Simulation-Secure Threshold PKE from LWE with Polynomial Modulus

Daniele Micciancio and Adam Suhl

UC San Diego, United States

Abstract. In LWE based cryptosystems, using small (polynomially bounded) ciphertext modulus improves both efficiency and security. In threshold encryption, one often needs *simulation security*: the ability to simulate decryption shares without the secret key. Existing lattice-based threshold encryption schemes provide one or the other but not both. Simulation security has seemed to require superpolynomial flooding noise, and the schemes with polynomial modulus use Rényi divergence based analyses that are sufficient for game-based but not simulation security.
In this work, we give the first construction of simulation-secure lattice-based threshold PKE with polynomially bounded modulus. The construction itself is relatively standard, but we use an improved analysis, proving that when the ciphertext noise and flooding noise are both Gaussian, simulation is possible even with very small flooding noise. Our modulus is small not just asymptotically but also concretely: this technique gives parameters roughly comparable to those of highly optimized non-threshold schemes like FrodoKEM. As part of our proof, we show that LWE remains hard in the presence of some types of leakage; these results and techniques may also be useful in other contexts where noise flooding is used.

1 Introduction

A *threshold cryptosystem* allows to share a secret (decryption) key among a set of servers, in such a way that the servers can collaboratively decrypt messages, and still no set of servers (below a given threshold) can gain any information about the encrypted messages. Threshold encryption schemes are a fundamental tool in cryptography, both in theory (e.g., as a building block used in the construction of secure multiparty computation protocols [AJL$^+$12]) and in practice (as an effective solution to avoid the single point of failure associated to the secret key.) In order to promote further progress in the area, NIST (the National Institute of Standards and Technology) has recently issued a call for Multi-Party Threshold Schemes [BP23], aimed both at the development of a standard for threshold versions of mature "pre-quantum" schemes already standardized by NIST, and also the exploratory investigation of other primitives that are not yet a NIST standard. Of special interest among the more advanced primitives (and the main focus of this paper) is lattice-based encryption, which has already been selected as a candidate for post-quantum cryptography, and is expected to become an official standard within a year or so.

A threshold version of lattice based encryption (or, more specifically, Regev's cryptosystem [Reg09]) was first given by Bendlin and Damgard [BD10] in 2010. Based on the linear key-homomorphic properties of lattice based encryption, the BD10 scheme [BD10] gives a

Research supported in part by Global Research Cluster program of Samsung Advanced Institute of Technology, Intel Cryptography Frontiers program, and NSF Award 1936703. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

E-mail: daniele@cs.ucsd.edu (Daniele Micciancio), asuhl@ucsd.edu (Adam Suhl)

Date of this document: 2024-08-29.

very elegant and relatively efficient (non-interactive) solution to the threshold decryption problem, where each server locally computes a "partial decryption", and then these partial decryptions are simply added up and rounded to the final output message. The main drawback of the BD10 scheme is that in order to protect the secret key during the partial decryption process, it uses *noise flooding.* This is a masking technique that requires the use of superpolynomially large noise, and correspondingly large "ciphertext modulus" q.[1] This has a negative impact both on efficiency (requiring computations modulo a large q, and the use of "big-int" large precision arithmetic libraries) and security (requiring the assumption that lattice problems are hard to solve within superpolynomial factors $\approx q$.) So, since the publication of [BD10], it has been an important open problem in the area to develop an efficient threshold cryptosystem using a polynomial modulus q and based on the hardness of approximating lattice problems within (small) polynomial factors.

Our contribution In this paper we give an efficient lattice-based threshold encryption scheme

- satisfying a strong (simulation based) notion of security (which is relevant to the use of threshold cryptography in MPC applications)
- supporting an arbitrary (polynomial) number of decryption queries, and
- using a polynomial modulus q (which results in a standard hardness assumption of approximating lattice problems within a polynomial approximation factor.)

Some progress towards these goals had recently been obtained in [CSS$^+$22, BS23], which used Renyi divergence techniques to analyze threshold encryption with polynomial modulus, but only achieving a weaker (game-based) definition of security and assuming an a-priori polynomial upper bound on the number of decryption queries.[2] So, to the best of our knowledge, our is the first work achieving these strong security properties with a polynomial modulus and inapproximability assumption.

On the technical side, we give a general construction and analysis technique that is applicable to a broad range of lattice-based encryption schemes, including Regev's encryption [Reg09] (as already done in [BD10] using superpolynomial noise flooding), and also more efficient variants like [LP11] that use the LWE problem both during key generation and encryption. In fact, for the case of Regev's cryptosystem, our scheme is essentially the same as BD10, and the main contribution is in the analysis technique, which may be of independent interest and find additional applications. This results in a very simple design, with a concrete efficiency (both in terms of running time, and key and ciphertext size) essentially the same as the basic (non-threshold) version of the schemes. We exemplify the practicality of the scheme considering a threshold version of a scheme similar to Frodo [NAB$^+$20], a highly optimized scheme that was among the leading NIST candidates for post-quantum cryptography based on the general (non-ring) Learning With Errors (LWE) problem.

In this paper we focus on (plain) LWE, where we can prove security of our scheme from standard assumptions. However, most of our results are easily adapted to the Ring LWE ([SSTX09, LPR10]) setting, which provides even more efficient constructions based on stronger (but still standard) hardness assumptions on algebraic lattices. In fact, the only lemma in this paper that does not easily adapt to the ring setting is the proof that

[1] In lattice-based cryptography, ciphertexts consists of vectors or matrices with integer entries modulo q, for some positive q called the ciphertext modulus.

[2] Specifically, the modulus q (and lattice inapproximability factor in the complexity assumption) in these works scales with $\sqrt{\ell}$, where ℓ is the number of decryption queries performed by the attacker. So, supporting an arbitrary polynomial number of queries ℓ still requires a superpolynomial modulus q. Moreover, this limitation seems intrinsic to the use of Renyi divergence techniques. See Section 1.2 for a more detailed comparison.

a certain variant of LWE with known error norm is equivalent to the standard LWE problem. (See Lemma 9.) In Section 5.3 we introduce the ring version of this problem as an assumption, which we call “Known-Covariance RLWE”, under which we sketch an RLWE-based PKE scheme from which we can construct threshold cryptosystems. We leave it as an open problem to either provide an attack showing that this generalization of LWE with known error norm is false in the ring setting, or demonstrate that it is equivalent to the standard Ring LWE problem and worst-case assumptions on the approximability of algebraic lattices.

1.1 Technical Overview

A common framework for Threshold LWE encryption is as follows: Ciphertexts are of the form $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e_{ct} + msg)$, so for decryption the parties need to compute $\langle \mathbf{a}, \mathbf{s}\rangle$ collectively. Each party gets a share $\mathbf{s_i}$ of $\mathbf{s}$ under some linear secret sharing scheme. (For example, for T-out-of-T threshold encryption one could have $\sum_i \mathbf{s_i} = \mathbf{s}$.) Party i can then compute $\langle \mathbf{a}, \mathbf{s_i}\rangle$ on its own, and by linearity the shares $\langle \mathbf{a}, \mathbf{s_i}\rangle$ can be combined to reconstruct $\langle \mathbf{a}, \mathbf{s}\rangle$. However, because revealing $\langle \mathbf{a}, \mathbf{s_i}\rangle$ would leak information about $\mathbf{s_i}$, each party instead reveals a noisy version $\langle \mathbf{a}, \mathbf{s_i}\rangle + \tilde{e}$ (where $\tilde{e}$ is called the “flooding” or “smudging” noise). This means when the decryption shares are combined to reconstruct $\langle \mathbf{a}, \mathbf{s}\rangle$, the result is noisy, but since the ciphertext is noisy anyway, the parameters of the scheme just need to be adjusted so that decryption succeeds even with the extra noise.

An adversary is able to learn $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e_{ct})$ from the ciphertext and $\langle \mathbf{a}, \mathbf{s}\rangle + \tilde{e}$ from the decryption share. The difficulty when proving security is simulating the latter given the former, but without knowing $\mathbf{s}$ or e_{ct}.

One approach is to have $\tilde{e}$ come from a very wide distribution, say a Gaussian $\mathcal{N}_\sigma$ for large σ, such that $\mathcal{N}_{0,\sigma}$ is statistically close to $\mathcal{N}_{e_{ct},\sigma}$. Then we can simulate partial decryption by sampling $e' \leftarrow \mathcal{N}_{0,\sigma}$ and returning $\langle \mathbf{a}, \mathbf{s}\rangle + e_{ct} + e'$. This is the approach taken by [BD10]. Unfortunately, it requires σ to be superpolynomially larger than e_{ct}.

One could use smaller σ anyway, and while $\mathcal{N}_{0,\sigma}$ and $\mathcal{N}_{e_{ct},\sigma}$ have non-negligible statistical distance, they have small Rényi divergence. Under the right conditions, this means an adversary’s advantage in a security game must remain small if real decryption shares are replaced with simulated ones. This is the approach taken by [CSS$^+$22, BS23]. Unfortunately, this is insufficient for simulation-based security, because the output of the simulator can still be distinguished from real decryption shares.

The main insight of this paper is that the real and simulated distributions just need to be *computationally* indistinguishable, not statistically indistinguishable, and for computationally bounded adversaries, the ciphertext looks uniform and e_{ct} is unknown. $\mathcal{N}_{0,\sigma}$ and $\mathcal{N}_{e_{ct},\sigma}$ being distinguishable is not necessarily a problem. To simulate $\langle \mathbf{a}, \mathbf{s}\rangle + \tilde{e}$, we need to know how $e_{ct} - \tilde{e}$ should be distributed. But rather than viewing e_{ct} as fixed (so that the distribution is $\mathcal{N}_{e_{ct},\sigma}$) we take the distribution over the uncertainty of e_{ct} as well as $\tilde{e}$.

For example, if $e_{ct} \leftarrow \mathcal{N}_t$ and $\tilde{e} \leftarrow \mathcal{N}_\sigma$, then the difference between $\langle \mathbf{a}, \mathbf{s}\rangle + e_{ct}$ and $\langle \mathbf{a}, \mathbf{s}\rangle + \tilde{e}$ is distributed as $\mathcal{N}_{\sqrt{\sigma^2+t^2}}$. As long as the ciphertext $(\mathbf{a}, b)$ is computationally uniform, we can sample $e' \leftarrow \mathcal{N}_{\sqrt{\sigma^2+t^2}}$ and simulate partial decryption as $b + e'$.

It’s not immediately obvious that $\langle \mathbf{a}, \mathbf{s}\rangle + e_{ct}$ will still look uniform if $\langle \mathbf{a}, \mathbf{s}\rangle + \tilde{e}$ has been revealed — $\mathbf{a}$ and $\mathbf{s}$ have both been reused, and LWE might no longer be hard in the presence of this sort of leakage. However, when e_{ct} and $\tilde{e}$ are both Gaussians, this “Reused-A LWE” problem is as hard as standard LWE with slightly smaller noise, as was already proved implicitly in [KY16] and [KLSS23]. This allows us to use smudging noise that is polynomially large, and in fact potentially smaller than the ciphertext noise!

One technical detail is that to properly simulate, we need to know the variance of the ciphertext noise — we can’t sample from $\mathcal{N}_{\sigma^2+t^2}$ without knowing t. In Section 5 we slightly modify the schemes of [Reg09] and [LP11] to produce ciphertexts with known noise variance.

We use continuous Gaussian noise for simplicity of exposition, but everything can be adapted to the discrete case using standard discrete Gaussian convolution theorems (e.g. [GMPW20]) with small increase in Gaussian parameter (typically a $\sqrt{2}$ factor).

1.2 Related Work

A number of works build Threshold PKE from lattices. In [BD10], Bendlin and Damgård build Threshold PKE that is UC-secure, but the modulus is superpolynomially large. Singh, Rangan, and Banerjee build Threshold PKE with polynomial modulus in [SRB13], but they achieve only a weak form of semantic security.

Other works use lattices to build threshold versions of stronger primitives like IBE. Bendlin, Krehbiel, and Peikert in [BKP13] build IBE with threshold key generation / extraction / delegation; their construction uses polynomial modulus and they prove the security of threshold key generation / extraction / delegation in the UC framework. However, that work does not consider threshold decryption. In [KM16], Kuchta and Markowitch build IBE with threshold decryption, but under a weaker security model that does not let the adversary see partial decryption shares. In [DDE$^+$23] Dahl et al. build a simulation-secure Threshold FHE (and thus also PKE) scheme that uses a relatively small modulus during evaluation, but during partial decryption they switch to an superpolynomially large modulus and then bootstrap. Boneh et al in [BGG$^+$18] build simulation secure Threshold FHE (which they use to build a "universal thresholdizer" and then numerous other threshold primitives) but with superpolynomial modulus. By applying their "universal thresholdizer" to a non-threshold scheme with small modulus, they can reduce their ciphertext modulus during evaluation, but at the cost of expensive partial decryption that performs homomorphic operations in the original large-modulus Threshold FHE scheme. Chowdhury et al [CSS$^+$22] and Boudgoust and Scholl [BS23] both build Threshold FHE with polynomial modulus (including during partial decryption), proving security under slightly different game-based security definitions; both proofs use Rényi divergence arguments that are sufficient for game-based security, but insufficient for simulation security. Furthermore, because of the Rényi divergence technique, both [CSS$^+$22] and [BS23] need a bound ℓ on the number of decryption queries to be known in advance, and the modulus scales with $\sqrt{\ell}$; in our scheme the modulus need not grow with the number of decryption queries.

The security notion we consider is CPA-like and assumes static corruptions. In [DLN$^+$21], Devevey et al. build lattice-based threshold PKE that achieves CCA2 security against adaptive corruptions. However, their construction uses noise flooding with super-polynomial modulus-to-noise ratio. Combining our techniques with those of [DLN$^+$21] to achieve CCA2 security with polynomially large modulus would be an interesting direction for future work.

One last technique for noise-flooding with polynomial modulus is "gentle noise flooding", first introduced in [BPMW16], later used for the analysis of entropic LWE in [BD20b, BD20a], and used in [dCHI$^+$22] to achieve (non-threshold) homomorphic encryption with circuit privacy. Here the goal is to avoid leakage from the plaintext, rather than the key, and the technique does not seem applicable to achieve threshold decryption.

Threshold decryption and key generation can also be performed using general MPC techniques [KLO$^+$19], without any noise flooding, and keeping the same LWE (polynomial) encryption modulus. It may be possible to adapt the multi-key FHE construction from MrNISC (reusable, non-interactive MPC) with polynomial modulus of [BJKL21, Shi22]. However, these methods are based on general MPC techniques and unlikely to be practical.

To the best of our knowledge, ignoring schemes based on generic MPC, ours is the first LWE-based Threshold PKE scheme with polynomially large modulus that achieves simulation security.

1.3 Outline

Section 2 recalls some results from previous works and proves some technical lemmas that will be used in our proofs.

In Section 3 we prove that LWE in the presence of certain kinds of leakage is as hard as standard LWE, which may be of independent interest.

In Section 4 we build a simulation-secure Threshold PKE scheme with polynomial modulus from any LWE-based PKE scheme satisfying certain properties. Informally, we require ciphertexts in the underlying scheme must look like fresh LWE samples (plus the encoded message), even when the secret key is known. Furthermore, the error distribution must be a (continuous) Gaussian whose standard deviation is publicly known.

In Section 5 we show a PKE scheme satisfying these conditions, namely a slightly modified version of Lindner and Peikert's scheme from [LP11]. We present another such scheme, a slight modification of Regev's scheme from [Reg09], in Appendix A.1; it is simpler but less practical. Both schemes are secure under standard assumptions. We also sketch a Ring-LWE based scheme whose security relies on a (plausible but nonstandard) assumption.

Finally in Section 6 we give example concrete parameters.

2 Preliminaries

2.1 Gaussians

Definition 1 (Continuous Gaussians)**.** The (one-dimensional) Gaussian measure $\rho_{c,s}$ with center c and width[3] s is defined as

$$\rho_{c,s}(x) = e^{-\pi(x-c)^2/s^2}.$$

More generally in n dimensions,

$$\rho_{\mathbf{c},s}(\mathbf{x}) = e^{-\pi\|\mathbf{x}-\mathbf{c}\|^2/s^2} = \prod_{i=1}^{n} \rho_{c_i,s}(x_i).$$

The n-dimensional spherical Gaussian distribution $\mathcal{N}^n_{\mathbf{c},s}$ is the distribution over $\mathbb{R}^n$ whose probability density is proportional to $\rho_{\mathbf{c},s}$. Equivalently,

$$\mathcal{N}^n_{\mathbf{c},s}(\mathbf{x}) = \frac{1}{s^n}\prod_{i=1}^{n} \rho_{c_i,s}(x_i).$$

For brevity, we will usually write $\mathcal{N}^n_s$ instead of $\mathcal{N}^n_{0,s}$ (and likewise ρ_s instead of $\rho_{0,s}$) for Gaussians centered at zero.

Definition 2 (Discrete Gaussian)**.** The discrete Gaussian distribution over an n-dimensional lattice Λ, written $\mathcal{D}_{\Lambda,\mathbf{c},s}$, is the distribution over Λ whose probability density is proportional to $\rho_{\mathbf{c},s}$. Equivalently,

$$\mathcal{D}_{\Lambda,\mathbf{c},s}(\mathbf{x}) = \frac{\rho_{\mathbf{c},s}(\mathbf{x})}{\rho_{\mathbf{c},s}(\Lambda)}.$$

Wide-enough discrete Gaussians behave in many ways like continuous Gaussians. To quantify how wide is wide enough, we use a lattice parameter known as the *smoothing parameter*. For a full definition of the smoothing parameter we refer the reader to [MR07]; for our purposes, the following fact will be sufficient:

[3] This is not the standard deviation; a Gaussian of width s has standard deviation $s/\sqrt{2\pi}$.

Lemma 1 ([MR07, Lemma 3.3])**.** *The smoothing parameter of $\mathbb{Z}^n$, written $\eta_\epsilon(\mathbb{Z}^n)$, is $\tilde{O}(1)$ for some negligible ϵ. More precisely,*

$$\eta_\epsilon(\mathbb{Z}^n) \leq \sqrt{\frac{\ln(2n(1+1/\epsilon))}{\pi}}$$

Shifting the center $\mathbf{c}$ of a discrete Gaussian $\mathcal{D}_{\mathbb{Z}^n,\mathbf{c},s}$ leaves the normalization factor $\rho_{\mathbf{c},s}(\mathbb{Z}^n)$ almost unchanged, assuming the width s is above the smoothing parameter:

Lemma 2 ([Reg09, Claim 3.8], special case where lattice is $\mathbb{Z}^n$)**.** *For any $\mathbf{c} \in \mathbb{R}^n$, $\epsilon > 0$, and $r \geq \eta_\epsilon(\mathbb{Z}^n)$,*

$$r^n(1-\epsilon) \leq \rho_{\mathbf{c},r}(\mathbb{Z}^n) \leq r^n(1+\epsilon).$$

The convolution of a wide-enough discrete Gaussian and a wide-enough continuous Gaussian is close to a continuous Gaussan:

Lemma 3 ([Reg09, Claim 3.9], special case where lattice is $\mathbb{Z}^n$)**.** *Let $r, s > 0$ be two reals, and let t denote $\sqrt{r^2+s^2}$. Assume that $rs/t = 1/\sqrt{1/r^2+1/s^2} \geq \eta_\epsilon(\mathbb{Z}^n) = \tilde{O}(1)$ for some $\epsilon < \frac{1}{2}$. Consider the continuous distribution Y on $\mathbb{R}^n$ obtained by sampling from $\mathcal{D}_{\mathbb{Z}^n,0,r}$ and then adding a noise vector taken from $\mathcal{N}^n_{0,s}$. Then, the statistical distance between Y and $\mathcal{N}^n_{0,t}$ is at most 4ϵ.*

We will in fact need a stronger result: suppose we have a discrete Gaussian vector $\mathbf{x}$ and add some continuous Gaussian noise $\mathbf{e}$. Assuming both Gaussians are wide enough, not only is $\mathbf{x}+\mathbf{e}$ close to a continuous Gaussian, but also if we condition on the value of $\mathbf{x}+\mathbf{e}$, then the conditional distribution of $\mathbf{x}$ will still be close to a discrete Gaussian:

Lemma 4 ([GMPW20, Corollary 4.2])**.** *Let $s, t \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^n)$ for some negligible ϵ. Then the following distributions are statistically close:*

$$D_1 = \{\mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^n,0,s};\ \mathbf{e} \leftarrow \mathcal{N}^n_t\ :\ (\mathbf{x}, \mathbf{x}+\mathbf{e})\}$$

$$D_2 = \{\mathbf{y} \leftarrow \mathcal{N}^n_\alpha;\ \mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\beta\mathbf{y},\gamma}\ :\ (\mathbf{x}, \mathbf{y})\}$$

where $\alpha = \sqrt{s^2+t^2}$, $\beta = \frac{s^2}{s^2+t^2}$ and $\gamma = \frac{st}{\sqrt{s^2+t^2}}$.

Proof. This is just a special case of Corollary 3 of [GMPW20] (Corollary 4.2 in the ePrint version), letting $\delta = t/\alpha$, $\delta' = s/\alpha$, $A_1 = \Lambda_1 = \frac{1}{s}\mathbb{Z}^n$, $A_2 = \Lambda_2 = \mathbb{R}^n$, and where our $\mathbf{x} = s\mathbf{x}_1$ and $\mathbf{y} = \alpha\mathbf{x}_2$. □

Letting $s = t$ immediately gives the following corollary:

Corollary 1. *Let $s \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^n)$. Then the following distributions are statistically close:*

$$D_1 = \{\mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}^n,0,s};\ \mathbf{e} \leftarrow \mathcal{N}^n_s\ :\ (\mathbf{r}, \mathbf{r}+\mathbf{e})\}$$

$$D_2 = \{\mathbf{v} \leftarrow \mathcal{N}^n_{0,\sqrt{2}s};\ \mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\mathbf{v}/2,s/\sqrt{2}}\ :\ (\mathbf{u}, \mathbf{v})\}.$$

We have the following tail bound on the norm of a discrete Gaussian vector:

Lemma 5 ([Ban93, Lemma 1.5])**.** *Let $L \subset \mathbb{R}^n$ be any lattice, let $s > 0$, and let $c \geq 1/\sqrt{2\pi}$. Then $\Pr_{\mathbf{x}\leftarrow\mathcal{D}_{L,0,s}}[\|\mathbf{x}\| > cs\sqrt{n}] < (c\sqrt{2\pi e}e^{-\pi c^2})^n$*

Plugging in $c = 0.8$ gives the following corollary:

Corollary 2. *If $\mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^n,0,s}$ then $\|\mathbf{x}\| < 0.8s\sqrt{n}$ with probability at least $1 - 2^{-n}$.*

2.2 Threshold PKE

Definition 3 (Threshold PKE)**.** A t-of-T *threshold PKE* scheme (also called a *PKE scheme with threshold decryption*) is a tuple of algorithms $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Rec})$:

$\mathsf{KeyGen}()$ outputs a list of T secret keys (one for each party) and a single public key pk.

$\mathsf{Enc}(\mathsf{pk}, \mathsf{msg})$ outputs a ciphertext ct.

$\mathsf{Dec}(\mathsf{sk}_i, \mathsf{ct})$ outputs a "partial decryption" (or "decryption share") d_i.

$\mathsf{Rec}(\mathsf{ct}, \mathsf{pk}, \{(i, d_i)\}_{i \in S})$, where $S \subseteq [T]$ is a set of t indices, reconstructs a decrypted message m'.

The scheme satisfies correctness if, for all messages msg and all sets $S \subseteq [T]$ of size t, the following experiment returns true with overwhelming probability:

```
$(\mathsf{sk}_1, \ldots, \mathsf{sk}_T, \mathsf{pk}) \leftarrow \mathsf{KeyGen}()$
$\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, \mathsf{msg})$
for $i \in S$ do
    $d_i \leftarrow \mathsf{Dec}(\mathsf{sk}_i, \mathsf{ct})$
$m' = \mathsf{Rec}(\mathsf{ct}, \mathsf{pk}, \{(i, d_i)\}_{i \in S})$
return $(\mathsf{msg} = m')$
```

2.2.1 Game-Based Security

The following security notion for Threshold PKE is very weak; it ignores threshold decryption entirely and merely requires the (non-threshold) encryption scheme remain IND-CPA secure even if the adversary compromises some of the secret keyshares.

Definition 4 ((Weak) Threshold IND-CPA)**.** Let $TPKE = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Rec})$ be a k-of-T threshold PKE scheme. We say the scheme is *(weakly) IND-CPA secure* if for all stateful PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, the following experiment returns true with probability negligibly far from $\frac{1}{2}$:

```
procedure $\mathsf{Expt}_{\mathsf{IND-CPA}}(\mathcal{A})$
    $b \xleftarrow{\$} \{0, 1\}$
    $(\mathsf{sk}_1, \ldots, \mathsf{sk}_T, \mathsf{pk}) \leftarrow \mathsf{KeyGen}()$
    $S_{mal} \leftarrow \mathcal{A}_1(\mathsf{pk})$ where $S_{mal} \subset [T]$ and $|S_{mal}| = k - 1$
    $b' \leftarrow \mathcal{A}_2^{\mathsf{LR}}(\{\mathsf{sk}_i\}_{i \in S_{mal}})$
    return $(b' = b)$
```

where the oracle $\mathsf{LR}(m_0, m_1)$ returns $\mathsf{Enc}(\mathsf{pk}, m_0)$ if $b = 0$ and $\mathsf{Enc}(\mathsf{pk}, m_1)$ if $b = 1$.

The following security notion is much stronger; it is analogous to the Threshold PKE security definition given in [BGG+18, full version, Definition 8.27] but for IND-CPA instead of IND-CCA, and is equivalent to the "ℓ-$\mathsf{IND\text{-}CPA}$ for ThPKE" definition of [BS23] but without the bound ℓ on the number of queries:

Definition 5 (Threshold IND-CPA-D)**.** Threshold IND-CPA-D security is defined as in Definition 4 except that $\mathcal{A}_2$ may also make polynomially many adaptive queries to the following oracle:

```
procedure $\mathsf{D}(m)$
    $ct \leftarrow \mathsf{Enc}(\mathsf{pk}, m)$
    for all $i \in [T] \setminus S_{mal}$ do
        $d_i \leftarrow \mathsf{Dec}(\mathsf{sk}_i, ct)$
    return $(ct, \{d_i\}_{i \in [T] \setminus S_{mal}})$
```

2.2.2 Simulation Security

We will use the following security notion for a T-out-of-T threshold PKE scheme. We assume an honest-but-curious adversary that knows all but one keyshare, and can ask the honest party to partially decrypt honestly-generated encryptions of adversarially chosen messages. Informally, simulation security means the adversary can by themself simulate the honest party's partial decryptions, without the help of the honest party or any access to the honest party's secret keyshare, as long as the adversary knows the underlying plaintext. This implies that the partial decryptions reveal nothing more than the underlying plaintexts and cannot help the adversary break the scheme.

More formally,

Definition 6. A T-out-of-T threshold PKE scheme is *simulation secure* if

- the scheme is weakly threshold IND-CPA secure, and
- there is an efficient algorithm Sim such that, with overwhelming probability over $(\mathsf{pk}, \mathsf{sk}^{hon}, \mathsf{sk}^{mal}) \leftarrow \mathsf{KeyGen}$ (letting sk^{mal} denote the set of compromised keys), and for all (possibly adaptively chosen) sequences of messages $\{m_i\}$,

  $$\big\{\mathsf{ct}_i \leftarrow \mathsf{Enc}_{\mathsf{pk}}(m_i)\,\forall i \quad : \quad (\mathsf{pk}, \mathsf{sk}^{mal}, \{m_i, \mathsf{ct}_i, \mathsf{Dec}_{\mathsf{sk}^{hon}}(\mathsf{ct}_i)\})\big\}$$

  is computationally indistinguishable from

  $$\big\{\mathsf{ct}_i \leftarrow \mathsf{Enc}_{\mathsf{pk}}(m_i)\,\forall i \quad : \quad (\mathsf{pk}, \mathsf{sk}^{mal}, \{m_i, \mathsf{ct}_i, \mathsf{Sim}_{\mathsf{pk},\mathsf{sk}^{mal}}(\mathsf{ct}_i, m_i)\})\big\}.$$

We remark that in general one would also let the adversary call Dec on the same ciphertext more than once, but in our constructions Dec will be deterministic, so this definition is equivalent.

Lemma 6. *Simulation security implies Threshold IND-CPA-D security.*

Proof. Let $\mathcal{A}$ be a Threshold IND-CPA-D adversary; make a Threshold IND-CPA adversary $\mathcal{B}$ by replacing each of $\mathcal{A}$'s D queries $D(m_i)$ with $\mathsf{ct}_i \leftarrow \mathsf{Enc}_{\mathsf{pk}}(m_i)$ followed by $d_i \leftarrow \mathsf{Sim}_{\mathsf{pk},\mathsf{sk}^{mal}}(\mathsf{ct}_i, m_i)$. By the simulatability property, $\mathcal{A}$ and $\mathcal{B}$ have computationally indistinguishable output. But the scheme is Threshold IND-CPA secure, so $\mathcal{B}$ (and thus $\mathcal{A}$) must have negligible advantage. □

In fact, simulation security is *strictly* stronger than Threshold IND-CPA-D security, as shown in Lemma 15 in the Appendix.

2.3 LWE-like ciphertexts

Definition 7. Fix a parameter σ_{dLWE}. We say a (non-threshold) PKE scheme with key generation KeyGen and encryption Enc has *LWE-like ciphertexts* if it satisfies the following conditions:

1. The secret key is a vector $\mathbf{s} \in \mathbb{Z}_q^n$.
2. Decisional LWE with secret $\mathbf{s}$ and Gaussian noise of width σ_{dLWE} remains hard when the public key is known. That is, for all polynomially large m,

   $$\{(\mathbf{s}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}();\ \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n};\ \mathbf{e} \xleftarrow{\$} \mathcal{N}^m_{\sigma_{dLWE}}\ :\ (\mathsf{pk}, \mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})\} \approx_c$$
   $$\{(\mathbf{s}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}();\ \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n};\ \mathbf{b} \xleftarrow{\$} \mathbb{R}_q^m\ :\ (\mathsf{pk}, \mathbf{A}, \mathbf{b})\}$$

3. Even conditioned on the secret key (and public key), the output of $\mathsf{Enc}(msg)$[4] (as a distribution over the randomness used in Enc, not in KeyGen) is indistinguishable from

   $$\{\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n;\ e_{ct} \leftarrow \mathcal{N}_\sigma\ :\ (\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e_{ct} + msg)\}$$

[4] Here msg is already scaled or otherwise encoded; we ignore the details of the encoding.

for some $\sigma \geq \sigma_{dLWE}$ that may depend on pk. In other words, each ciphertext looks like the message plus a fresh LWE sample with continuous Gaussian noise of width σ.

Lemma 7. *If a PKE scheme has LWE-like ciphertexts, then its ciphertexts are pseudorandom conditioned on the public key, and thus the scheme is IND-CPA secure.*

Proof. This follows immediately from conditions 3 and 2. □

Definition 8. A scheme with LWE-like ciphertexts has *public error width* if the width σ of the distribution of e_{ct} is publicly, efficiently computable from pk.

Note that having σ-LWE-like ciphertexts (and specifically Condition 3) is a much stronger property than is generally needed from (or proved about) a PKE scheme, even for schemes with ciphertexts that can be written as $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e + msg)$. It is unusual to require anything beyond correctness when $\mathbf{s}$ is known, but here we require that $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e_{ct})$ look like a *fresh, random* LWE sample even given $\mathbf{s}$. For example, if $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e_{ct})$ is generated by combining LWE samples from the public key, perhaps knowing $\mathbf{s}$ (and thus also the errors in the public key LWE samples) would allow finding correlations between $\mathbf{a}$ and e_{ct}; such a scheme could still be IND-CPA secure but would not have σ-LWE-like ciphertexts. Likewise, a scheme where the ciphertext noise distribution depends on the secret key could be IND-CPA secure but not have public error width. We will show in Section 5 that, with slight modifications, both standard Regev PKE ([Reg09]) and the compact variant of Lindner and Peikert ([LP11]) have LWE-like ciphertexts with public error width, with polynomially large modulus.

3 LWE Variants

For security we will rely on a few variants of the usual LWE assumption. All are provably as hard as the standard LWE assumption.

3.1 Reused-A LWE

Definition 9. The Reused-A LWE distribution with parameters n, m, σ_1, σ_2 is defined as

$$\left\{ \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n};\ \mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n;\ \mathbf{e_1} \leftarrow \mathcal{N}_{\sigma_1}^m;\ \mathbf{e_2} \leftarrow \mathcal{N}_{\sigma_2}^m\ :\ (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e_1}, \mathbf{A}\mathbf{s} + \mathbf{e_2}) \right\}.$$

The Search Reused-A LWE problem is to recover $\mathbf{s}$.

The Decision Reused-A LWE problem is to distinguish the distribution from

$$\left\{ \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n};\ \mathbf{b}' \xleftarrow{\$} \mathbb{R}_q^m;\ \mathbf{c} \leftarrow \mathcal{N}_{\sqrt{\sigma_1^2 + \sigma_2^2}}^m\ :\ (\mathbf{A}, \mathbf{b}', \mathbf{b}' + \mathbf{c}) \right\}.$$

For notational convenience we present this in "matrix form" where $\mathbf{A}$ is a matrix of m vectors $\mathbf{a_i}$. We can also consider the case where m is polynomially large but not known in advance – i.e., the adversary can make polynomially-many queries to an oracle that outputs samples $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e_1, \langle \mathbf{a}, \mathbf{s} \rangle + e_2)$ for fixed unknown $\mathbf{s}$.

Lemma 8. *Search Reused-*$\mathbf{A}$ *LWE is as hard as search LWE with Gaussian noise of width* $\frac{1}{\sqrt{\sigma_1^{-2} + \sigma_2^{-2}}}$. *Furthermore, if decision LWE is hard with Gaussian noise of width* $\frac{1}{\sqrt{\sigma_1^{-2} + \sigma_2^{-2}}}$, *then decision Reused-*$\mathbf{A}$ *LWE is also hard.*

This result is not new; more general versions (for discrete Gaussian noise) are proved in prior works. It follows from [KY16, Lemma 1] setting $\mathbf{V} = [\mathbf{I}|\mathbf{I}]$, and can also be viewed as a special case of the Hint MLWE problem of [KLSS23]. The special case we use here is simpler to state and suffices for our work; for completeness we give a direct proof with continuous Gaussian noise in the appendix.

As an aside, we remark that the hardness proofs for Reused-$\mathbf{A}$ LWE do not go through if the noise is non-Gaussian. In fact, for polynomially large bounded uniform noise, there is an explicit attack!

Claim. *There is a poly-time adversary that solves Search Reused-$\mathbf{A}$ LWE if errors, instead of Gaussian, are uniform in the interval $[-B, B]$ for polynomially large B, given $m = \Omega(B^2 n)$ samples.*

Proof. For each sample $(\mathbf{a}, b_1 = \langle \mathbf{a}, \mathbf{s} \rangle + e_1, b_2 = \langle \mathbf{a}, \mathbf{s} \rangle + e_2)$, check whether $b_2 - b_1 = 2B$. If so, it must be the case that $e_2 = B$ and $e_1 = -B$, and we learn the value of $\langle \mathbf{a}, \mathbf{s} \rangle$. This will happen with probability $\frac{1}{(2B+1)^2} = O(1/B^2)$ — non-negligible because B is only polynomially large. With $\Omega(B^2 n)$ samples we expect to recover the exact value of $\langle \mathbf{a}, \mathbf{s} \rangle$ for n different values of $\mathbf{a}$, from which we can recover $\mathbf{s}$. □

3.2 Known-Norm LWE

Definition 10 (LWE with known norm)**.** Given integers n, m, q, with m and q polynomially large in n, and an error distribution χ whose support is in $\mathbb{Z}_q$, the (small-secret) Known-Norm LWE Distribution is defined as

$$\left\{ \mathbf{s} \leftarrow \chi^n;\ \mathbf{e} \leftarrow \chi^m;\ \mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}\ :\ (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, \|\mathbf{s}\|^2 + \|\mathbf{e}\|^2) \right\}.$$

In other words, Known-Norm LWE is small-secret LWE except the adversary is also given the ℓ_2 norm of the vector $[\mathbf{s}|\mathbf{e}] \in \mathbb{Z}^{n+m}$.

The decisional Known-Norm LWE assumption is that this distribution is computationally indistinguishable from

$$\left\{ \mathbf{s} \leftarrow \chi^n;\ \mathbf{e} \leftarrow \chi^m;\ \mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}\ :\ (\mathbf{A}, \mathrm{Unif}, \|\mathbf{s}\|^2 + \|\mathbf{e}\|^2) \right\}.$$

Lemma 9. *The (decisional) Known-Norm LWE problem is as hard as (decisional) small-secret LWE with the same parameters, up to a polynomial factor in the advantage.*

Proof. $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2$ is no more than $q^2(m+n)$, which is polynomially large. For search, given a small-secret LWE instance, an adversary that solves the Known-Norm LWE problem can just guess the value of $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2$, and will guess correctly with at least $1/poly(n)$ probability. For the decision variant, the search-to-decision reduction of [MM11] extends to Known-Norm LWE, as described in Appendix B. □

The above argument is extremely loose. Concretely, when χ is a discrete Gaussian of small width, $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2$ will with overwheming probability be much smaller than $q^2(m + n)$, allowing better concrete parameters for a given security level.

We can also define a version of Known-Norm LWE where the secret $\mathbf{s}$ is uniform instead of from the error distribution, and the norm given is $\|\mathbf{e}\|^2$ instead of $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2$. The same proof shows that this version is also as hard as standard LWE (again up to a polynomial factor in the advantage).

We remark that while this problem is reminiscent of Hint-MLWE [KLSS23], Known-Norm LWE is not a special case of Hint-MLWE: here the leakage $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2$ is quadratic in the secrets, whereas the hints in Hint-MLWE are linear.

3.3 Fixed-Matrix Shifted LWE

Definition 11 (Fixed-Matrix Shifted LWE)**.** Let $n \in \mathbb{Z}$; $q \in poly(n)$; $\gamma \in \mathbb{R}$; and let Ψ be an arbitrary distribution over $\mathbb{R}^n$. The Fixed-Matrix Shifted LWE problem is as follows:

Fix a public random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times n}$. Given $\mathbf{A}$, and given polynomially many samples all from either

$$D_{real} = \{\mathbf{c} \leftarrow \Psi;\ \mathbf{d} \leftarrow \Psi;\ \mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\mathbf{c},\gamma};\ \mathbf{f} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\mathbf{d},\gamma};\ :\ (\mathbf{rA}+\mathbf{f},\mathbf{c},\mathbf{d})\}$$

or

$$D_{random} = \left\{\mathbf{c} \leftarrow \Psi;\ \mathbf{d} \leftarrow \Psi;\ \mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n;\ :\ (\mathbf{a},\mathbf{c},\mathbf{d})\right\},$$

determine whether the samples were from D_{real} or D_{random}.

We remark that if the $\mathbf{r}$ and $\mathbf{f}$ vectors came from zero-centered discrete Gaussians, this would be Matrix LWE: the problem of distinguishing $(\mathbf{A}, \mathbf{SA}+\mathbf{E})$ from uniform. Matrix LWE is known to be as hard as standard LWE by a hybrid argument[LP11]. If the shifts $\mathbf{c}$ and $\mathbf{d}$ were integer vectors, we could add or subtract $\mathbf{cA}+\mathbf{d}$ ourselves to shift the centers. The only subtlety is dealing with the fractional parts of the $\mathbf{c}$ and $\mathbf{d}$, which we can do by adding small noise with the appropriate mean.

Lemma 10. *If Matrix LWE is hard in dimension n with the secret and noise taken from discrete Gaussians of parameter σ, then Fixed-Matrix Shifted LWE is hard with noise parameter $\sqrt{\sigma^2 + \eta_\epsilon(\mathbb{Z}^n)^2}$.*

Proof. We are given a Matrix LWE instance $(\mathbf{A}, \mathbf{B} = \mathbf{AS}+\mathbf{E}) \in \mathbb{Z}_q^{n\times n} \times \mathbb{Z}_q^{n\times m}$, where columns of $\mathbf{S}$ and $\mathbf{E}$ come from $\mathcal{D}_{\mathbb{Z}^n,0,\sigma}$. Sample m pairs of vectors $(\mathbf{c_i}, \mathbf{d_i})$ from Ψ. Then, for $i \in [m]$, sample $\bar{\mathbf{s_i}} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\mathbf{c_i},\eta_\epsilon(\mathbb{Z}^n)}$ and $\bar{\mathbf{e_i}} \leftarrow \mathcal{D}_{\mathbb{Z}^n,\mathbf{d_i},\eta_\epsilon(\mathbb{Z}^n)}$.

Concatenate the $\bar{\mathbf{s_i}}$ into an $n \times m$ matrix $\bar{\mathbf{S}}$ and the $\bar{\mathbf{e_i}}$ into $\bar{\mathbf{E}}$. Now

$$(\mathbf{A}, \mathbf{B}+\mathbf{A}\bar{\mathbf{S}}+\bar{\mathbf{E}}, \{(\mathbf{c_i},\mathbf{d_i})\}) = (\mathbf{A}, \mathbf{A}(\mathbf{S}+\bar{\mathbf{S}}) + (\mathbf{E}+\bar{\mathbf{E}}), \{(\mathbf{c_i},\mathbf{d_i})\})$$

which is a Fixed-Matrix Shifted LWE instance with $\gamma = \sqrt{\sigma^2 + \eta_\epsilon(\mathbb{Z}^n)^2}$. If $\mathbf{B}$ is instead uniform, then the result of this transformation is still uniform. Thus Fixed-Matrix Shifted LWE with noise parameter $\gamma = \sqrt{\sigma^2 + \eta_\epsilon(\mathbb{Z}^n)^2}$ is as hard as Matrix LWE with noise parameter σ. □

4 Threshold PKE from PKE

We now show how to transform any PKE scheme satisfying certain properties into a Threshold PKE scheme where the amount of "smudging noise" can be extremely small and the modulus can be polynomially large.

Theorem 1. *Let* PKE *be a public key encryption scheme with LWE-like ciphertexts and public error width, and assume the error width is at least $\sqrt{2}\sigma_{dLWE}$. Assume further (for correctness) that when the error width is σ_{ct}, decryption is possible with noise slightly wider than σ_{ct}; in particular, to build T-out-of-T threshold PKE, we will have error width $\sqrt{\sigma_{ct}^2 + 2T\sigma_{dLWE}^2}$.*

Then the following construction of (T, T)-threshold PKE satisfies simulation security:

Key generation: $(\mathsf{s}, \mathsf{pk}) \leftarrow \mathsf{PKE.KeyGen}$ *as in the underlying (non-threshold) PKE scheme. Each party's secret key $\mathbf{s}_i$ is an additive share of $\mathbf{s}$: $\mathbf{s} = \sum_i \mathbf{s_i}$.*

Encryption: *Same as the underlying PKE scheme.*

Partial Decryption: *To decrypt $(\mathbf{a}, b)$ using keyshare $\mathbf{s_i}$: Sample $\tilde{e} \leftarrow \mathcal{N}_{\sigma_{sm}}$, where $\sigma_{sm} = \sqrt{2}\sigma_{dLWE}$. Output $\langle \mathbf{a}, \mathbf{s_i}\rangle + \tilde{e}$. If given the same input $\mathbf{a}$ multiple times, give the same output every time (e.g., by keeping a table of previous inputs and outputs, or by keeping some secret PRF key k and using $\mathrm{PRF}_k(\mathbf{a})$ as a PRG seed for sampling $\tilde{e}$.)*

Reconstruction: *Add the partial decryptions* $\{\langle \mathbf{a}, \mathbf{s_i}\rangle + \tilde{e}_i\}_{i\in[T]}$ *to recover* $\langle \mathbf{a}, \mathbf{s}\rangle + \sum_i \tilde{e}_i$. *Subtract this from the* b *component of the ciphertext to recover*

$$b - \langle \mathbf{a}, \mathbf{s}\rangle - \sum_i \tilde{e}_i = msg + e_{ct} - \sum_i \tilde{e}_i$$

and then error-correct to recover the underlying message.

Proof. The underlying PKE scheme PKE by assumption has LWE-like ciphertexts and thus is already IND-CPA secure; since giving the adversary $T-1$ shares of $\mathbf{s}$ information theoretically reveals nothing about $\mathbf{s}$, we immediately get (weak) Threshold IND-CPA security for our construction. What remains to be shown is that partial decryption queries can be simulated without knowing the full secret key, and that the adversary's view with a real Dec oracle is computationally indistinguishable from the view with the simulated Dec oracle. Let $\mathbf{s}^{hon}$ be the honest party's secret keyshare, and let $\mathbf{s}^{mal}$ be the sum of the adversary's keyshares, such that $\mathbf{s}^{hon} + \mathbf{s}^{mal} = \mathbf{s}$.

We simulate decryption to a given message as follows:

$$\mathsf{Sim}_{\mathsf{pk},\mathsf{sk}^{mal}}((\mathbf{a}, b), \mathsf{msg}) = b - \langle \mathbf{a}, \mathbf{s}^{mal}\rangle - \mathsf{msg} + c$$

where $c \leftarrow \mathcal{N}_{\sqrt{\sigma_{ct}^2+\sigma_{sm}^2}}$. However, Sim will repeat the same output if queried repeatedly on the same input ciphertext, like the actual Dec. (Without loss of generality we will assume each ciphertext is unique and is decrypted exactly once.)

The adversary's view in the real world is

$$\begin{aligned}&\left\{\mathsf{s}^{mal}, \mathsf{pk}, \{\mathsf{msg}_i, \mathsf{ctxt}_i, \mathsf{Dec}(\mathsf{ctxt}_i)\}_i\right\}\\ =&\left\{\mathsf{s}^{mal}, \mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a_i}, \langle \mathbf{a_i}, \mathbf{s}\rangle + e_i + \mathsf{msg}_i, \langle \mathbf{a_i}, \mathbf{s}^{hon}\rangle + \tilde{e}_i\}_i\right\}\end{aligned}$$

Subtracting the known values $\{\mathsf{msg}_i + \langle \mathbf{a_i}, \mathbf{s}^{mal}\rangle\}$ from each ciphertext (but not from the partial decryptions), we can equivalently write the view as

$$= \left\{\mathsf{s}^{mal}, \mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a_i}, \langle \mathbf{a_i}, \mathbf{s}^{hon}\rangle + e_i, \langle \mathbf{a_i}, \mathbf{s}^{hon}\rangle + \tilde{e}_i\}_i\right\}$$

The $\{\mathbf{a_i}\}$ are indistinguishable from i.i.d. uniform, and the $\{e_i\}$ are indistinguishable from i.i.d. samples from $\mathcal{N}_{\sigma_{ct}}$. (This is true even in the presence of pk and the partial decryptions: since ciphertexts are LWE-like, it would be true even if $\mathbf{s}$ were known, which would be enough for the adversary to compute Dec on their own.) The $\{\tilde{e}_i\}$ are distributed as $\mathcal{N}_{\sigma_{sm}}$.

So $\{\mathbf{a_i}, \langle \mathbf{a_i}, \mathbf{s}^{hon}\rangle + e_i, \langle \mathbf{a_i}, \mathbf{s}^{hon}\rangle + \tilde{e}_i\}_i$ follows the Reused-A LWE distribution, with $\sigma_1 = \sigma_{ct}$ and $\sigma_2 = \sigma_{sm}$. We apply Lemma 8: both σ_{ct} and σ_{sm} are at least $\sqrt{2}\sigma_{dLWE}$, so $\frac{1}{\sigma_{sm}^{-2}+\sigma_{ct}^{-2}} \geq \sigma_{dLWE}^2$, large enough that decision LWE is hard. The conditions for Lemma 8 are thus satisfied, and distribution of this part of the view is computationally indistinguishable from

$$\{\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n;\ b \xleftarrow{\$} \mathbb{R}_q^m;\ c \leftarrow \mathcal{N}_{\sqrt{\sigma_{ct}^2+\sigma_{sm}^2}}\ :\ (\mathbf{a}, b, b+c)\}.$$

Recall that we earlier subtracted $\langle \mathbf{a}, \mathbf{s}^{mal}\rangle + \mathsf{msg}$ from the b component of each ciphertext. We add it back now. $b' = b + \langle \mathbf{a}, \mathbf{s}^{mal}\rangle + \mathsf{msg}$ remains uniform.

$$\{\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n;\ b' \xleftarrow{\$} \mathbb{R}_q^m;\ c \leftarrow \mathcal{N}_{\sqrt{\sigma_{ct}^2+\sigma_{sm}^2}}\ :\ (\mathbf{a}, b', b' - \langle \mathbf{a}, \mathbf{s}^{mal}\rangle - \mathsf{msg} + c)\}.$$

We have shown that the adversary's view in the real world is indistinguishable from

$$V_{real} = \left\{\mathsf{s}^{mal}, \mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a_i}, b_i, b_i - \langle \mathbf{a_i}, \mathbf{s}^{mal}\rangle - \mathsf{msg} + c_i\}_i\right\}$$

where each b_i is uniform. The simulated view V_{sim} is the same except each $b_i = \langle \mathbf{a_i}, \mathbf{s} \rangle + e_i + \mathsf{msg}_i$. These views are seen to be computationally indistinguishable as follows:

Assume $\mathcal{A}$ is a distinguisher between V_{real} and V_{sim}, and define a new $\mathcal{A}'$

$$\mathcal{A}'(\mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a_i}, b_i\}_i) = \mathcal{A}(\mathbf{s}^{mal}, \mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a_i}, b_i, b_i - \langle \mathbf{a_i}, \mathbf{s}^{mal} \rangle - \mathsf{msg}_i + c_i\})$$

where $\mathbf{s}^{mal}$ and c_i are chosen at random. Then $\mathcal{A}'$ distinguishes between

$$X_1 = (\mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a}_i, Unif\}_i) \quad \text{and}$$
$$X_2 = (\mathsf{pk}, \{\mathsf{msg}_i, \mathbf{a}_i, \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i + \mathsf{msg}_i\}_i),$$

which violates the assumption that Decisional LWE with secret $\mathbf{s}$ remains hard given pk. Thus the real world and simulated world are indistinguishable, and the construction is simulation secure. □

We make a few remarks:

- The smudging noise has width $\sigma_{sm} = \sqrt{2}\sigma_{dLWE}$ — only $\sqrt{2}$ times wider than the smallest possible noise under which LWE can be hard. In contrast to previous schemes where smuding noise was superpolynomially larger than ciphertext noise, here the smudging noise is *smaller* than the ciphertext noise. σ_{sm} being so small is especially helpful for supporting a large number of parties, because when partial decryptions from all T parties are added together during reconstruction, the width of the noise will be $\sqrt{\sigma_{ct}^2 + T\sigma_{sm}^2}$.

- The number of decryption queries need not be known in advance, as Reused-A LWE remains hard with arbitrary polynomially many samples.

- The proof requires that σ_{ct} (or more precisely, $\sigma_{ct}^2 + \sigma_{sm}^2$) be a publicly known value, as otherwise the simulator won't know how much noise to add. A natural question is whether a bound on σ_{ct} is sufficient, rather than the actual value, as long as the noise is still Gaussian. This would be a useful property for building threshold homomorphic encryption: the precise noise variance after some operations (like bootstrapping) depends on secret information, but it can be publicly bounded.

  Unfortunately, this is not a proof artifact: simulation security does require the actual value. In the real world, given a ciphertext $(a, b = \langle a, s \rangle + e)$ and partial decryption $c = \langle a, s^{hon} \rangle + \tilde{e}$, the adversary can compute $b - \langle a, s^{mal} \rangle - c = e - \tilde{e}$, which is distributed as a Gaussian with width $\sqrt{\sigma_{ct}^2 + \sigma_{sm}^2}$. The adversary could repeat this for many ciphertexts and measure the variance of the resulting distribution to learn $\sigma_{ct}^2 + \sigma_{sm}^2$.

- The proof of Theorem 1 easily adapts to the Ring LWE setting, where the ciphertext noise should now be a spherical Gaussian of known variance.

- If Dec is called on the same ciphertext twice, it must always give the same output. Otherwise, an adversary could ask for many partial decryptions of the same ciphertext to gather a large number of samples $\{\langle \mathbf{a}, \mathbf{s} \rangle + \tilde{e}_i\}$ and average them, effectively reducing the width of the smudging noise below the $\sqrt{2}\sigma_{dLWE}$ needed for security.

- This Threshold PKE construction is linearly homomorphic (assuming the underlying PKE scheme is) but is not secure as-is as a Threshold Linearly Homomorphic Encryption scheme. In particular, if we allow homomorphic evaluation, the $\mathbf{a}$ vectors in each ciphertext may no longer be independent. For example, $\mathsf{Dec}(ct_1) + \mathsf{Dec}(ct_2) \approx \mathsf{Dec}(ct_1 + ct_2)$ but with different smudging noise. This effectively makes Dec no longer deterministic: the adversary can see many samples $\{\langle \mathbf{a}, \mathbf{s} \rangle + \tilde{e}_i\}$ with the same $\mathbf{a}$ but different $\tilde{e}_i$ and average them.

5 Schemes Satisfying the Conditions of Section 4

5.1 Scheme based on LP11

We now present a PKE scheme (a slight modification of the scheme of [LP11]) that satisfies the conditions needed for the Theorem 1 construction. (For a simpler but less practical alternative PKE scheme based on [Reg09] instead of [LP11], see Appendix A.1.)

Parameters Let q be a polynomially large prime; let n and σ_{dLWE} be chosen such that small-secret LWE in dimension n with modulus q and discrete Gaussian noise of parameter σ_{dLWE} is secure. (We remark that for provable security σ_{dLWE} can be $O(\sqrt{n})$ [Reg09] but in practice $\sigma_{dLWE} \sim O(1)$ is often used.[ACC+18]) Let $\sigma_{pk} = \sigma_{dLWE}$, and $\sigma_e = 2\max(\eta_\epsilon(\mathbb{Z}^{2n}), \sigma_{dLWE})$.

Key generation Sample a square matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times n}$; sample vectors $\mathbf{s}$ and $\mathbf{e_{pk}}$ from $\mathcal{D}_{\mathbb{Z}^n,0,\sigma_{pk}}$. The secret key is $\mathbf{s} \in \mathbb{Z}_q^n$, and the public key is $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e_{pk}}, c = \|\mathbf{s}\|^2 + \|\mathbf{e_{pk}}\|^2) \in \mathbb{Z}_q^{n\times n} \times \mathbb{Z}_q^n \times \mathbb{Z}$.

Encryption On encoded input message msg and public key $(\mathbf{A}, \mathbf{b}, c)$, Enc first samples vectors $\mathbf{r}$ and $\mathbf{f}$ from $\mathcal{D}_{\mathbb{Z}^n,0,\sigma_e}$ and a (continuous) scalar $e' \leftarrow \mathcal{N}_{\sigma_e\sqrt{c}}$. The ciphertext is $(\mathbf{r}^\top\mathbf{A} + \mathbf{f}, \mathbf{r}^\top\mathbf{b} + e' + msg)$.

Theorem 2. *The above scheme has LWE-like ciphertexts (Definition 7) with public error width (Definition 8).*

Proof. For the first and second conditions of Definition 7, the secret key is a vector $\mathbf{s}$, and the public key consists of LWE samples and $\|\mathbf{s}\|^2 + \|\mathbf{e_{pk}}\|^2$, so decisional LWE with secret $\mathbf{s}$ when the public key is known is Known-Norm LWE, which is hard. We now show that Condition 3 of Definition 7 is satisfied, namely, that conditioned on the secret key $\mathbf{s}$ and the public key $(\mathbf{A}, \mathbf{b}, c)$, the output of $\mathsf{Enc}(msg)$ is indistinguishable from

$$\{\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n;\ e_{ct} \leftarrow \mathcal{N}_{\sigma_{ct}};\ (\mathbf{a}, \langle\mathbf{a}, \mathbf{s}\rangle + e_{ct} + msg)\}$$

where $\sigma_{ct} = \sqrt{2c}\sigma_e$. (Observe that σ_{ct} is public.)

We can rewrite the output of Enc as

$$\begin{aligned}
&(\mathbf{r}^\top\mathbf{A} + \mathbf{f},\ \mathbf{r}^\top\mathbf{b} + e' + msg) \\
=\ &(\mathbf{r}^\top\mathbf{A} + \mathbf{f},\ \mathbf{r}^\top\mathbf{A}\mathbf{s} + \langle\mathbf{r}, \mathbf{e_{pk}}\rangle + e' + msg) \\
=\ &(\mathbf{r}^\top\mathbf{A} + \mathbf{f},\ (\mathbf{r}^\top\mathbf{A} + \mathbf{f})\mathbf{s} - \langle\mathbf{f}, \mathbf{s}\rangle + \langle\mathbf{r}, \mathbf{e_{pk}}\rangle + e' + msg).
\end{aligned}$$

We can't yet say that $\mathbf{r}^\top\mathbf{A} + \mathbf{f}$ is uniform by decisional LWE, as $\mathbf{r}$ and $\mathbf{f}$ appear elsewhere.

The vectors $\mathbf{s}$ and $\mathbf{e_{pk}}$ are fixed; $\mathbf{r}$ and $\mathbf{f}$ each come from $\mathcal{D}_{\mathbb{Z}^n,0,\sigma_e}$; e' comes from $\mathcal{N}_{\sigma_e\sqrt{c}}$. We can view the concatenation of $\mathbf{r}$ and $\mathbf{f}$ as a single vector from $\mathcal{D}_{\mathbb{Z}^{2n},0,\sigma_e}$, and the concatenation of $\mathbf{e_{pk}}$ and $\mathbf{s}$ as a vector of norm $\sqrt{\|\mathbf{s}\|^2 + \|\mathbf{e_{pk}}\|^2} = \sqrt{c}$. Our ciphertext distribution is:

$$\left\{[\mathbf{r}|\mathbf{f}] \leftarrow \mathcal{D}_{\mathbb{Z}^{2n},0,\sigma_e};\ e' \leftarrow \mathcal{N}_{\sigma_e\sqrt{c}}\ :\ \Big(\mathbf{rA} + \mathbf{f}, (\mathbf{rA} + \mathbf{f})\mathbf{s} + \Big\langle[\mathbf{r}|\mathbf{f}],\ [\mathbf{e_{pk}} \mid -\mathbf{s}]\Big\rangle + e' + msg\Big)\right\}$$

$$= \left\{[\mathbf{r}|\mathbf{f}] \leftarrow \mathcal{D}_{\mathbb{Z}^{2n},0,\sigma_e};\ \mathbf{e} \leftarrow \mathcal{N}_{\sigma_e}^{2n}\ :\ \Big(\mathbf{rA} + \mathbf{f}, (\mathbf{rA} + \mathbf{f})\mathbf{s} + \Big\langle[\mathbf{r}|\mathbf{f}] + \mathbf{e},\ [\mathbf{e_{pk}} \mid -\mathbf{s}]\Big\rangle + msg\Big)\right\}$$

We now apply Corollary 1 to the joint distribution of $([\mathbf{r}|\mathbf{f}],\ [\mathbf{r}|\mathbf{f}] + \mathbf{e})$, where the value of s in the corollary is σ_e. (σ_e is more than $\sqrt{2}\eta_\epsilon(\mathbb{Z}^{2n})$ so the corollary applies.)

$$\approx_s \left\{\mathbf{y} \leftarrow \mathcal{N}_{\sqrt{2}\sigma_e}^{2n};\ [\mathbf{r}|\mathbf{f}] \leftarrow \mathcal{D}_{\mathbb{Z}^{2n},\mathbf{y}/2,\sigma_e/\sqrt{2}}\ :\ \Big(\mathbf{rA} + \mathbf{f}, (\mathbf{rA} + \mathbf{f})\mathbf{s} + \Big\langle\mathbf{y},\ [\mathbf{e_{pk}} \mid -\mathbf{s}]\Big\rangle + msg\Big)\right\}$$

Now that $\mathbf{r}$ and $\mathbf{f}$ are used only in the expression $\mathbf{rA}+\mathbf{f}$, we can argue by decisional LWE that $\mathbf{rA}+\mathbf{f}$ is computationally indistinguishable from uniform even given $\mathbf{A}$. In particular, this is Fixed-Matrix Shifted LWE, and we apply Lemma 10. (In particular, for our chosen parameters, we have noise width $\sigma_e/\sqrt{2}=\sqrt{2}\max(\eta_\epsilon(\mathbb{Z}^{2n}),\sigma_{dLWE})$, which is larger than the $\sqrt{\sigma_{dLWE}^2+\eta_\epsilon(\mathbb{Z}^n)}$ necessary for Fixed-Matrix Shifted LWE to be hard.)

$$\approx_c \left\{\mathbf{y}\leftarrow\mathcal{N}^{2n}_{\sqrt{2}\sigma_e};\ \mathbf{a}\overset{\$}{\leftarrow}\mathbb{Z}_q^n\ :\ \Big(\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+\Big\langle\mathbf{y},\ [\mathbf{e_{pk}}\mid-\mathbf{s}]\Big\rangle+msg\Big)\right\}$$

Finally, now that $\mathbf{y}$ appears nowhere else, we can replace $\langle\mathbf{y},[\mathbf{e_{pk}}|\mathbf{s}]\rangle$ with a one-dimensional continuous Gaussian:

$$=\left\{e_{ct}\leftarrow\mathcal{N}_{\sqrt{2c}\sigma_e};\ \mathbf{a}\overset{\$}{\leftarrow}\mathbb{Z}_q^n\ :\ (\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+e_{ct}+msg)\right\}$$

In particular, even conditioned on the secret (and public) keys, the output of $\mathsf{Enc}(msg)$ is a fresh LWE sample (plus the message) where the ciphertext noise comes from a continuous Gaussian of width $\sqrt{2c}\sigma_e=\sigma_{ct}$.

□

Theorem 3. *In the above scheme, $\sigma_{ct}\geq\sqrt{2}\sigma_{dLWE}$, and with overwhelming probability, σ_{ct} is polynomially large. In particular, we have correctness for polynomially large modulus q.*

Proof. For the first inequality, $\sigma_{ct}=\sqrt{2c}\sigma_e\geq\sqrt{2}\sigma_{dLWE}$. For the second, using the extremely loose bound of Lemma 2, we have $\sqrt{c}<0.8\sqrt{2n}\sigma_{dLWE}$ with overwhelming probability, so

$$\begin{aligned}\sigma_{ct}&<0.8\sqrt{2n}\sigma_{dLWE}(2\sqrt{2}\max(\eta_\epsilon(\mathbb{Z}^{2n}),\sigma_{dLWE}))\\&=3.2\sqrt{n}\sigma_{dLWE}\max(\eta_\epsilon(\mathbb{Z}^{2n}),\sigma_{dLWE})\end{aligned}$$

which is polynomially large. □

5.2 Scheme based on Regev09

Regev's scheme of [Reg09] can also be modified to have LWE-like ciphertexts with public error width; we defer the details to Appendex A.1.

5.3 RLWE-based PKE

We now sketch a ring-based PKE scheme conjectured to satisfy the properties needed in Theorem 1. Its security depends on the (conjectured but unproven) hardness of a ring analog of the Known-Norm LWE problem that we call the Known-Covariance RLWE problem.

5.3.1 Ring Background

We use the following facts about power-of-2 cyclotomics. Let n be a power of 2, and let $\mathcal{R}=\frac{\mathbb{Z}[x]}{(x^n+1)}$.

- A ring element is a degree $<n$ polynomial a, which can be viewed as an n-dimensional coefficient vector.
- Multiplication by a ring element a is a linear operation, which can be described as multiplication by an n-by-n negacyclic matrix whose top row is the coefficients of a. In a slight abuse of notation, we identify a ring element with its corresponding matrix.

- If A is the matrix for multiplication by $a(x)$, then its transpose $A^\top$ corresponds to multiplication by the ring element $a(x^{-1})$. We write $\overline{a(x)} = a(x^{-1})$.
- Embed $\mathcal{R}$ into $\mathbb{R}^n$. For any ring element a, multiplying a spherical Gaussian by a gives a (non-spherical) Gaussian with covariance matrix $a\overline{a}$. That is, the distribution of $\{e \leftarrow \mathcal{N}_1^n \ : \ ae\}$ is a continuous Gaussian whose covariance matrix is (the matrix corresponding to) $a\overline{a}$.

5.3.2 Known-Covariance RLWE

Definition 12. Let $\mathcal{R} = \frac{\mathbb{Z}[x]}{(x^n+1)}$ be a power-of-2 cyclotomic ring. Let χ be an error distribution whose support is in $\mathcal{R}$. Let q be a (polynomially large) modulus. Let m be polynomially large.

The Known-Covariance RLWE Distribution is defined as

$$\left\{ s \leftarrow \chi;\ \mathbf{e} \leftarrow \chi^m;\ \mathbf{a} \xleftarrow{\$} \mathcal{R}^m \ : \ (\mathbf{a}, s\mathbf{a} + \mathbf{e}, s\overline{s} + \sum_{i=1}^{m} e_i\overline{e_i}) \right\}.$$

The decisional Known-Covariance RLWE assumption is that this distribution is computationally indistinguishable from

$$\left\{ s \leftarrow \chi;\ \mathbf{e} \leftarrow \chi^m;\ \mathbf{a} \xleftarrow{\$} \mathcal{R}^m \ : \ (\mathbf{a}, Unif, s\overline{s} + \sum_{i=1}^{m} e_i\overline{e_i}) \right\}.$$

In Known-Norm LWE we revealed $\|\mathbf{s}\|^2 + \|\mathbf{e}\|^2$, which gives the variance of $\langle [\mathbf{s}|\mathbf{e}], \mathbf{z}\rangle$ when $\mathbf{z} \leftarrow \mathcal{N}_1^{n+m}$.

Analogously in the ring setting, $s\overline{s} + \sum_i e_i\overline{e_i}$ gives the covariance matrix of $\langle [s|\mathbf{e}], \mathbf{z}\rangle$ when $\mathbf{z} \leftarrow (\mathcal{N}_1^n)^{1+m}$. So this is a natural generalization of Known-Norm LWE to the ring setting.

In the plain LWE case Lemma 9 shows Known-Norm LWE is as hard as standard LWE. Unfortunately, that proof does not carry over to the ring setting. It seems plausible that the problem remains hard (for appropriate values of m), but proving or disproving the Known-Covariance RLWE assumption is left open for future work.

5.3.3 Construction

We now sketch our RLWE-based PKE scheme. At a high level, it is similar to the LP11-like scheme from Section 5.1. But directly mapping that construction to the ring setting would give non-spherical ciphertext noise – its covariance matrix will depend on the secret key. So we will add extra noise whose covariance matrix is chosen to cancel this out and make the result spherical, and reveal in the public key the information necessary to compute this covariance matrix. Assuming Known-Covariance RLWE is hard, the scheme remains secure with this leakage.

Let m be such that Known-Covariance RLWE is (conjectured) hard. Let the distribution χ be a discrete Gaussian on $\mathcal{R}$ with width σ_r. Let α be a fixed public parameter.

The secret key is a short ring element s. The public key contains a vector of m RLWE samples $(\mathbf{a}, \mathbf{b} = s\mathbf{a} + \mathbf{e})$. (We remark that $m = 1$ is sufficient if we assume Known-Covariance RLWE is hard when $m = 1$.) The public key also includes a ring element F and an integer λ computed as follows: Let $E = s\overline{s} + \sum_i e_i\overline{e_i}$. Let $\lambda \geq (\sigma_r^2 + \alpha)\lambda_1$, where λ_1 is the largest eigenvalue of E (viewing E as a matrix). Let $F = \lambda\mathbf{I}_n - \sigma_r^2 E$. The public key is $(\mathbf{a}, \mathbf{b}, F, \lambda)$.

Observe that F has the same eigenvectors as E, and all its eigenvalues are non-negative (so F is positive semi-definite). Moreover, observe that the sum of a Gaussian with covariance $\sigma_r^2 E$ and an independent Gaussian with covariance F will be spherical, having covariance $\lambda \mathbf{I}_n$.

Like in the rest of the paper, for ease of presentation ciphertexts will use continuous rather than discrete Gaussian noise.[5]

Encryption is as follows: sample $\mathbf{r} \leftarrow \chi^m$ and $r_0 \leftarrow \chi$. Sample f from an n-dimensional non-spherical Gaussian with covariance matrix F.[6] Output ciphertext $(a' = \langle \mathbf{r}, \mathbf{a} \rangle - r_0, b' = \langle \mathbf{r}, \mathbf{b} \rangle + f + msg)$.

Claim. *In the above scheme, conditioned on the secret key, the distribution of a fresh ciphertext is indistinguishable from $(a', a's + e' + msg)$ where a' is a uniform $\mathcal{R}$ element and e' is an (independent) spherical Gaussian of variance λ.*

Proof Sketch.

$$\begin{aligned}
b' &= \langle \mathbf{r}, \mathbf{b} \rangle + f + msg \\
&= s\langle \mathbf{r}, \mathbf{a} \rangle + \langle \mathbf{r}, \mathbf{e} \rangle + f + msg \\
&= sa' + sr_0 + \langle \mathbf{r}, \mathbf{e} \rangle + f + msg \\
e' &= b' - a's - msg = \langle \mathbf{r}, \mathbf{e} \rangle + sr_0 + f \\
&= \langle [\mathbf{r}|r_0],\ [\mathbf{e}|s] \rangle + f
\end{aligned}$$

Since $\langle [\mathbf{r}|r_0], [\mathbf{e}|s] \rangle$ has covariance matrix $\sigma_r^2 E$, and f is independent with covariance matrix F, the resulting ciphertext noise e' has covariance matrix $\sigma_r^2 E + F = \lambda I_n$, so is a spherical Gaussian of variance λ (all over the randomness of $\mathbf{r}$, r_0, and f).

On its own, $a' = \langle \mathbf{r}, \mathbf{a} \rangle - r_0$ would be computationally indistinguishable from uniform even given $\mathbf{a}$ under Decisional RLWE.[7] However, this is insufficient, because it might be noticeably correlated with e'; we need to analyze the joint distribution of (a', e').

We can write $e = e_1 + e_2$ where e_1 has covariance αE and e_2 has covariance $F - \alpha E$.

$$(a', e') = (\langle \mathbf{r}, \mathbf{a} \rangle - r_0, \langle [r_0|\mathbf{r}],\ [s|\mathbf{e}] \rangle + e_1 + e_2)$$

e_1 has covariance αE, and $\langle \mathbf{r}, \mathbf{e} \rangle$ has covariance $\sigma_r^2 E$; we can combine:

$$= (\langle \mathbf{r}, \mathbf{a} \rangle - r_0, \langle [r_0|\mathbf{r}] + \mathbf{y},\ \mathbf{e} \rangle + e_2)$$

where $\mathbf{y}$ is a spherical Gaussian with variance α.

The rest is similar to the proof of Theorem 2 for the plain case. We apply Lemma 4 to argue that the joint distribution of $([\mathbf{r}|r_0], [\mathbf{r}|r_0] + \mathbf{y})$ is the same as $(\mathbf{r}', \mathbf{y}')$ where $\mathbf{y}'$ is a zero-centered Gaussian and $\mathbf{r}'$ is a discrete Gaussian whose center depends on $\mathbf{y}'$. As long as α and σ_r aren't too small, the distribution of $\mathbf{r}'$ is wide enough to argue that a' is computationally uniform under a ring version of Fixed-Matrix Shifted LWE (Lemma 10), which essentially says that RLWE remains hard if the secrets and error come from discrete Gaussians that aren't centered at zero. This makes a' and e' computationally indistinguishable from independent samples, completing the proof. □

6 Example concrete parameters

We now present some example concrete parameters for the threshold version of the LP11-like scheme from Section 5.1, showing that simulation-secure Threshold PKE is possible with parameters roughly similar to those of FrodoKEM[NAB+20].

[5] The public key still uses discrete noise, like in the rest of the paper, and it does not need to be Gaussian.

[6] For example, sample from $\mathcal{N}_\sigma^n$ and multiply by the matrix square root $\sqrt{\mathbf{F}}$.

[7] If $m > 1$ we could instead assume Module LWE, but Ring LWE suffices.

Frodo-640 targets NIST Level 1, “matching or exceeding the brute-force security of AES-128.” It uses LWE secret dimension $n = 640$ and ciphertext modulus $q = 32768$. The plaintext modulus is 4 — Frodo encodes a 128-bit message as an 8-by-8 matrix of 2-bit values.

For our comparison we will use plaintext modulus 4 as well. Our construction from Section 5.1 encrypts individual scalars rather than 8-by-8 matrices; adapting the construction to matrices we believe should be straightforward but we leave it for future work. Our example parameters are not highly optimized and are for rough comparison only.

First we set n, q, and σ_{dLWE}. For the 128-bit security level, we set $n = 640$ (to match Frodo-640) and $q = 65537$ (twice the modulus of Frodo-640) and $\sigma_{dLWE} = 5$ (giving standard deviation ≈ 1.99, somewhat smaller than Frodo-640’s standard deviation of 2.8, although Frodo’s noise is not exactly Gaussian). According to Albrecht, Player, and Scott’s Lattice Estimator[APS15][8], the estimated complexity of attacking small-secret LWE with these parameters is roughly $2^{141.4}$ operations. Applying Lemma 5 with $c = 0.509$, we get that $\sqrt{\|\mathbf{s}\|^2 + \|\mathbf{e_{pk}}\|^2}$ is less than 91.053 with probability at least $1-2^{-128}$. In particular, if we assume we lose $2\lg(91.053) \approx 13.02$ bits of security in the reduction from Known-Norm LWE to LWE, then the complexity of attacking the public key is still at least $2^{128.3}$. So $\sigma_{dLWE} = 5$ is sufficient for the 128-bit security level.

By Lemma 1, for $\epsilon = 2^{-128}$, we have $\eta_\epsilon(\mathbb{Z}^{2n}) \leq 5.545$. We set $\sigma_e = 2\max(\eta_\epsilon(\mathbb{Z}^{2n}), \sigma_{dLWE}) \approx 11.09$, and we set $\sigma_{pk} = \sigma_{dLWE} = 5$. With these parameters, and using the fact that $\sqrt{\|\mathbf{s}\|^2 + \|\mathbf{e_{pk}}\|^2} < 91.053$ with overwhelming probability, we have

$$\sigma_{ct} < \sqrt{91.053 \cdot 2}\sigma_e \approx 1428.$$

Now we verify that decryption will be correct (for 2-bit messages) with these parameters. For threshold decryption with T parties, the noise during reconstruction will have width $\sigma_d = \sqrt{\sigma_{ct}^2 + 2T\sigma_{dLWE}^2} < \sqrt{1428^2 + 50T}$. To successfully decrypt 2-bit messages, we need the size of the noise to be less than $65536/8 = 8192$ with high probability.

$$\Pr_{x \leftarrow \mathcal{N}_{\sigma_d}}[|x| > 8192] = \Pr_{x \leftarrow \mathcal{N}_1}[|x| > 8192/\sigma_d].$$

We can compute numerically that this decryption failure probability will be less than 2^{-128} as long as $\sigma_d < 1566$. This means these parameters will support as many as $T = 8263$ parties.

We remark that while Frodo-640 ostensibly targets 2^{128} security, its parameters are chosen to withstand quantum attacks and include a large security margin — the Lattice Estimator estimates 2^{163} operations to break Frodo-640. Still, this rough comparison to a highly optimized scheme like FrodoKEM shows that our technique allows Threshold PKE for large numbers of parties with concrete parameters of practical size.

References

[ACC+18] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. Homomorphic encryption security standard. Technical report, HomomorphicEncryption.org, Toronto, Canada, November 2018.

[AJL+12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval

[8] https://github.com/malb/lattice-estimator, commit fd4a460c

and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 483–501. Springer, 2012. doi:10.1007/978-3-642-29011-4_29.

[APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015. URL: https://doi.org/10.1515/jmc-2015-0016 [cited 2023-10-05], doi:doi:10.1515/jmc-2015-0016.

[Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993. doi:10.1007/BF01445125.

[BD10] Rikke Bendlin and Ivan Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In Daniele Micciancio, editor, *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, volume 5978 of *Lecture Notes in Computer Science*, pages 201–218. Springer, 2010. doi:10.1007/978-3-642-11799-2_13.

[BD20a] Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part II*, volume 12106 of *Lecture Notes in Computer Science*, pages 551–575. Springer, 2020. doi:10.1007/978-3-030-45724-2_19.

[BD20b] Zvika Brakerski and Nico Döttling. Lossiness and entropic hardness for ring-lwe. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part I*, volume 12550 of *Lecture Notes in Computer Science*, pages 1–27. Springer, 2020. doi:10.1007/978-3-030-64375-1_1.

[BGG$^+$18] Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 565–596. Springer, 2018. doi:10.1007/978-3-319-96884-1_19.

[BJKL21] Fabrice Benhamouda, Aayush Jain, Ilan Komargodski, and Huijia Lin. Multiparty reusable non-interactive secure computation from LWE. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 724–753. Springer, 2021. doi:10.1007/978-3-030-77886-6_25.

[BKP13] Rikke Bendlin, Sara Krehbiel, and Chris Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *Applied Cryptography and Network Security - 11th International Conference,* *ACNS 2013, Banff, AB, Canada, June 25-28, 2013. Proceedings*, volume 7954 of *Lecture Notes in Computer Science*, pages 218–236. Springer, 2013. doi:10.1007/978-3-642-38980-1_14.

[BP23] Luís T. A. N. Brandão and René Peralta. NIST first call for multi-party threshold schemes (initial public draft). Technical report, National Institute of Standards and Technology, 2023. URL: https://doi.org/10.6028/NIST.IR.8214C.ipd, doi:doi:10.6028/NIST.IR.8214C.ipd.

[BPMW16] Florian Bourse, Rafaël Del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 62–89. Springer, 2016. doi:10.1007/978-3-662-53008-5_3.

[BS23] Katharina Boudgoust and Peter Scholl. Simple threshold (fully homomorphic) encryption from LWE with polynomial modulus. *IACR Cryptol. ePrint Arch.*, page 16, 2023. URL: https://eprint.iacr.org/2023/016.

[CSS+22] Siddhartha Chowdhury, Sayani Sinha, Animesh Singh, Shubham Mishra, Chandan Chaudhary, Sikhar Patranabis, Pratyay Mukherjee, Ayantika Chatterjee, and Debdeep Mukhopadhyay. Efficient threshold FHE with application to real-time systems. *IACR Cryptol. ePrint Arch.*, page 1625, 2022. URL: https://eprint.iacr.org/2022/1625.

[dCHI+22] Leo de Castro, Carmit Hazay, Yuval Ishai, Vinod Vaikuntanathan, and Muthu Venkitasubramaniam. Asymptotically quasi-optimal cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part I*, volume 13275 of *Lecture Notes in Computer Science*, pages 303–334. Springer, 2022. doi:10.1007/978-3-031-06944-4_11.

[DDE+23] Morten Dahl, Daniel Demmler, Sarah Elkazdadi, Arthur Meyre, Jean-Baptiste Orfila, Dragos Rotaru, Nigel P. Smart, Samuel Tap, and Michael Walter. Noah's ark: Efficient threshold-fhe using noise flooding. *IACR Cryptol. ePrint Arch.*, page 815, 2023. URL: https://eprint.iacr.org/2023/815.

[DLN+21] Julien Devevey, Benoît Libert, Khoa Nguyen, Thomas Peters, and Moti Yung. Non-interactive CCA2-secure threshold cryptosystems: Achieving adaptive security in the standard model without pairings. Cryptology ePrint Archive, Paper 2021/630, 2021. https://eprint.iacr.org/2021/630. URL: https://eprint.iacr.org/2021/630, doi:10.1007/978-3-030-75245-3_24.

[GMPW20] Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part I*, volume 12110 of *Lecture Notes in Computer Science*, pages 623–651. Springer, 2020. doi:10.1007/978-3-030-45374-9_21.

[KLO$^+$19] Michael Kraitsberg, Yehuda Lindell, Valery Osheter, Nigel P. Smart, and Younes Talibi Alaoui. Adding distributed decryption and key generation to a ring-lwe based CCA encryption scheme. In Julian Jang-Jaccard and Fuchun Guo, editors, *Information Security and Privacy - 24th Australasian Conference, ACISP 2019, Christchurch, New Zealand, July 3-5, 2019, Proceedings*, volume 11547 of *Lecture Notes in Computer Science*, pages 192–210. Springer, 2019. doi:10.1007/978-3-030-21548-4_11.

[KLSS23] Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-mlwe. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 549–580. Springer, 2023. doi:10.1007/978-3-031-38554-4_18.

[KM16] Veronika Kuchta and Olivier Markowitch. Identity-based threshold encryption on lattices with application to searchable encryption. In Lynn Batten and Gang Li, editors, *Applications and Techniques in Information Security - 6th International Conference, ATIS 2016, Cairns, QLD, Australia, October 26-28, 2016, Proceedings*, volume 651 of *Communications in Computer and Information Science*, pages 117–129, 2016. doi:10.1007/978-981-10-2741-3_10.

[KY16] Shuichi Katsumata and Shota Yamada. Partitioning via non-linear polynomial functions: More compact ibes from ideal lattices and bilinear maps. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part II*, volume 10032 of *Lecture Notes in Computer Science*, pages 682–712, 2016. doi:10.1007/978-3-662-53890-6_23.

[LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for lwe-based encryption. In Aggelos Kiayias, editor, *Topics in Cryptology - CT-RSA 2011 - The Cryptographers' Track at the RSA Conference 2011, San Francisco, CA, USA, February 14-18, 2011. Proceedings*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011. doi:10.1007/978-3-642-19074-2_21.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010. doi:10.1007/978-3-642-13190-5_1.

[MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer, 2011. doi:10.1007/978-3-642-22792-9_26.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. URL:

https://cseweb.ucsd.edu/~daniele/papers/Gaussian.pdf, doi:10.1137/S0097539705447360.

[NAB+20] Michael Naehrig, Erdem Alkim, Joppe Bos, Léo Ducas, Karen Easterbrook, Brian LaMacchia, Patrick Longa, Ilya Mironov, Valeria Nikolaenko, Christopher Peikert, Ananth Raghunathan, and Douglas Stebila. FrodoKEM. Technical report, National Institute of Standards and Technology, 2020. available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009. doi:10.1145/1568318.1568324.

[Shi22] Sina Shiehian. mrnisc from LWE with polynomial modulus. In Clemente Galdi and Stanislaw Jarecki, editors, *Security and Cryptography for Networks - 13th International Conference, SCN 2022, Amalfi, Italy, September 12-14, 2022, Proceedings*, volume 13409 of *Lecture Notes in Computer Science*, pages 481–493. Springer, 2022. doi:10.1007/978-3-031-14791-3_21.

[SRB13] Kunwar Singh, C. Pandu Rangan, and A. K. Banerjee. Lattice based efficient threshold public key encryption scheme. *J. Wirel. Mob. Networks Ubiquitous Comput. Dependable Appl.*, 4(4):93–107, 2013. doi:10.22667/JOWUA.2013.12.31.093.

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635. Springer, 2009. doi:10.1007/978-3-642-10366-7_36.

A Appendix

A.1 Regev-like PKE

A.1.1 Preliminaries: Min-Entropy and the Leftover Hash Lemma

Definition 13 (Min-entropy)**.** The min-entropy of a discrete distribution P is defined as

$$H_\infty(P) = \log \min_{x \in \mathrm{Supp}(P)} \frac{1}{P(x)}$$

The following lemma shows that the min-entropy of a discrete Gaussian distribution over a lattice is minimized when the Gaussian is centered around the origin. Notice that this statement is different from the well known fact (usually proved using Poisson summation formula) that the Gaussian sum $\rho_{\mathbf{c},s}(L)$ is maximized when $\mathbf{c} = 0$ because both the numerator and denominator of $\mathcal{D}_{L,\mathbf{c},s}(\mathbf{x}) = \rho_{\mathbf{c},s}(\mathbf{x})/\rho_{\mathbf{c},s}(L)$ depend on $\mathbf{c}$.

Lemma 11. *The min-entropy of a discrete Gaussian over a lattice L is lowest when centered at the origin. That is, for all $s > 0$ and $\mathbf{c} \in \mathbb{R}^n$, we have*

$$H_\infty(\mathcal{D}_{L,\mathbf{c},s}) \geq H_\infty(\mathcal{D}_{L,0,s}).$$

For simplicity, we prove the statement only for $L = \mathbb{Z}^n$, as this is all that we need in this paper.

Proof. Since $\mathbb{Z}^n$ has an orthogonal basis, we have $\mathcal{D}_{\mathbb{Z}^n,\mathbf{c},s}(\mathbf{x}) = \prod_{i=1}^n \mathcal{D}_{\mathbb{Z},c_i,s}(x_i)$, and it suffices to prove the lemma for $n = 1$. Without loss of generality assume $c \in [-\frac{1}{2}, \frac{1}{2}]$, so that $\mathcal{D}_{\mathbb{Z},c,s}(x)$ is maximized at $x = 0$, and the min-entropy $H_\infty(\mathcal{D}_{\mathbb{Z},c,s})$ is $\log(1/\mathcal{D}_{\mathbb{Z},c,s}(0))$. We wish to show that (for fixed s) $1/\mathcal{D}_{\mathbb{Z},c,s}(0)$ is minimized at $c = 0$.

$$\begin{aligned}\frac{1}{\mathcal{D}_{\mathbb{Z},c,s}(0)} = \frac{\rho_{c,s}(\mathbb{Z})}{\rho_{c,s}(0)} &= \frac{\sum_{x\in\mathbb{Z}} e^{-\pi(x-c)^2/s^2}}{e^{-\pi c^2/s^2}} \\ &= \sum_{x\in\mathbb{Z}} e^{(-\pi/s^2)((x-c)^2-c^2)} \\ &= \sum_{x\in\mathbb{Z}} e^{(-\pi/s^2)(x^2-2xc)}\end{aligned}$$

Merging the summands for $+x$ and $-x$,

$$\begin{aligned}&= 1 + \sum_{x=1}^{\infty}(e^{(-\pi/s^2)(x^2-2xc)} + e^{(-\pi/s^2)(x^2+2xc)}) \\ &= 1 + \sum_{x=1}^{\infty} e^{(-\pi/s^2)x^2}\left(e^{2\pi xc/s^2} + e^{-2\pi xc/s^2}\right) \\ &= 1 + \sum_{x=1}^{\infty} e^{(-\pi/s^2)x^2}\left(y + \frac{1}{y}\right)\end{aligned}$$

where $y = e^{2\pi xc/s^2}$. By symmetry/convexity, this quantity is minimized when $y = 1$, i.e., $c = 0$. Thus $1/\mathcal{D}_{\mathbb{Z},c,s}(0)$ is minimized at $c = 0$, and so is the min-entropy of $\mathcal{D}_{\mathbb{Z},c,s}$. □

The proof can be generalized to arbitrary lattices as follows: without loss of generality, rotate the lattice so that the optimal $\mathbf{c}$ is on the line through the origin and $(1, 0, \ldots, 0)$, so that we can treat c as a scalar. Then take the derivative of $1/\mathcal{D}_{L,c,s}(0)$ with respect to c to show that it is minimized at $c = 0$. We omit the details since $L = \mathbb{Z}^n$ suffices for this paper.

Lemma 12. *Let $s \geq 3$, and let $\mathbf{c} \in \mathbb{R}^n$ be any vector. Then $\mathcal{D}_{\mathbb{Z}^n,\mathbf{c},s}$ has at least n bits of min-entropy.*

Proof. By Lemma 11 we can assume $\mathbf{c} = 0$. In one dimension we can compute numerically that $\rho_{0,3}(0)/\rho_{0,3}(\mathbb{Z}) < 1/2$. So in n dimensions we have $\mathcal{D}_{\mathbb{Z}^n,\mathbf{0},3}(0) < 1/2^n$, and so $\mathcal{D}_{\mathbb{Z}^n,\mathbf{0},3}$ has more than n bits of min-entropy. □

We will use the following formulation of the Leftover Hash Lemma:

Lemma 13 ([BD20a, Lemma 2.1])**.** *Let q be prime and let m, n be integers. Let $\mathbf{r}$ be a random variable defined on $\mathbb{Z}_q^m$ and let $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$ be chosen uniformly at random. Then*

$$\Delta((\mathbf{A}, \mathbf{r}^T\mathbf{A}), (\mathbf{A}, \mathrm{Unif}(\mathbb{Z}_q^n))) \leq \sqrt{q^n \cdot 2^{-H_\infty(\mathbf{r})}}$$

A.1.2 PKE Construction

The following PKE scheme, similar to the scheme of [Reg09], satisfies the properties needed for Theorem 1, and so can be transformed into a simulation secure threshold PKE scheme. Like in the original scheme, here the $\mathbf{a}$ vector of each ciphertext is statistically close to uniform by the Leftover Hash Lemma. We add some additional noise to the b component of the ciphertext to make the ciphertext noise distribution a continuous Gaussian. We also reveal the ℓ_2 norm of the error vector of the public key, so that the width of the ciphertext noise can be publicly known.

Parameters: Let λ be the security parameter. Let q be a polynomially large prime. Let n and σ_{dLWE} be such that LWE in dimension n with Gaussian noise of width σ_{dLWE} is secure. (In theory, σ_{dLWE} is $O(\sqrt{n})$; in practice it is $O(1)$.) Let $\sigma_{pk} = \sigma_{dLWE}$, and $\sigma_r = \sigma_e = \sqrt{2} \cdot \max(3, \eta_\epsilon(\mathbb{Z}^n)) = \tilde{O}(1)$. Let $m = n \log q + \lambda$.

KeyGen: Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}, \mathbf{e_{pk}} \leftarrow \mathcal{D}^m_{\mathbb{Z}^m,0,\sigma_{pk}}$.

The public key is $(\mathbf{A}, \mathbf{b_{pk}} = \mathbf{As} + \mathbf{e_{pk}}, \|\mathbf{e_{pk}}\|)$.

$\mathsf{Enc}_{\mathsf{pk}}(msg)$: Sample $\mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}^m,0,\sigma_r}, e' \leftarrow \mathcal{N}_{\sigma_e \cdot \|e_{pk}\|}$. Output $(\mathbf{r}^\top \mathbf{A}, \mathbf{r}^\top \mathbf{b_{pk}} + e' + msg)$

We now verify the conditions of Theorem 1, namely that the scheme has LWE-like ciphertexts (Definition 7) and public error width (Definition 8).

Theorem 4. *The above construction has LWE-like ciphertexts with public error width.*

Proof. For Condition 1 of Definition 7, the secret key is indeed a vector $\mathbf{s}$. For Condition 2, the public key consists of LWE samples and $\|\mathbf{e_{pk}}\|$, so Decisional LWE when the public key is known is just Known-Norm LWE with uniform secrets, which is as hard as standard LWE (up to a polynomial factor in the advantage). For Condition 3, we will now show that the output distribution of Enc is statistically close to (the message plus) a fresh LWE sample with continuous Gaussian noise of width $\sigma_{ct} = \sigma_e\sqrt{2}\|\mathbf{e_{pk}}\|$. (Note that this width can be computed from the public key, so the scheme has public error width.)

$$\begin{aligned}
\{\mathsf{Enc}(msg)\} &= \{\mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}^m,0,\sigma_r};\ e' \leftarrow \mathcal{N}_{\sigma_e\cdot\|e_{pk}\|} \ :\ (\mathbf{r}^T\mathbf{A}, \mathbf{r}^T(\mathbf{As} + \mathbf{e_{pk}}) + e' + msg)\} \\
&= \{\mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}^m,0,\sigma_r};\ \mathbf{e} \leftarrow \mathcal{N}^m_{0,\sigma_e} \ :\ (\mathbf{r}^T\mathbf{A}, \mathbf{r}^T\mathbf{As} + \langle \mathbf{r}, \mathbf{e_{pk}}\rangle + \langle \mathbf{e}, \mathbf{e_{pk}}\rangle + msg)\} \\
&= \{\mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}^m,0,\sigma_r};\ \mathbf{e} \leftarrow \mathcal{N}^m_{0,\sigma_e} \ :\ (\mathbf{r}^T\mathbf{A}, \mathbf{r}^T\mathbf{As} + \langle \mathbf{r} + \mathbf{e}, \mathbf{e_{pk}}\rangle + msg)\}
\end{aligned}$$

We apply Corollary 1:

$$\approx_s \{\mathbf{t} \leftarrow \mathcal{N}^m_{0,\sqrt{2}\sigma_e};\ \mathbf{r}' \leftarrow \mathcal{D}_{\mathbb{Z}^m,\mathbf{t}/2,\sigma_e/\sqrt{2}} \ :\ (\mathbf{r}'^T\mathbf{A}, \mathbf{r}'^T\mathbf{As} + \langle \mathbf{t}, \mathbf{e_{pk}}\rangle + msg)\}$$

By Lemma 12, $\mathbf{r}'$ has at least m bits of min-entropy; by the Leftover Hash Lemma (Lemma 13) the statistical distance between $(\mathbf{A}, \mathbf{r}'^T\mathbf{A})$ and $(\mathbf{A}, \mathrm{Unif}(\mathbb{Z}_q^n))$ is no more than $2^{-\lambda}$.

$$\begin{aligned}
\{\mathsf{Enc}(msg)\} &\approx_s \{\mathbf{a}' \xleftarrow{\$} \mathbb{Z}_q^n;\ \mathbf{t} \leftarrow \mathcal{N}^m_{0,\sqrt{2}\sigma_e} \ :\ (\mathbf{a}', \langle \mathbf{a}', \mathbf{s}\rangle + \langle \mathbf{t}, \mathbf{e_{pk}}\rangle + msg)\} \\
&= \{\mathbf{a}' \xleftarrow{\$} \mathbb{Z}_q^n;\ \tilde{e} \leftarrow \mathcal{N}_{0,\sqrt{2}\sigma_e\cdot\|\mathbf{e_{pk}}\|} \ :\ (\mathbf{a}', \langle \mathbf{a}', \mathbf{s}\rangle + \tilde{e} + msg)\}.
\end{aligned}$$

So the output of encryption is statistically close to an LWE sample (plus the message) where the noise is a continuous Gaussian with width $\sigma_{ct} = \sigma_e \cdot \sqrt{2}\|\mathbf{e_{pk}}\|$. □

Theorem 5. *For the above scheme, $\sigma_{ct} > \sqrt{2}\sigma_{dLWE}$. Furthermore, for any polynomially large T, polynomially large q is sufficient to allow correct decryption of ciphertexts with noise of width $\sqrt{\sigma_{ct}^2 + 2T\sigma_{dLWE}^2}$.*

Proof. Since $\sigma_e > \sigma_{dLWE}$, and $\sigma_{ct} = \sqrt{2}\sigma_e\|\mathbf{e_{pk}}\|$, we have $\sigma_{ct} \geq \sqrt{2}\sigma_{dLWE}$ as desired.

Since $\mathbf{e}_{\mathsf{pk}}$ comes from $\mathcal{D}_{\mathbb{Z}^m,0,\sigma_{dLWE}}$, by Corollary 2 we have that $\|\mathbf{e_{pk}}\| < 0.8\sqrt{m}\sigma_{dLWE}$ with overwhelming probability. Now $\sigma_{ct} = \sigma_e\sqrt{2}\|\mathbf{e_{pk}}\| \in \tilde{O}(\sqrt{n}\sigma_{dLWE})$. Since σ_{dLWE}, n, and T are all polynomially large, so is $\sigma_{ct}^2 + 2T\sigma_{dLWE}^2$, and a polynomially large modulus suffices to allow decryption to succeed with overwhelming probability. □

A.2 Proof of Lemma 8

First we will need the following lemma:

Lemma 14. *For all* $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ *and* $\mathbf{s} \in \mathbb{Z}_q^n$*, the following two distributions are identical:*

$$\left\{\mathbf{e_1} \leftarrow \mathcal{N}_{\sigma_1}^m;\ \mathbf{e_2} \leftarrow \mathcal{N}_{\sigma_2}^m\ :\ (\mathbf{A}, \mathbf{As}+\mathbf{e_1}, \mathbf{As}+\mathbf{e_2})\right\}$$

and

$$\left\{\mathbf{e} \leftarrow \mathcal{N}_{\sigma_b}^m;\ \mathbf{e}' \leftarrow \mathcal{N}_1^m;\ \mathbf{b} \leftarrow \mathbf{As}+\mathbf{e}\ :\ (\mathbf{A}, \mathbf{b}+\sigma_3\mathbf{e}', \mathbf{b}-\sigma_4\mathbf{e}')\right\}$$

where

$$\sigma_b = \frac{\sigma_1\sigma_2}{\sqrt{\sigma_1^2+\sigma_2^2}} = \frac{1}{\sqrt{\sigma_1^{-2}+\sigma_2^{-2}}}$$
$$\sigma_3 = \frac{\sigma_1^2}{\sqrt{\sigma_1^2+\sigma_2^2}}$$
$$\sigma_4 = \frac{\sigma_2^2}{\sqrt{\sigma_1^2+\sigma_2^2}}$$

Proof. It suffices to show that the joint distribution D_1 of (e_1, e_2) is identical to the joint distribution D_2 of $(e+\sigma_3 e', e-\sigma_4 e')$. Since these are all linear transformations of continuous Gaussians, the distributions are completely specified by their covariance matrices.

The covariance matrix of D_1 is

$$\begin{pmatrix}\sigma_1^2 & \\ & \sigma_2^2\end{pmatrix}.$$

For D_2, we start with (e, e') which has covariance matrix

$$\begin{pmatrix}\sigma_b^2 & \\ & 1\end{pmatrix}$$

and then apply the transformation

$$T = \begin{pmatrix}1 & \sigma_3 \\ 1 & -\sigma_4\end{pmatrix}.$$

The resulting covariance matrix of D_2 is

$$T\begin{pmatrix}\sigma_b^2 & \\ & 1\end{pmatrix}T^\top = \begin{pmatrix}\sigma_b^2+\sigma_3^2 & \sigma_b^2-\sigma_3\sigma_4 \\ \sigma_b^2-\sigma_3\sigma_4 & \sigma_b^2+\sigma_4^2\end{pmatrix} = \begin{pmatrix}\sigma_1^2 & 0 \\ 0 & \sigma_2^2\end{pmatrix}$$

which is the same as the covariance matrix of D_1. So the two distributions are identical. □

With this lemma, we show hardness of Reused-**A** LWE.

Proof of Lemma 8. To reduce LWE to Reused-A LWE, we take an LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{As}+\mathbf{e})$ where $\mathbf{e}$ is continuous Gaussian with width $\sigma_b = \frac{1}{\sqrt{\sigma_1^{-2}+\sigma_2^{-2}}}$. We sample $\mathbf{e}' \leftarrow \mathcal{N}_1^m$ ourselves and compute $(\mathbf{A}, \mathbf{b}+\sigma_3\mathbf{e}', \mathbf{b}-\sigma_4\mathbf{e}')$, with σ_3 and σ_4 defined as in Lemma 14. By the lemma, this is distributed exactly as a Reused-**A** LWE sample with noise parameters σ_1 and σ_2, with the same secret as our input LWE instance. This gives a reduction from search LWE to search Reused-**A** LWE.

For the decision version, by Lemma 14 the Reused-A LWE distribution is identical to $(\mathbf{A}, \mathbf{b}+\sigma_3\mathbf{e}', \mathbf{b}-\sigma_4\mathbf{e}')$, where $\mathbf{e}' \leftarrow \mathcal{N}_1^m$ and $(\mathbf{A}, \mathbf{b})$ are an LWE instance with noise width $\sigma_b = \frac{1}{\sqrt{\sigma_1^{-2}+\sigma_2^{-2}}}$. But $(\mathbf{A}, \mathbf{b})$ look uniform by decision LWE. Letting $\mathbf{b}' = \mathbf{b}+\sigma_3\mathbf{e}'$ (which also looks uniform), our distribution is $(\mathbf{A}, \mathbf{b}', \mathbf{b}'-(\sigma_3+\sigma_4)\mathbf{e}')$, and $\sigma_3+\sigma_4 = \sqrt{\sigma_1^2+\sigma_2^2}$, completing the proof. □

A.3 Simulation Security vs. Threshold IND-CPA-D

Lemma 15. *Threshold IND-CPA-D security does not imply simulation security.*

Proof. Let $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Rec})$ be a Threshold IND-CPA-D secure scheme. Modify the scheme as follows:

- Generate a signing keypair for each party, including each signing key as part of the corresponding party's secret key sk_i, and including all the verification keys in the public key pk.
- Modify Dec so that $\mathsf{Dec}_{\mathsf{sk}_i}(\mathsf{ct})$ also outputs a signature of ct under party i's signing key. (Rec is unmodified and simply ignores the signatures.)

Clearly these modifications do not affect the Threshold IND-CPA-D security of the scheme. But real decryption shares will include valid signatures of the ciphertext, which cannot be simulated without breaking the security of the signature scheme, and so this modified scheme is not simulation secure. □

B Extending [MM11]'s search-to-decision reduction to Known-Norm LWE

Micciancio and Mol[MM11] give a sample-preserving search-to-decision reduction for LWE with uniform secrets, assuming the parameters satisfy some conditions. We extend the results of [MM11] in two ways.

First, we observe that their reduction works not only when the LWE error vector $\mathbf{e} \in \mathbb{Z}_q^m$ is sampled from χ^m for some one-dimensional error distribution χ, but also when $\mathbf{e}$ comes from an arbitrary m-dimensional distribution $\mathcal{X}$ over $\mathbb{Z}_q^m$. In particular, we can let $\mathcal{X}$ be χ^m conditioned on having norm d (for any d), giving a search-to-decision reduction for "Uniform-Secret Known-Norm LWE conditioned on the norm being d".

Next we give a search-to-decision reduction for LWE where secret and error vectors $[\mathbf{s}|\mathbf{e}]$ come from an arbitrary $(m+n)$-dimensional distribution $\mathcal{Y}$ over $\mathbb{Z}_q^{m+n}$, provided the distribution $\mathcal{Y}$ is invariant under permutations. In particular, we can let $\mathcal{Y}$ be χ^{m+n} conditioned on having norm d (for any d), giving a search-to-decision reduction for "Small-Secret Known-Norm LWE conditioned on the norm being d".

When the number of possible norms d is polynomially large, we immediately get search-to-decision reductions for both the small-secret and uniform-secret versions of Known Norm LWE.

We recall the following notations and theorems from [MM11]:

Definition 14 (knapsacks)**.** For any group G and input distribution $\mathcal{X}$ over $\mathbb{Z}^m$, the knapsack family $\mathcal{K}(G, \mathcal{X})$ is the function family with input distribution $\mathcal{X}$ and set of functions $f_{\mathbf{g}} : [\mathcal{X}] \to G$ indexed by $\mathbf{g} \in G^m$ and defined as $f_{\mathbf{g}}(\mathbf{x}) = \mathbf{g} \cdot \mathbf{x} \in G$.

Definition 15 ([MM11])**.** For a function family and input distribution $(F, \mathcal{X})$, define

$$\mathcal{F}(F, \mathcal{X}) = \left\{ f \xleftarrow{\$} F,\, x \leftarrow \mathcal{X} \;:\; (f, f(x)) \right\}$$

The core of [MM11] is a search-to-decision reduction for knapsacks:

Lemma 16 ([MM11, Lemma 4.2])**.** *Let G be a finite abelian group. Let p be the smallest prime factor of $|G|$ and $\mathcal{X}$ be such that $[\mathcal{X}] \subseteq [s]^m$ where $s = \mathsf{poly}(n)$ such that $s \leq p$. If $\mathcal{K}(G, \mathcal{X})$ is one-way, then it is also pseudorandom.*

The LWE search-to-decision reduction combines the knapsack search-to-decision reduction of Lemma 16 with a reduction from search LWE to knapsack search, and a reduction from knapsack decision to decision LWE. (Note that [MM11] uses uniform secrets for LWE.)

Lemma 17 ([MM11, Lemma 4.8])**.** *For any n, $m \geq n + \omega(\log n)$, q, and χ, there is a polynomial time reduction from the problem of inverting* $\mathrm{LWE}(n, m, q, \chi)$ *with probability ϵ, to the problem of inverting $\mathcal{K}(\mathbb{Z}_q^{m-n}, \chi^m)$ with probability $\epsilon' = \epsilon + \mathsf{negl}(n)$.*

Lemma 18 ([MM11, Lemma 4.9])**.** *For any n, $m \geq n + \omega(\log n)$, q, and χ, there is a polynomial time reduction from the problem of distinguishing $\mathcal{F}(\mathcal{K}(\mathbb{Z}_q^{m-n}, \chi^m))$ from uniform with advantage ϵ to the problem of distinguishing $\mathcal{F}(LWE(n, m, q, \chi))$ from uniform with advantage $\epsilon' = \epsilon + \mathsf{negl}(n)$.*

By invoking Lemma 16 with $G = \mathbb{Z}_q^{m-n}$ and $\mathcal{X} = \chi^m$, and combining with Lemma 17 and Lemma 18, [MM11] arrives at the following result:

Lemma 19 ([MM11, Proposition 4.10(ii)])**.** *Let $q = \mathsf{poly}(n)$ be prime and let χ be any distribution over $\mathbb{Z}_q$. If there exists an algorithm for decision* $\mathrm{LWE}_{n,m,q,\chi}$ *that has noticeable advantage, then there exists an efficient algorithm that solves search* $\mathrm{LWE}_{n,m,q,\chi}$ *with noticeable success probability.*

We will use the fact that [MM11]'s reductions between LWE and knapsacks map the error vector of the LWE instance to the solution vector of the knapsack instance and vice-versa:

Claim. *The reduction of Lemma 17, on input an LWE instance* $(\mathbf{A}, \mathbf{As} + \mathbf{e})$*, outputs a knapsack instance with solution vector* $\mathbf{e}$*. In particular, for any m-dimensional distribution $\mathcal{X}$ over $\mathbb{Z}_q^m$, Lemma 17 reduces search* $\mathrm{LWE}(n, m, q, \mathcal{X})$ *to search $\mathcal{K}(\mathbb{Z}_q^{m-n}, \mathcal{X})$.*

Proof. On input LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{As} + \mathbf{e})$, the reduction computes a parity-check matrix $\mathbf{H}$ for $\mathbf{A}$ and outputs knapsack instance $(\mathbf{H}, \mathbf{Hb} = \mathbf{HAs} + \mathbf{He} = \mathbf{He})$, which has solution $\mathbf{e}$. We refer to [MM11] for full details. □

Claim. *The reduction of Lemma 18, on input a knapsack instance* $(\mathbf{G}, \mathbf{c} = \mathbf{Ge})$*, outputs an LWE instance with error vector* $\mathbf{e}$*. In particular, for any m-dimensional distribution $\mathcal{X}$ over $\mathbb{Z}_q^m$, Lemma 18 reduces decision $\mathcal{K}(\mathbb{Z}_q^{m-n}, \mathcal{X})$ to decision $LWE(n, m, q, \mathcal{X})$.*

Proof. On input $(\mathbf{G}, \mathbf{c})$, the reduction computes a random matrix $\mathbf{A}$ whose columns generate the nullspace of $\mathbf{G}$, computes an arbitrary vector $\mathbf{r}$ such that $\mathbf{Gr} \equiv \mathbf{c} \pmod{q}$, picks uniform $\mathbf{s}' \overset{\$}{\leftarrow} \mathbb{Z}_q^n$, and outputs $(\mathbf{A}, \mathbf{As}' + \mathbf{r})$. If the input was a knapsack instance $\mathbf{c} = \mathbf{Ge}$, then the output will be $(\mathbf{A}, \mathbf{As} + \mathbf{e})$ for some unknown $\mathbf{s}$, and so the LWE error vector is exactly the knapsack solution $\mathbf{e}$. We again refer to [MM11] for full details. □

Thus when q is a polynomially large prime, Lemma 19 applies to LWE where the error vector is distributed as any $\mathcal{X}$ over $\mathbb{Z}_q^m$, and not just when the error vector is distributed as χ^m for some χ.

Theorem 6. *Let $q = \mathsf{poly}(n)$ be a polynomially large prime, and let $\mathcal{X}$ be an arbitrary distribution over $\mathbb{Z}_q^m$. If there exists an algorithm for decision* $\mathrm{LWE}_{n,m,q,\mathcal{X}}$ *that has noticeable advantage, then there exists an efficient algorithm that solves search* $\mathrm{LWE}_{n,m,q,\mathcal{X}}$ *with noticeable success probability.*

Proof. Combine the reductions of Lemma 17 and Lemma 18 with Lemma 16. □

B.1 Small Secrets

We now give alternatives of Lemma 17 and Lemma 18 that can be applied to LWE with small secrets, rather than uniform secrets. Whereas in the uniform-secret case the knapsack solution vector was the LWE error $\mathbf{e}$ exactly, here the knapsack solution vector corresponds to the LWE secret and error $[\mathbf{s} \mid \mathbf{e}]$, but with the elements of this vector randomly permuted.

We start by defining a generalization of small-secret LWE:

Definition 16. Let $\mathcal{Y}$ be a distribution over $\mathbb{Z}_q^{n+m}$. The $\mathsf{ssLWE}(n, m, q, \mathcal{Y})$ distribution is

$$\left\{ [\mathbf{s} \mid \mathbf{e}] \leftarrow \mathcal{Y};\ \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}\ :\ (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \right\}.$$

The search ssLWE problem is to recover $\mathbf{s}$ (or equivalently, $[\mathbf{s} \mid \mathbf{e}]$) from such a sample. The decision ssLWE problem is to distinguish a sample from the ssLWE distribution from uniform.

The usual small-secret LWE with one-dimensional error distribution χ is obtained by letting $\mathcal{Y} = \chi^{n+m}$.

Definition 17. Let $\mathcal{Y}$ be a distribution over $\mathbb{Z}_q^{n+m}$. We say that $\mathcal{Y}$ is *invariant under permutations* if, for all $\mathbf{v} \in \mathbb{Z}_q^{m+n}$ and all permutation matrices $\mathbf{P} \in \mathbb{Z}_q^{(m+n)\times(m+n)}$, $\mathcal{Y}$ assigns the same probability to $\mathbf{v}$ as to $\mathbf{P}\mathbf{v}$.

Lemma 20. *For any n, m, and q, and for any distribution $\mathcal{Y}$ over $\mathbb{Z}_q^{m+n}$ that is invariant under permutations, there is a polynomial time reduction from the problem of inverting $\mathsf{ssLWE}(n, m, q, \mathcal{Y})$ with probability ϵ, to the problem of inverting $\mathcal{K}(\mathbb{Z}_q^m, \mathcal{Y})$ with probability $\epsilon' = \epsilon + \mathsf{negl}(n)$.*

Furthermore, on input small-secret LWE instance $(A, A\mathbf{s} + \mathbf{e})$, the reduction outputs a knapsack instance with solution vector $\mathbf{t} = \mathbf{P}[\mathbf{s} \mid \mathbf{e}]$ for some $(n + m)$-dimensional permutation matrix $\mathbf{P}$.

Proof. Let the input LWE instance be $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$. Let $\mathbf{B_0} = \mathbf{U}\,[\mathbf{A}|\mathbf{I}_m] \in \mathbb{Z}_q^{m\times(m+n)}$, where $\mathbf{U}$ is a random invertible m-by-m matrix. Let $\mathbf{t}_0 = [\mathbf{s}|\mathbf{e}]$. Observe that $(\mathbf{B_0}, \mathbf{U}\mathbf{b})$ is a knapsack instance with solution vector $\mathbf{t_0}$, because $\mathbf{B_0}\mathbf{t_0} = \mathbf{U}(\mathbf{A}\mathbf{s}+\mathbf{e}) = \mathbf{U}\mathbf{b}$.

Permute the columns of $\mathbf{B_0}$ and the elements of $\mathbf{t_0}$: let $\mathbf{B} = \mathbf{B_0}\mathbf{P}^\top$ and $\mathbf{t} = \mathbf{P}\mathbf{t_0}$, where $\mathbf{P}$ is a random $(m + n)$-by-$(m + n)$ permutation matrix; observe that $\mathbf{B}\mathbf{t} = \mathbf{B_0}\mathbf{t_0}$. The output knapsack instance is $(\mathbf{B}, \mathbf{U}\mathbf{b})$, which has solution vector $\mathbf{t}$, and $\mathbf{t}$ is a permutation of $[\mathbf{s} \mid \mathbf{e}]$ as desired.

Finally, we show the output knapsack instance follows the correct distribution: $\mathbf{B}$ must be (statistically close to) uniform in $\mathbb{Z}_q^{m\times(m+n)}$, and $\mathbf{t}$ must follow $\mathcal{Y}$. $[\mathbf{s} \mid \mathbf{e}] \sim \mathcal{Y}$, and $\mathbf{t}$ is a permutation of $[\mathbf{s} \mid \mathbf{e}]$. Since $\mathcal{Y}$ is invariant under permutation, the distribution of $\mathbf{t}$ is also $\mathcal{Y}$, as desired.

Since $\mathbf{A}$ is uniformly distributed, $\mathbf{B_0}$ is a uniformly distributed m-by-$(m + n)$ matrix conditioned on its last m columns forming an invertible matrix. Then $\mathbf{B}$ is uniformly distributed conditioned on having any set of m columns that form an invertible matrix, i.e., conditioned on being nonsingular. A random matrix in $\mathbb{Z}_q^{m\times(m+n)}$ is singular with probability at most $1/p^{n-1}$ where p is the smallest prime factor of q.[MM11, proof of Lemma 4.8] Thus the distribution of $\mathbf{B}$ is statistically close to the uniform distribution on $\mathbb{Z}_q^{m\times(m+n)}$. □

Lemma 21. *For any n, $m \geq n + \omega(\log n)$, q, and for any distribution $\mathcal{Y}$ over $\mathbb{Z}_q^{m+n}$ that is invariant under permutation, there is a polynomial time reduction from the problem of distinguishing $\mathcal{F}(\mathcal{K}(\mathbb{Z}_q^m, \mathcal{Y}))$ from uniform with advantage ϵ to the problem of distinguishing $\mathcal{F}(\mathsf{ssLWE}(n, m, q, \mathcal{Y}))$ from uniform with advantage $\epsilon' = \epsilon + \mathsf{negl}(n)$.*

Furthermore, on input knapsack instance $(G, \mathbf{c} = G\mathbf{t})$, the reduction outputs an LWE instance $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ where $[\mathbf{s} \mid \mathbf{e}]$ is some permutation of $\mathbf{t}$.

Proof. Let the input be $(\mathbf{G}, \mathbf{c})$, where $\mathbf{G}$ is uniformly distributed in $\mathbb{Z}_q^{m\times(m+n)}$, and $\mathbf{c}$ is either uniform in $\mathbb{Z}_q^{m+n}$ or is $\mathbf{Gt}$ for some $\mathbf{t}$ sampled from $\mathcal{Y}$. We will output a pair $(\mathbf{A}, \mathbf{b})$ where $\mathbf{A}$ is uniform in $\mathbb{Z}_q^{m\times n}$ and $\mathbf{b}$ is either uniform in $\mathbb{Z}_q^m$ or is $\mathbf{As} + \mathbf{e}$ for some $[\mathbf{s}|\mathbf{e}]$ distributed as χ^{m+n}.

As $\mathbf{G}$ is m-by-$(m+n)$, with all but negligible probability, $\mathbf{G}$ has some set of m columns that form an invertible matrix. Let $\mathbf{P}$ be a random permutation matrix such that the last m columns of $\mathbf{G}' = \mathbf{GP}^\top$ form an invertible matrix $\mathbf{U}$.

Let $[\mathbf{A} \mid \mathbf{I}] = \mathbf{U}^{-1}\mathbf{G}'$, noting that if $\mathbf{G}$ is uniformly distributed then so is $\mathbf{A}$. Finally let $\mathbf{b} = \mathbf{U}^{-1}\mathbf{c}$. The reduction outputs $(\mathbf{A}, \mathbf{b})$.

If $\mathbf{c}$ is uniform, then clearly so is $\mathbf{b}$, so the reduction maps uniform to uniform. Now suppose instead $\mathbf{c} = \mathbf{Gt}$, with $\mathbf{t}$ distributed as $\mathcal{Y}$. Let $\mathbf{t}' = \mathbf{Pt}$; $\mathbf{t}'$ also follows $\mathcal{Y}$ because $\mathcal{Y}$ is invariant under permutations. Now

$$\mathbf{b} = \mathbf{U}^{-1}\mathbf{c} = \mathbf{U}^{-1}\mathbf{Gt} = \mathbf{U}^{-1}\mathbf{G}'\mathbf{t}' = [\mathbf{A} \mid \mathbf{I}]\,\mathbf{t}'.$$

Letting $[\mathbf{s} \mid \mathbf{e}] = \mathbf{t}'$, we have $\mathbf{b} = \mathbf{As} + \mathbf{e}$ with $[\mathbf{s} \mid \mathbf{e}] \sim \mathcal{Y}$, so the output is a random instance of $\mathsf{ssLWE}_{n,m,q,\mathcal{Y}}$ with $[\mathbf{s} \mid \mathbf{e}]$ a permutation of $\mathbf{t}$ as desired. □

Theorem 7. *Let $q = \mathsf{poly}(n)$ be a polynomially large prime. Let $\mathcal{Y}$ be an arbitrary distribution over $\mathbb{Z}_q^{m+n}$ that is invariant under permutations. If there exists an algorithm for decision $\mathsf{ssLWE}_{n,m,q,\mathcal{Y}}$ that has noticeable advantage, then there exists an efficient algorithm that solves search $\mathsf{ssLWE}_{n,m,q,\mathcal{Y}}$ with noticeable success probability.*

Proof. Lemma 20 reduces search $\mathsf{ssLWE}_{n,m,q,\mathcal{Y}}$ to search $\mathcal{K}(\mathbb{Z}_q^m, \mathcal{Y})$.

Lemma 16 reduces search $\mathcal{K}(\mathbb{Z}_q^m, \mathcal{Y})$ to decision $\mathcal{K}(\mathbb{Z}_q^m, \mathcal{Y})$.

Lemma 21 reduces decision $\mathcal{K}(\mathbb{Z}_q^m, \mathcal{Y})$ to decision $\mathsf{ssLWE}_{n,m,q,\mathcal{Y}}$. □

B.2 Application to Known-Norm LWE

Finally, we prove a search-to-decision reduction for Known-Norm LWE, assuming q is a polynomially large prime.

For all d, let $\mathcal{X}_d$ be χ^m conditioned on having norm d, and let $\rho(d) = \Pr_{\mathbf{e}\leftarrow\chi^m}[\|\mathbf{e}\| = d]$.

Suppose $\mathcal{D}$ is a distinguisher for uniform-secret known-norm LWE with non-negligible advantage ϵ. Let ϵ_d be $\mathcal{D}$'s advantage conditioned on the norm being d. Then $\epsilon = \sum_d \rho(d)\epsilon_d$. Since q is polynomially large and so there are only polynomially many possibilities for d, there must be some d for which $\rho(d)\epsilon_d$ is non-negligible. Applying Theorem 6 with $\mathcal{X} = \mathcal{X}_d$ for this particular d gives an algorithm for search LWE conditioned on the norm being d that succeeds with non-negligible probability. Since $\rho(d)$ is non-negligible, this gives an algorithm for search LWE with non-negligible success probability overall.

For small-secret Known-Norm LWE, we let $\mathcal{Y}_d$ be χ^{m+n} conditioned on having norm d; as this is permutation invariant, we can apply Theorem 7. The rest of the proof is as in the uniform-secret case.