}'_a/q\mathcal{I}'_a$ is a bijection. Also, since $q\mathcal{I}_b \subseteq q\mathcal{I}'_b$, the transform sends uniformly random $\vec{b}$ to uniformly random $\vec{b}'$.

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

$$s \cdot r \cdot a + q\mathcal{I}_b' = (s + q\mathcal{I}_s') \cdot (r + q\mathcal{O}') \cdot (a + q\mathcal{I}_a') + q\mathcal{I}_b' = s' \cdot r' \cdot a' + q\mathcal{I}_b'.$$

To see that the transform is efficiently invertible (under one of the additional hypotheses), first recall that the natural inclusion map $\mathcal{I}_a/q\mathcal{I}_a \to \mathcal{I}_a'/q\mathcal{I}_a'$ is efficiently invertible. Thus, it suffices to show that the transform $\vec{b}' = \vec{b} \bmod (q\mathcal{I}_b')^{k_b}$ is efficiently invertible. This is clearly the case if $\mathcal{I}_b' = \mathcal{I}_b$, because $\vec{b}' = \vec{b}$. Alternatively, if ψ is over $\mathcal{I}_b^{k_b}$, then the function mapping each entry of $\vec{b}$ to the corresponding entry of $\vec{b}'$ is the natural inclusion map $\mathcal{I}_b/q\mathcal{I}_b \to \mathcal{I}_b'/q\mathcal{I}_b'$, which is an efficiently invertible bijection by Lemma 2.12.

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 $\mathcal{I}_s/q\mathcal{I}_s \to \mathcal{I}_s'/q\mathcal{I}_s'$. □

Remark 4.5. Theorem 4.4 (together with Theorem 4.1) shows that the special choices of order and ideals we made in defining $\mathcal{L}$-LWE (Definition 3.4) as an instantiation of generalized LWE are canonical ones, under mild conditions. Specifically, start from any generalized LWE instantiation over an order $\mathcal{O}$ where $\mathcal{I}_a = \mathcal{O}$ (following Theorem 4.1), and let $\mathcal{L} = \mathcal{I}_s^\vee$. Then if $\mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$ and $q\mathcal{O}$ are coprime as $\mathcal{O}$-ideals, then by Theorem 4.4 we can change the order to the canonical choice $\mathcal{O}' = \mathcal{O}^{\mathcal{I}_s} = \mathcal{O}^{\mathcal{L}}$, and take $\mathcal{I}_a' = \mathcal{I}_a\mathcal{O}' = \mathcal{O}'$ and $\mathcal{I}_s' = \mathcal{I}_s\mathcal{O}' = \mathcal{I}_s = \mathcal{L}^\vee$, which are exactly the choices made in $\mathcal{L}$-LWE. Note that a sufficient condition for both Theorem 4.1 and Theorem 4.4 is $\mathcal{C}_{\mathcal{O}} + q\mathcal{O} = \mathcal{O}$, by Lemma 2.8 and because $\mathcal{C}_{\mathcal{O}} \subseteq \mathcal{C}_{\mathcal{O}}^{\mathcal{O}'}$.

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

As a straightforward corollary to Theorem 4.4 we get the following efficient, deterministic reduction from $\mathcal{L}$-$\mathsf{LWE}_{q,\psi,\ell}$ to $\mathcal{L}'$-$\mathsf{LWE}_{q,\psi,\ell}$, under suitable conditions on the lattices $\mathcal{L}' \subseteq \mathcal{L}$.

Theorem 4.6. *Let $\mathcal{L}' \subseteq \mathcal{L}$ be lattices in a number field K with $\mathcal{O}^{\mathcal{L}'} \subseteq \mathcal{O}^{\mathcal{L}}$, ψ be a distribution over $K_{\mathbb{R}}$, and q be a positive integer. If $|\mathcal{L}/\mathcal{L}'|$ is coprime with q, and bases of $\mathcal{L}', \mathcal{O}^{\mathcal{L}'}$ relative to bases of $\mathcal{L}, \mathcal{O}^{\mathcal{L}}$ (respectively) are known, then there is an efficient deterministic reduction from (search or decision) $\mathcal{L}$-$\mathsf{LWE}_{q,\psi,\ell}$ to (search or decision, respectively) $\mathcal{L}'$-$\mathsf{LWE}_{q,\psi,\ell}$.*

Proof. First, the hypothesis that $|\mathcal{L}/\mathcal{L}'|$ is coprime with q, combined with the lemmas from Section 2.6, implies that the natural inclusion $\mathcal{L}^\vee/q\mathcal{L}^\vee \to (\mathcal{L}')^\vee/q(\mathcal{L}')^\vee$ is a bijection and $\mathcal{C}_{\mathcal{O}^{\mathcal{L}'}}^{\mathcal{O}^{\mathcal{L}}}$ is coprime with $q\mathcal{O}^{\mathcal{L}'}$.

Because $\mathcal{L}^\vee \subseteq \mathcal{L}^\vee\mathcal{O}^{\mathcal{L}'} \subseteq \mathcal{L}^\vee\mathcal{O}^{\mathcal{L}} = \mathcal{L}^\vee$ (and hence all the inclusions are equalities), Theorem 4.4 gives a deterministic polynomial-time equivalence between (search or decision) $\mathcal{L}$-$\mathsf{LWE}_{q,\psi,\ell} = \mathsf{LWE}_{T,\mathcal{L}^\vee,\mathcal{O}^{\mathcal{L}},\psi,\ell}$ and (search or decision, respectively) $\mathsf{LWE}_{T',\mathcal{L}^\vee,\mathcal{O}^{\mathcal{L}'},\psi,\ell}$, where T, T' correspond to identity matrices over $\mathcal{O}_q^{\mathcal{L}}, \mathcal{O}_q^{\mathcal{L}'}$, respectively. Then because $\mathcal{L}^\vee \subseteq (\mathcal{L}')^\vee$, Theorem 4.3 gives a deterministic reduction from the latter problems to (search or decision, respectively) $\mathsf{LWE}_{T',(\mathcal{L}')^\vee,\mathcal{O}^{\mathcal{L}'},\psi,\ell} = \mathcal{L}'$-$\mathsf{LWE}_{q,\psi,\ell}$, as desired. □

Remark 4.7. A main case of interest for Theorem 4.6 is when $\mathcal{L} = \mathcal{O} = \mathcal{O}^{\mathcal{L}}$ and $\mathcal{L}' = \mathcal{O}' = \mathcal{O}^{\mathcal{L}'}$ are themselves orders, in which case the coprimality hypothesis is equivalent (by Lemma 2.14) to $\mathcal{C}_{\mathcal{O}'}^{\mathcal{O}} + q\mathcal{O}' = \mathcal{O}'$. We remark that [RSW18, Theorem 4.2] proves a result similar to Theorem 4.6 for any order $\mathcal{O}'$ and $\mathcal{O} = \mathcal{O}_K$, under the hypothesis $\mathcal{C}_{\mathcal{O}'}^{\mathcal{O}_K} + q\mathcal{O}_K = \mathcal{O}_K$ (among others). By [Con09, Theorem 3.8], this hypothesis is equivalent to ours (for this choice of orders), so our result applies at least as generally as the one of [RSW18].

4.4 Changing the Tensor

We now give a reduction from one generalized LWE problem to another, when their associated tensors are suitably related.

Theorem 4.8. *Let*

- $\mathcal{O}$ *be an order in a number field* K*, and* $\mathcal{I}_s$ *and* $\mathcal{I}_a$ *be fractional* $\mathcal{O}$*-ideals with* $\mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$*;*
- $T \in \mathcal{O}_q^{k_s \times k_a \times k_b}$ *and* $T' \in \mathcal{O}_q^{k'_s \times k'_a \times k'_b}$ *be tensors over* $\mathcal{O}_q$ *for positive integers* $k_s, k_a, k_b, k'_s, k'_a, k'_b, q$*;*
- $\mathbf{S} \in \mathcal{O}_q^{k'_s \times k_s}$*,* $\mathbf{A} \in \mathcal{O}_q^{k_a \times k'_a}$*, and* $\mathbf{B} \in \mathcal{O}^{k'_b \times k_b}$ *be matrices where* $\mathbf{A}$ *and* $\mathbf{B}$ *are right invertible over* $\mathcal{O}_q$ *and* K *(respectively) and* $\sum_{jk} T_{ijk}\, \mathbf{A}_{jj'}\, \mathbf{B}_{k'k} = \sum_{i'} T'_{i'j'k'}\, \mathbf{S}_{i'i}$ *for all* i, j', k'*; and*
- ψ *be a distribution over* $K_{\mathbb{R}}^{k_b}$*.*

There is an efficient randomized transform which:

1. *maps distribution* $U_{T,\mathcal{I}_s,\mathcal{I}_a}$ *to distribution* $U_{T',\mathcal{I}_s,\mathcal{I}_a}$*, and*
2. *maps the distribution* $A_{T,\mathcal{I}_s,\mathcal{I}_a,\psi}(\vec{s})$ *to* $A_{T',\mathcal{I}_s,\mathcal{I}_a,\psi'}(\vec{s}')$*, 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 an efficient randomized reduction from decision-$\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$ *to decision-*$\mathsf{LWE}_{T',\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$*, and similarly for the search problems if* $\mathbf{S}$ *is left invertible (over* $\mathcal{O}_q$*). 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 $\nu\colon (\mathcal{I}_a/q\mathcal{I}_a)^{k'_a} \to (\mathcal{I}_a/q\mathcal{I}_a)^{k_a}$ 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 efficiently computed using standard linear-algebraic techniques. Therefore, we can efficiently sample uniformly from $\nu^{-1}(\vec{a})$, and such a sample is uniformly random over $(\mathcal{I}_a/q\mathcal{I}_a)^{k'_a}$ when $\vec{a}$ is uniformly random over $(\mathcal{I}_a/q\mathcal{I}_a)^{k_a}$. 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 $(\vec{a}, \vec{b}) \in (\mathcal{I}_a/q\mathcal{I}_a)^{k_a} \times (K_{\mathbb{R}}/q\mathcal{I}_b)^{k_b}$, we output

$$(\vec{a}' \leftarrow \nu^{-1}(\vec{a})\ ,\ \vec{b}' = \mathbf{B}\vec{b}) \in (\mathcal{I}_a/q\mathcal{I}_a)^{k'_a} \times (K_{\mathbb{R}}/q\mathcal{I}_b)^{k'_b}.$$

As already observed, this maps uniformly random $\vec{a}$ to uniformly random $\vec{a}'$. It also maps uniformly random $\vec{b} \in (K_{\mathbb{R}}/q\mathcal{I}_b)^{k_b}$ to uniformly random $\vec{b}' \in (K_{\mathbb{R}}/q\mathcal{I}_b)^{k'_b}$, 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 $\mathbf{B}(q\mathcal{I}_b)^{k_b} \subseteq (q\mathcal{I}_b)^{k'_b}$.

It remains to show that if $\vec{b} = T(\vec{s}, \vec{a}) + \vec{e} \bmod (q\mathcal{I}_b)^{k_b}$ for some $\vec{e} \leftarrow \psi$, then $\vec{b}' = T'(\vec{s}', \vec{a}') +$

$\vec{e}' \bmod (q\mathcal{I}_b)^{k'_b}$ where $\vec{e}' \leftarrow \psi'$. To see this, let $\vec{e}' = \mathbf{B}\vec{e}$ (which has distribution ψ'), and observe that

$$\begin{aligned} b'_{k'} = \sum_k \mathbf{B}_{k'k} b_k &= \sum_{i,j,k} \mathbf{B}_{k'k}(T_{ijk}\, s_i\, a_j + e_k) \\ &= \sum_{i,j,j',k} \mathbf{B}_{k'k}\, T_{ijk}\, s_i\, \mathbf{A}_{jj'}\, a'_{j'} + e'_{k'} \\ &= \sum_{i,i',j'} T'_{i'j'k'}\, \mathbf{S}_{i'i}\, s_i\, a'_{j'} + e'_{k'} \\ &= \sum_{i',j'} T'_{i'j'k'}\, s'_{i'} + e'_{k'} \\ &= T'(\vec{s}', \vec{a}')_{k'} + \vec{e}'_{k'}, \end{aligned}$$

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 $\vec{r}' \in (\mathcal{I}_s/q\mathcal{I}_s)^{k'_s}$ 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 $\vec{s} = \mathbf{S}^+\vec{s}'$. □

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 q, i.e., for which there exists some $u \in \mathcal{O}$ such that $t \cdot u = 1 \bmod q\mathcal{O}$. This is equivalent to the condition that t and q are coprime in $\mathcal{O}$, i.e., $t\mathcal{O} + q\mathcal{O} = \mathcal{O}$. 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 defined over number fields $K \subseteq K'$ of different degrees, for appropriately related orders, ideals, and tensors. Despite the rather technical nature of the theorem statement, the core idea is relatively straightforward: essentially, the tensor T' over an order $\mathcal{O}'$ (modulo q) of K' is expanded into a tensor T over an order $\mathcal{O}$ (modulo q) of K by replacing each entry with a block representing multiplication by that entry, relative to suitable bases. Formally, this is done by letting T be the entry-wise trace of the Kronecker product $T' \otimes C$, where C is the Kronecker product of the bases.

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';*
- *$\mathcal{M}'_s, \mathcal{M}'_a, \mathcal{M}'_b = \mathcal{M}'_s\mathcal{M}'_a$ 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 are hence K-bases of K');*
- *$\mathcal{I}_s, \mathcal{I}_a, \mathcal{I}_b = \mathcal{I}_s\mathcal{I}_a$ be fractional $\mathcal{O}$-ideals and $\mathcal{I}'_s = \mathcal{I}_s\mathcal{M}'_s$, $\mathcal{I}'_a = \mathcal{I}_a\mathcal{M}'_a$, $\mathcal{I}'_b = \mathcal{I}_b\mathcal{M}'_b = \mathcal{I}'_s\mathcal{I}'_a$ be the corresponding fractional $\mathcal{O}'$-ideals;*

- *q, k'_s, k'_a, k'_b be positive integers, with $k_s = k \cdot k'_s$, $k_a = k \cdot k'_a$, and $k_b = k \cdot k'_b$;*
- *$T' \in (\mathcal{O}'_q)^{k'_s \times k'_a \times k'_b}$ be an order-three tensor over $\mathcal{O}'_q$, and $T = \mathrm{Tr}(T' \otimes C) \in \mathcal{O}_q^{k_s \times k_a \times k_b}$ for the order-three tensor $C = \vec{b}_s \otimes \vec{b}_a \otimes \vec{b}_b^\vee$;*[6] *and*
- *ψ' be a distribution over $K'_\mathbb{R}$.*

Then there is an efficiently computable and invertible transform which:

1. *maps distribution $U_{T',\mathcal{I}'_s,\mathcal{I}'_a}$ to $U_{T,\mathcal{I}_s,\mathcal{I}_a}$, and*
2. *maps distribution $A_{T',\mathcal{I}'_s,\mathcal{I}'_a,\psi'}(\vec{s}')$ to $A_{T,\mathcal{I}_s,\mathcal{I}_a,\psi}(\vec{s})$, 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) $\mathsf{LWE}_{T',\mathcal{I}'_s,\mathcal{I}'_a,\psi',\ell}$ is polynomial-time equivalent to (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{I}_s,\mathcal{I}_a,\psi,\ell}$.

Proof. We begin by showing that the bases $\vec{b}_s, \vec{b}_a, \vec{b}_b$ yield a number of efficiently computable bijections, which we will utilize throughout the proof. We show this in detail for basis $\vec{b}_s$; the reasoning for the other bases $\vec{b}_a, \vec{b}_b$ is similar. Observe that since $\vec{b}_s$ is an $\mathcal{O}$-basis of a (full-rank) 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 $\mathcal{I}'_s = \mathcal{I}_s \mathcal{M}'_s$, it follows that φ induces a bijection from $\mathcal{I}'_s$ to $(\mathcal{I}_s)^k$: for if $x \in \mathcal{I}'_s$, then $\mathrm{Tr}(x \cdot \vec{b}_s^\vee) \in \mathcal{I}_s \cdot \mathrm{Tr}(\mathcal{M}'_s \vec{b}_s) = \mathcal{I}_s \cdot \mathcal{O}^k = (\mathcal{I}_s)^k$, and if $\vec{x} \in (\mathcal{I}_s)^k$, then $\langle \vec{x}, \vec{b}_s \rangle \in \mathcal{I}_s \mathcal{M}'_s = \mathcal{I}'_s$.

The claimed transform is as follows: given a sample $(\vec{a}', \vec{b}') \in (\mathcal{I}'_a / q\mathcal{I}'_a)^{k'_a} \times (K'_\mathbb{R} / q\mathcal{I}'_b)^{k'_b}$, we output

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

By what we showed above, $\mathrm{Tr}(\vec{a}' \otimes \vec{b}_a^\vee)$ simply extracts the unique 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 index T by (i', i), (j', j), and (ℓ', ℓ) where i', j', ℓ' are the indices of T', and i, j, ℓ are respectively the 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

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

[6]Note that $T \in \mathcal{O}_q^{k_s \times k_a \times k_b}$ 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 $\mathrm{Tr}(\mathcal{O}'_q \mathcal{M}'_s \mathcal{M}'_a (\mathcal{M}'_b)^{\vee_\mathcal{O}}) \subseteq \mathrm{Tr}(\mathcal{J}_q) \subseteq \mathcal{O}_q$, where $\mathcal{J} = \mathcal{M}'_b (\mathcal{M}'_b)^{\vee_\mathcal{O}}$.

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

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

as desired. Finally, the transform is efficiently invertible because the maps applied to $\vec{a}'$ and $\vec{b}'$ are both efficiently invertible (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 recover $\vec{s}$ from $\vec{s}'$, or vice versa, via the bijection φ or its inverse (from the first paragraph of the proof). $\square$

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 4.6, 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 $\mathbb{Z}_q$.

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 $\mathcal{I}_s = \mathcal{I}_a = \mathcal{I}_b = \mathbb{Z}$. Let $k_s = n + d - 1$ and $k_a = n$, and respectively identify $(\mathcal{I}_s/q\mathcal{I}_s)^{k_s} = \mathbb{Z}_q^{n+d-1}$ and $(\mathcal{I}_a/q\mathcal{I}_a)^n = \mathbb{Z}_q^n$ 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 $\vec{s} = (1, x, \dots, x^{n+d-2})$ and $\vec{a} = (x^{n-1}, x^{n-2}, \dots, 1)$. (Basis $\vec{a}$ is in decreasing order by degree for reasons that will become clear shortly.) Finally, let $k_b = d$ and identify $(\mathcal{I}_b/q\mathcal{I}_b)^{k_b} = \mathbb{Z}_q^d$ with $x^{n-1} \cdot \mathbb{Z}_q^{<d}[x]$ via the basis $\vec{b} = (x^{n-1}, x^n, \dots, x^{n+d-2})$.

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

$$M_{ijk} = \begin{cases} 1 & \text{if } i = j + k \\ 0 & \text{otherwise.} \end{cases}$$

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).[7] 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.

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) $\mathsf{MP\text{-}LWE}_{n,d,q,\psi,\ell}$ problem is simply the (search or decision, respectively) $\mathsf{LWE}_{M,\mathbb{Z},\mathbb{Z},\psi,\ell}$ problem, where M is the order-three tensor from Section 5.1.

We remark 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" $M_{(n+d-2)\cdot\cdot}$ and $M_{\cdot\cdot(d-1)}$, 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 $\mathcal{O}$ of a number field, a *power basis* of $\mathcal{L}$ is a $\mathbb{Z}$-basis $\vec{p}$ of $\mathcal{O}$ of the form $\vec{p} = (1, x, x^2, \ldots, x^{d-1})$ for some $x \in \mathcal{O}$.

Theorem 5.3. *Let $d \le n$ be positive integers; $\mathcal{O}$ be an order of a degree-d number field K having a known power basis $\vec{p} = (x^j)_{j=0,\ldots,d-1}$; ψ be a distribution over $K_{\mathbb{R}}$; and q be a positive integer. There is an efficient randomized reduction from (search or decision) $\mathcal{O}$-$\mathsf{LWE}_{q,\psi,\ell}$ to (search or decision, respectively) $\mathsf{MP\text{-}LWE}_{n,d,q,\psi',\ell}$, 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 an efficient deterministic reduction from (search or decision) $\mathcal{O}$-$\mathsf{LWE}_{q,\psi,\ell} = \mathsf{LWE}_{1,\mathcal{O}^\vee,\mathcal{O},q,\psi,\ell}$ to (search or decision, respectively) $\mathsf{LWE}_{T,\mathbb{Z},\mathbb{Z},\psi',\ell}$, where $T_{ijk} = \mathrm{Tr}(p_i^\vee p_j p_k) \in \mathbb{Z}_q^{d\times d\times d}$ for $\mathrm{Tr} = \mathrm{Tr}_{K/\mathbb{Q}}$. Observe that each "slice" $T_{i\cdot\cdot}$ is a $d \times d$ Hankel matrix, because $p_j p_k = x^{j+k}$ depends only on $j + k$. 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

[7] Recall that a matrix H is Hankel if each entry H_{jk} is determined by $j + k$ (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_{jk} for $i = j + k$.

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),

$$\sum_j T_{ijk'} \mathbf{A}_{jj'} = \sum_{i'} M_{i'j'k'} \mathbf{S}_{i'i} \,.$$

First, extend the power basis $\vec{p} = (x^j)_{j=0,\ldots,d-1}$ of $\mathcal{O}$ to $\vec{p}' = (x^{j'})_{j'=0,\ldots,n+d-2}$, by including more powers of x. 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, by linearity of the trace and the properties of dual bases, the left-hand side of Proof 18 is

$$\begin{aligned} T'_{ij'k'} &:= \sum_j T_{ijk'} \mathbf{A}_{jj'} \\ &= \sum_j \mathrm{Tr}(p_i^\vee p_j p_{k'}) \cdot \mathrm{Tr}(p_j^\vee \cdot p'_{j'}) \\ &= \mathrm{Tr}\Big(p_i^\vee \cdot \sum_j p_j \,\mathrm{Tr}(p_j^\vee \cdot p'_{j'}) \cdot p_{k'}\Big) \\ &= \mathrm{Tr}(p_i^\vee \cdot p'_{j'} \cdot p_{k'}) = \mathrm{Tr}(p_i^\vee \cdot x^{j'+k'}). \end{aligned}$$

Observe that each "slice" $T'_{i\cdot\cdot}$ 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 $T'_{i\cdot\cdot}$ can be written as an (efficiently computable) $\mathbb{Z}_q$-linear combination of the slices $M_{i\cdot\cdot}$, 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 $T'_{ij'k'} = \sum_{i'} M_{i'j'k'} \mathbf{S}_{i'i}$ 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 Proof 18 with suitable $\mathbf{A}, \mathbf{S}$, and hence Theorem 4.8 yields an efficient 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. □

Corollary 5.4. *Adopt the notation from Theorem 5.3, and let $\mathcal{O}' \subseteq \mathcal{O}$ be a suborder which has a known power basis $\vec{p}$ and for which $|\mathcal{O}/\mathcal{O}'|$ is coprime with q. There is an efficient randomized reduction from $\mathcal{O}$-$\mathsf{LWE}_{q,\psi,\ell}$ to $\mathsf{MP\text{-}LWE}_{n,d,q,\psi',\ell}$, where $\psi' = \mathrm{Tr}_{K_\mathbb{R}/\mathbb{R}}(\psi \cdot \vec{p})$.*

Proof. We reduce $\mathcal{O}$-$\mathsf{LWE}_{q,\psi,\ell}$ to $\mathcal{O}'$-$\mathsf{LWE}_{q,\psi,\ell}$ by Theorem 4.6, then to $\mathsf{MP\text{-}LWE}_{n,d,q,\psi',\ell}$ by Theorem 5.3. □

Remark 5.5. 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 can define $\mathsf{MP\text{-}LWE}_{\mathcal{O},n,d,q,\psi,\ell}$ as $\mathsf{LWE}_{M,\mathcal{O}^\vee,\mathcal{O},\psi,\ell}$, where $M \in \mathcal{O}_q^{(n+d-1)\times n\times d}$ is the order-three Hankel tensor whose entries are given by Section 5.1. Then for any degree-d extension ring $\mathcal{O}'/\mathcal{O}$ having a power $\mathcal{O}$-basis $\vec{p}$, Theorem 5.3 easily extends to give a reduction from $\mathcal{O}'$-$\mathsf{LWE}_{q,\psi,\ell}$ to $\mathsf{MP\text{-}LWE}_{\mathcal{O},n,d,q,\psi',\ell}$, where $\psi' = \mathrm{Tr}(\psi \cdot \vec{p})$.

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 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 5.6. *Let $d \leq n$ be positive integers, $\mathcal{O}$ be an order of a degree-d number field K with a power 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 an efficient randomized reduction from (search or decision) $\mathcal{O}$-$\mathsf{LWE}_{q,D_{\sqrt{\Sigma}},\ell}$ to (search or decision, respectively) MP-$\mathsf{LWE}_{n,d,q,D_{\sqrt{\Sigma'}},\ell}$.*

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

Proof. By applying Theorem 5.3 we obtain an efficient randomized reduction from $\mathcal{O}$-$\mathsf{LWE}_{q,D_{\sqrt{\Sigma}},\ell}$ to MP-$\mathsf{LWE}_{n,d,q,\psi',\ell}$, 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 $\mathbf{x}$. ☐

5.4 Example Instantiations

Corollary 5.6 bounds the expansion of the error distribution by the square root of the spectral norm of the Gram matrix $\mathbf{P}$ of a power 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 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 $i \in \{0, 1, \ldots, d-1\}$. 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 5.7. *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. $\square$

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 q be a positive integer. Then for any positive integer d', there is an efficient deterministic reduction from (search or decision) $\mathcal{O}'$-$\mathsf{LWE}^{d'}_{q,\psi',\ell}$ to (search or decision, respectively) $\mathcal{O}$-$\mathsf{LWE}^{d}_{q,\psi,\ell}$, where $d = rd'$ and $\psi = \mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}(\psi')$.*

Proof. First note that by Lemma 2.17, we have $(\mathcal{O}')^\vee = \mathcal{O}^\vee (\mathcal{O}')^{\vee_\mathcal{O}}$. Furthermore, $\vec{b}^\vee$ (where the dual is taken relative to K) 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 $\mathcal{I}_a = \mathcal{O}, \mathcal{I}_s = \mathcal{I}_b = \mathcal{O}^\vee$, which yield $\mathcal{I}'_a = \mathcal{O}', \mathcal{I}'_s = \mathcal{I}'_b = (\mathcal{O}')^\vee$. This gives an efficient deterministic reduction from (search or decision) $\mathcal{O}'$-$\mathsf{LWE}^{d'}_{q,\psi',\ell} = \mathsf{LWE}_{T',(\mathcal{O}')^\vee,\mathcal{O}',q,\psi',\ell}$, where $T'_{i'j'1} = \delta_{i'j'}$ is the $d' \times d' \times 1$ identity-matrix tensor over $\mathcal{O}'_q$, to (search or decision, respectively) $\mathsf{LWE}_{\tilde{T},\mathcal{O}^\vee,\mathcal{O},\tilde{\psi},\ell}$, where $\tilde{T}_{(i',i)(j',j)(1,k)} = \delta_{i'j'} \cdot \mathrm{Tr}_{K'/K}(b_i^\vee b_j b_k) \in \mathcal{O}_q^{d \times d \times r}$ and $\tilde{\psi} = \mathrm{Tr}_{K'_{\mathbb{R}}/K_{\mathbb{R}}}(\psi' \cdot \vec{b})$.

To complete the reduction to $\mathcal{O}$-$\mathsf{LWE}^d_{q,\psi,\ell}$, we will use Theorem 4.8 with an appropriate linear combination of the layers of $\tilde{T}$ to obtain the $d \times d \times 1$ identity-matrix tensor. Specifically, consider the $1 \times r$ matrix

$\mathbf{B} = \mathrm{Tr}_{K'/K}(\vec{b}^\vee)^t$, which has right inverse $r^{-1} \cdot \mathrm{Tr}_{K'/K}(\vec{b}) \in K^{r \times 1}$, and define

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

so T is the identity-matrix tensor, as desired. Similarly, $\mathbf{B}\tilde{\psi} = \psi$. Then, taking $\mathbf{A}$ and $\mathbf{S}$ to be identity matrices and invoking Theorem 4.8, we get an efficient deterministic reduction from (search or decision) $\mathsf{LWE}_{\tilde{T},\mathcal{O}^\vee,\mathcal{O},\tilde{\psi},\ell}$ to (search or decision, respectively) $\mathsf{LWE}_{T,\mathcal{O}^\vee,\mathcal{O},\psi,\ell} = \mathcal{O}\text{-}\mathsf{LWE}^d_{q,\psi,\ell}$. □

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^d 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 $\mathbf{A} = \mathrm{Tr}_{K'_\mathbb{R}/\mathbb{R}}(\vec{c}' \cdot \tau(\vec{c})^t)$ 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 an efficient randomized reduction from (search or decision) $\mathcal{O}'$-$\mathsf{LWE}^{d'}_{q,\psi',\ell}$ to (search or decision, respectively) $\mathcal{O}$-$\mathsf{LWE}^d_{q,\psi,\ell}$.*

Moreover, for $s = s'\sqrt{r}$, there is an efficient deterministic *reduction from (search or decision) $\mathcal{O}'$-$\mathsf{LWE}^{d'}_{q,D_{s'},\ell}$ to (search or decision, respectively) $\mathcal{O}$-$\mathsf{LWE}^d_{q,D_s,\ell}$.*

Proof. By Theorem 6.1, there exists an efficient deterministic reduction from $\mathcal{O}'$-$\mathsf{LWE}^{d'}_{q,\psi',\ell}$ to $\mathcal{O}$-$\mathsf{LWE}^d_{q,\psi_1,\ell}$ where $\psi_1 = \mathrm{Tr}_{K'_\mathbb{R}/K_\mathbb{R}}(\psi')$ 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,

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

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. ☐

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 $(1, \alpha, \ldots, \alpha^{r-1})$ 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 an efficient deterministic reduction from $\mathcal{O}'\text{-}\mathsf{LWE}^{d'}_{q,\psi',\ell}$ to $\mathcal{O}\text{-}\mathsf{LWE}^{d}_{q,\psi,\ell}$, where $d = rd'$.

References

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

[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`.

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

[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.

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

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

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

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

[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.

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

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

[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.

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

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

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

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

[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.

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

[RSSS17] M. Rosca, A. Sakzad, D. Stehlé, and R. Steinfeld. Middle-product learning with errors. In *CRYPTO*, pages 283–297. 2017.

[RSW18] M. Rosca, D. Stehlé, and A. Wallet. On the Ring-LWE and Polynomial-LWE problems. In *EUROCRYPT*, pages 146–173. 2018.

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