

Key-Homomorphic Pseudorandom Functions from LWE with Small Modulus

Sam Kim

Stanford University

Abstract. Pseudorandom functions (PRFs) are fundamental objects in cryptography that play a central role in symmetric-key cryptography. Although PRFs can be constructed from one-way functions generically, these black-box constructions are usually inefficient and require deep circuits to evaluate compared to direct PRF constructions that rely on specific algebraic assumptions. From lattices, one can directly construct PRFs from the Learning with Errors (LWE) assumption (or its ring variant) using the result of Banerjee, Peikert, and Rosen (Eurocrypt 2012) and its subsequent works. However, all existing PRFs in this line of work rely on the hardness of the LWE problem where the associated modulus is super-polynomial in the security parameter.

In this work, we provide two new PRF constructions from the LWE problem. In each of these constructions, each focuses on either minimizing the depth of its evaluation circuit or providing key-homomorphism while relying on the hardness of the LWE problem with either a polynomial modulus or nearly polynomial modulus. Along the way, we introduce a new variant of the LWE problem called the Learning with Rounding and Errors (LWRE) problem. We show that for certain settings of parameters, the LWRE problem is as hard as the LWE problem. We then show that the hardness of the LWRE problem naturally induces a pseudorandom synthesizer that can be used to construct a low-depth PRF. The techniques that we introduce to study the LWRE problem can then be used to derive variants of existing key-homomorphic PRFs whose security can be reduced from the hardness of the LWE problem with a much smaller modulus.

1 Introduction

A pseudorandom function (PRF) [32] is a deterministic function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ that satisfies a specific security property: for a randomly sampled key $k \xleftarrow{R} \mathcal{K}$, the output of the function $F(k, \cdot)$ is computationally indistinguishable from those of a truly random function. PRFs are fundamental objects in cryptography that serve as a basis for symmetric cryptography. Even beyond symmetric cryptography, PRFs serve as one of the core building blocks for many advanced cryptographic constructions and protocols.

In theory, PRFs can be constructed from any one-way function via the transformation of [34, 32]. However, the PRFs that are constructed from one-way functions in a blackbox way are generally inefficient. Furthermore, the transformation of [34, 32] is inherently *sequential* and therefore, the resulting PRFs

require deep circuits to evaluate. For practical deployment, this is problematic as symmetric objects like PRFs are often deployed inside designated hardware devices similar to how modern blockciphers, such as AES, are incorporated into many modern processors. For these settings, it is important for PRFs to exhibit *low depth* evaluation circuits that require few computing cycles to evaluate using multiple cores or processing units.

For these reasons, constructing low-depth pseudorandom functions from standard cryptographic assumptions have been a highly active area of research. Starting from the seminal work of Naor and Reingold [45], there have been great progress in constructing low-depth pseudorandom functions from *group-based* assumptions like the Decisional Diffie-Hellman (DDH) assumption [45, 1, 2]. However, constructing low-depth PRFs from standard lattice assumptions such as the Learning with Errors (LWE) assumption [48] has been surprisingly more difficult. Indeed, a low-depth PRF from LWE was constructed by a breakthrough result of Banerjee, Peikert, and Rosen [11], but only after the realization of seemingly more powerful primitives such as (lattice-based) identity-based encryption [31, 3, 4, 27] and fully homomorphic encryption [30, 23, 22].

Key-homomorphic PRFs. Since the work of [11], the study of lattice-based PRFs has become a highly productive area of research. There have been a sequence of results that further improve the constructions of [11] with various trade-offs in the parameters [17, 10, 28, 43, 35]. A long sequence of results also show how to construct PRFs with useful algebraic properties such as key-homomorphic PRFs [17, 10, 25], constrained PRFs [25, 16, 26, 21], and even watermarkable PRFs [36, 26, 47, 37] from LWE.

A special family of PRFs that are particularly useful for practical applications are key-homomorphic PRFs. The concept was first introduced by Naor, Pinkas, and Reingold [44] and it was first formalized as a cryptographic primitive by Boneh et al. [17]. We say that a pseudorandom function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ is *key-homomorphic* if the key-space (\mathcal{K}, \oplus) and the range of the PRF (\mathcal{Y}, \otimes) exhibit group structures such that for any two keys k_1, k_2 and input $x \in \mathcal{X}$, we have $F(k_1 \oplus k_2, x) = F(k_1, x) \otimes F(k_2, x)$. Key-homomorphic PRFs have many useful applications in symmetric cryptography and give rise to distributed PRFs, symmetric-key proxy re-encryption, and updatable encryption. The study of updatable encryption, in particular, have recently gained a considerable amount of traction [17, 29, 40, 38, 19]. Most of these existing proposals for updatable encryption rely on key-homomorphic PRFs or use direct updatable encryption constructions that take advantage of similar algebraic structures.

LWE modulus. Despite significant progress in our understanding of lattice-based PRFs as described above, all existing direct PRF constructions suffer from one caveat: the modulus q , which defines the underlying ring for the PRF, must be set to be super-polynomial in the security parameter. The need for a large modulus q has several disadvantages. The first and immediate disadvantage is *efficiency*. A lattice-based PRF is generally defined with respect to a set of public matrices in $\mathbb{Z}_q^{n \times m}$ and a secret vector in \mathbb{Z}_q^n for some suitable choice of parameters

n and m . A bigger modulus q means that more space is required to store these values and more time is needed to evaluate the PRF.

Another disadvantage of a large modulus q is related to the *quality* of the LWE assumption that is needed to prove security. Generally, PRFs that are defined with a super-polynomial modulus q relies on the hardness of the LWE problem with a super-polynomial noise-to-modulus ratio. This means that the security of the PRF can only be based on the hardness of solving worst-case lattice problems with a super-polynomial approximation factor, which is a significantly stronger assumption than what is required by many other lattice-based cryptographic constructions. In fact, today, we can base seemingly stronger primitives like fully-homomorphic encryption [24, 7] and attribute-based encryption [33] (when restricted to NC^1 computations) on the hardness of approximating worst-case lattice problems with only polynomial approximation factors. This clearly demonstrates that our current understanding of lattice-based PRFs is still quite limited.

Many existing lattice-based PRFs including the original constructions of [11] are, in fact, conjectured to be secure even when they are instantiated with much smaller values of the modulus q . However, their formal proof of security has been elusive for many years. An important question in the study of lattice-based PRFs is whether there exist direct lattice-based PRF constructions that rely on a polynomial modulus q that still exhibit some of the useful features not satisfied by the generic constructions [34, 32] such as low-depth evaluation circuits or key-homomorphism.

1.1 Our Contributions

In this work, we present two PRF constructions from the hardness of the LWE problem with a small modulus q . For our first construction, we focus on minimizing the depth of the evaluation circuit while in our second construction, we focus on constructing a key-homomorphic PRF. In both settings, our goal is to construct lattice-based PRFs that work over small moduli q .

1.1.1 Low-depth PRF

In our first PRF construction, our main focus is on minimizing the size of the modulus q while also minimizing the depth of the PRF evaluation circuit. We provide an overview of the construction in Section 2 and briefly discuss our approach below. The resulting PRF can be instantiated with a range of parameter settings that are determined by a trade-off between the depth of the evaluation circuit and the associated modulus q . We consider two types of parameter settings that provide different levels of security.

- *Theoretical security*: In this setting, we guarantee that any adversary has at most a negligible advantage in breaking the PRF. For this level of security, we can set the parameters of our PRF such that the modulus is polynomial in the security parameter $q = \tilde{O}(\lambda^{2.5})$ and the depth of the evaluation circuit is in NC^2 .

- 2^λ -*security*: In this setting, we guarantee that an adversary’s advantage in breaking the PRF degrades exponentially in the security parameter. For this level of security, we can set the parameters of our PRF such that the depth of the evaluation circuit is $\tilde{O}(\lambda/\log q)$. In particular, setting $q = 2^{\tilde{O}(\sqrt{\lambda})}$, the PRF evaluation can be done in depth $\tilde{O}(\sqrt{\lambda})$. Previously, all lattice-based PRFs either required that the depth of the evaluation circuit is at least $\tilde{O}(\lambda)$ or the modulus q to be at least $2^{\tilde{O}(\lambda)}$.

We provide a comparison of the size of the modulus q and the depth of the evaluation circuit that is needed for our PRF with those of existing LWE-based PRFs in Table 1. We further discuss how to interpret our parameter settings in Section 1.2 and how to concretely instantiate them in Section 4.3.

Synthesizers and LWRE. The main intermediate object that we use to construct our PRF is a pseudorandom synthesizer (Definition 4.5), which was first introduced by Naor and Reingold [45]. They showed that a pseudorandom synthesizer that can be computed by a low-depth circuit can be used to construct a PRF that can also be computed by a low-depth circuit. The work of Banerjee, Peikert, and Rosen [11] first showed that such pseudorandom synthesizers can be constructed from a natural variant of the LWE problem called the Learning with Rounding (LWR) problem. They showed that the hardness of the LWR problem can be reduced from the hardness of the LWE problem when the modulus q is set to be very large.

To construct a pseudorandom synthesizer from a small-modulus LWE problem, we introduce yet another variant of the LWE problem called the Learning with Rounding *and* Errors (LWRE) problem whose hardness naturally induces a pseudorandom synthesizer. The challenger for an LWRE problem chains multiple samples of the LWR and LWE problems together such that the error terms that are involved in each of the LWE samples are derived from the “noiseless” LWR samples. This specific way of chaining LWR and LWE samples together allows us to reduce the hardness of LWRE from the hardness of the LWE problem with a much smaller modulus q . We provide an overview of the LWRE problem and the reduction from LWE in Section 2. We precisely formulate the LWRE problem in Definition 4.1 and provide the formal reduction in in the full version of this work.

The LWRE problem and our synthesizer construction naturally extend to the ring setting as well. In the full version of this work, we formulate the Ring-LWRE problem similarly to how the Ring-LWE and Ring-LWR problems are defined. Then, we show how to construct a pseudorandom synthesizer from the Ring-LWRE problem.

1.1.2 Key-homomorphic PRF

For our second construction, we focus on constructing a key-homomorphic PRF with a small modulus q . Specifically, we provide a key-homomorphic PRF whose security (either theoretical or 2^λ -security) can be based on the hardness of LWE

Construction	Size of Modulus	Evaluation Depth
[11, GGM]	$\lambda^{\Omega(1)}$	$\Omega(\lambda \log \lambda)$
[11, Synthesizer]	$2^{\Omega(\lambda)}$	$\Omega(\log^2 \lambda)$
[11, Direct]	$2^{\Omega(\lambda \log \lambda)}$	$\Omega(\log^2 \lambda)$
[17]	$2^{\Omega(\lambda \log \lambda)}$	$\Omega(\log^2 \lambda)$
[10]	$2^{\Omega(\lambda)}$	$\Omega(\log^2 \lambda)$
[28]	$2^{\Omega(\lambda)}$	$\Omega(\log^{1+o(1)} \lambda)$
[43]	$2^{\Omega(\lambda)}$	$\Omega(\log^2 \lambda)$
[35]	$2^{\Omega(\lambda)}$	$\Omega(\log^{1+o(1)} \lambda)$
This work: Synthesizer-based	$2^{\Omega(\sqrt{\lambda})}$	$\Omega(\sqrt{\lambda} \log \lambda)$
This work: BP-based	$\lambda^{\Omega(1)}$	$\Omega(\lambda^2 \log \lambda)$

Table 1: Comparison of the PRF constructions in this work and the existing PRF constructions based on LWE. For each of the PRF constructions, the table denotes the size of the modulus q that is needed to prove 2^λ -security from LWE and the depth of the evaluation circuit that is needed to evaluate the PRFs.

with a polynomial size q without relying on random oracles. All previous key-homomorphic PRFs from lattices either relied on LWE with a super-polynomial modulus q [17, 10, 25] or random oracles [44, 17]. As in previous LWE-based key-homomorphic PRFs, our construction is “almost” key-homomorphic in that the homomorphism on the PRF keys hold subject to some small rounding error, which does not significantly impact its usefulness to applications.

Our construction relies on the same chaining technique that is used to construct our first PRF. This time, instead of chaining multiple LWE and LWR samples together as was done in our first construction, we chain multiple instances of the existing lattice-based PRFs themselves. For most existing PRF constructions that are based on LWE, their proof of security proceeds in two steps:

1. Define a “noisy” variant of the deterministic PRF function whose security can be based on the hardness of the LWE problem.
2. Show that the deterministic PRF function and its “noisy” variant have the same input-output behavior with overwhelming probability over the randomness used to sample the noise.

Generally, in order to show that the deterministic PRF and its “noisy” variant are statistically indistinguishable in step 2 above, the modulus q has to be set to be super-polynomial in the security parameter.

To reduce the need for a large modulus q in step 2, we chain multiple instances of the deterministic PRF and its noisy variant. Namely, our PRF construction consists of many noisy variants of these PRFs that are chained together such that the noise that is needed to evaluate the noisy PRF in a chain is derived from a PRF in the previous level of the chain. By setting the initial PRF at the

start of the chain to be the original deterministic PRF, the entire evaluation of the chained PRF can be made to be deterministic.

This simple way of chaining multiple instances of deterministic and noisy variants of PRFs allows us to prove the security of the final PRF from the hardness of LWE with a much smaller modulus q . In fact, when we chain multiple instances of a key-homomorphic PRF, the resulting PRF is also key-homomorphic. Instantiating the chain with the Banerjee-Peikert key-homomorphic PRF [10] results in a key-homomorphic PRF that works over a polynomial modulus q . We provide a detailed overview of our technique in Section 2 and provide the formal construction and its security proof in Section 5.

1.2 Discussions

Regarding theoretical and 2^λ -security. The reason why we present our results with respect to both theoretical and 2^λ -security is due to the fact that the generic PRF constructions [32] can already be used to construct a low-depth PRF that provides asymptotically equivalent level of security. Note that using a length-doubling pseudorandom generator, the Goldwasser, Goldreich, and Micali [32] construction can be used to provide a PRF that can be evaluated in depth linear in the input size of the PRF. One way to achieve a poly-logarithmic depth PRF using the GGM construction is to first hash the input using a universal hash function into a domain of size $2^{\omega(\log \lambda)}$ and then apply the PRF on the hash of the message. As the hashed outputs are poly-logarithmic in length, the PRF can be evaluated in poly-logarithmic depth. At the same time, as long as the adversary is bounded to making a polynomial number of evaluation queries on the PRF, a collision on the hash function occurs with negligible probability and therefore, any efficient adversary can have at most negligible advantage in breaking the PRF. As a length-doubling PRG can be constructed from LWE with a polynomial modulus q , this gives an LWE-based PRF with both small evaluation depth and small modulus q . Of course, the actual security of this final PRF is quite poor since the probability that an adversary forces a collision on the hash function is barely negligible.

Therefore, the way to view our low-depth PRF is to consider its parameters when they are set to provide 2^λ -security. In this setting, our PRF provides security under the condition that $d \log q = \tilde{\Omega}(\lambda)$ where d is the depth of the evaluation circuit. When setting $\tilde{\Omega}(\log q) = \sqrt{\lambda}$, the evaluation circuit has depth that scales with $\sqrt{\lambda}$. This means that setting $\lambda = 128$ and ignoring arithmetic and vector operations, our PRF can be evaluated by a circuit with depth ≈ 11 that works over a ≈ 11 -bit modulus. The GGM PRF, on the other hand, requires a circuit with depth at least $\lambda = 128$, which is prohibitive for practical use, while the existing lattice-based PRFs require $7 \approx \log \lambda$ circuit depth, but must operate over at least a 128-bit modulus. We discuss concrete instantiation of our scheme in Section 4.3.

We note that for key-homomorphic PRFs, no construction that works over a polynomial modulus was previously known. Therefore, our second PRF con-

struction can be viewed as the first key-homomorphic PRF that works over a polynomial modulus independent of whether it provides theoretical or 2^λ -security.

On the chaining method and blockciphers. The pseudorandom synthesizers or PRFs in this work consist of many repeated rounds of computation that are chained together in such a way that the output of each round of computation is affected by the output of the previous round. This way of chaining multiple rounds of computation is reminiscent of the structure of many existing blockciphers such as DES or AES, which also consist of many rounds of bit transformations that are chained together. Interestingly, chaining in blockciphers and chaining in our work seem to serve completely opposite roles. In blockcipher design, chaining is generally used to achieve the effect of *diffusion*, which guarantees that a small change to the input to the blockcipher significantly alters its final output. This assures that no correlation can be efficiently detected between the input and the output of the blockcipher. In this work, chaining is used to actually *prevent* diffusion. Namely, chaining guarantees that some small error that is induced by the PRF evaluation does not affect the final output of the PRF. This allows us to switch from the real PRF evaluation to the “noisy” PRF evaluation in our hybrid security argument such that we can embed an LWE challenge to the output of the PRF.

1.3 Other Related Work

PRF cascades. The techniques that are used in this work are conceptually similar to PRF cascading, which is the process of chaining multiple small-domain PRFs to construct large-domain PRFs. The technique was first introduced by Bellare et al. [14] and was further studied by Boneh et al. [18]. PRF cascades serve as a basis for NMAC and HMAC PRFs [13, 12].

LWR for bounded number of samples. There have been a sequence of results that study the hardness of the Learning with Rounding problem when the number of samples are a priori bounded. Alwen et al. [8] first showed that such variant of LWR is as hard as LWE even when the modulus q is set to be polynomial in the security parameter. Bogdanov et al. [15] improved the statistical analysis of this reduction using Rényi divergence. Alperin-Sherif and Apon [6] further improved these previous results such that the reduction from LWE to LWR is sample-preserving.

2 Overview

In this section, we provide a high level overview of the main techniques that we use in this work. For the full details of our main results and proofs, we refer the readers to Sections 4 and 5.

We divide the overview into three parts. In Section 2.1, we provide additional background on existing results on constructing lattice-based PRFs. In Section 2.2, we provide an overview of our synthesizer construction from a new computational

problem called the Learning with Rounding and Errors (LWRE) problem. Then, in Section 2.3, we show how the technique that we use to prove the security of our synthesizer-based PRF can be applied to the parameters for existing lattice-based key-homomorphic PRFs.

2.1 Background on Lattice PRFs via Synthesizers

The main intermediate primitive that we use to construct our first lattice-based PRF is a special family of pseudorandom generators called *pseudorandom synthesizers* [45]. A pseudorandom synthesizer over a domain \mathcal{D} is a two-to-one function $S : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$ such that for any (a priori *unbounded*) polynomial number of inputs $a_1, \dots, a_\ell \stackrel{\text{R}}{\leftarrow} \mathcal{D}$, and $b_1, \dots, b_\ell \stackrel{\text{R}}{\leftarrow} \mathcal{D}$, the set of ℓ^2 elements $\{S(a_i, b_j)\}_{i,j \in [\ell]}$ are computationally indistinguishable from uniformly random elements $\{u_{i,j}\}_{i,j \in [\ell]} \stackrel{\text{R}}{\leftarrow} \mathcal{D}^{\ell^2}$.

A pseudorandom synthesizer $S : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$ induces a PRF F with key space $\mathcal{D}^{2\ell}$, domain $\{0, 1\}^\ell$, and range \mathcal{D} for any positive power-of-two integer ℓ as follows:

- Define a PRF key to consist of 2ℓ uniformly random elements in \mathcal{D} :

$$\begin{pmatrix} s_{1,0} & s_{2,0} & \dots & s_{\ell,0} \\ s_{1,1} & s_{2,1} & \dots & s_{\ell,1} \end{pmatrix}.$$

- To evaluate the PRF on an input $x \in \{0, 1\}^\ell$, compress the subset of the elements $s_{1,x_1}, \dots, s_{\ell,x_\ell}$ into a single element of \mathcal{D} by iteratively applying the synthesizer:

$$S\left(\dots S(S(s_{1,x_1}, s_{2,x_2}), S(s_{3,x_3}, s_{4,x_4})), \dots S(s_{\ell-1,x_{\ell-1}}, s_{\ell,x_\ell}) \dots\right).$$

The pseudorandomness of the output of the PRF can roughly be argued as follows. Since each of the ℓ elements $s_{1,x_1}, \dots, s_{\ell,x_\ell} \in \mathcal{D}$ that are part of the PRF key are sampled uniformly at random, the compressed $\ell/2$ elements $S(s_{1,x_1}, s_{2,x_2}), \dots, S(s_{\ell-1,x_{\ell-1}}, s_{\ell,x_\ell})$ are computationally indistinguishable from random elements in \mathcal{D} . This, in turn, implies that the compression of these $\ell/2$ elements into $\ell/4$ elements are pseudorandom. This argument can be applied iteratively to show that the final output of the PRF is computationally indistinguishable from uniform in \mathcal{D} .

LWE and Synthesizers. As pseudorandom synthesizers imply pseudorandom functions, one can naturally hope to construct a pseudorandom synthesizer from the LWE problem [48]. Recall that the $\text{LWE}_{n,q,\chi}$ assumption, parameterized by positive integers n, q and a B -bounded error distribution χ , states that for any (a priori unbounded) $m = \text{poly}(\lambda)$, if we sample a uniformly random secret vector $\mathbf{s} \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q^n$, uniformly random public vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q^n$, “small” error terms $e_1, \dots, e_\ell \leftarrow \chi^m$, and uniformly random values $u_1, \dots, u_m \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q$, the following

distributions are computationally indistinguishable:

$$\begin{aligned} (\mathbf{a}_1, \langle \mathbf{a}_1, \mathbf{s} \rangle + e_1) &\approx_c (\mathbf{a}_1, u_1) \\ (\mathbf{a}_2, \langle \mathbf{a}_2, \mathbf{s} \rangle + e_2) &\approx_c (\mathbf{a}_2, u_2) \\ &\vdots \\ (\mathbf{a}_m, \langle \mathbf{a}_m, \mathbf{s} \rangle + e_m) &\approx_c (\mathbf{a}_m, u_m). \end{aligned}$$

Given the $\text{LWE}_{n,q,\chi}$ assumption, it is natural to define a (randomized) pseudo-random synthesizer $S : \mathbb{Z}_q^{n \times n} \times \mathbb{Z}_q^{n \times n} \rightarrow \mathbb{Z}_q^{n \times n}$ as follows

$$S(\mathbf{S}, \mathbf{A}) = \mathbf{S} \cdot \mathbf{A} + \mathbf{E},$$

where the error matrix $\mathbf{E} \leftarrow \chi^{n \times n}$ is sampled randomly by the synthesizer S . It is easy to show via a standard hybrid argument that for any set of matrices $\mathbf{S}_1, \dots, \mathbf{S}_\ell \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times n}$, $\mathbf{A}_1, \dots, \mathbf{A}_\ell \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_q^{n \times n}$, and $\mathbf{E}_{1,1}, \dots, \mathbf{E}_{\ell,\ell} \leftarrow \chi^{n \times n}$, the pairwise applications of the synthesizer $S(\mathbf{S}_i, \mathbf{A}_j) = \mathbf{S}_i \cdot \mathbf{A}_j + \mathbf{E}_{i,j}$ for all $i, j \in [\ell]$ result in pseudorandom matrices.

Learning with Rounding. The problem with the synthesizer construction above is that the synthesizer must be randomized. Namely, in order to argue that the synthesizer's outputs are pseudorandom, the evaluation algorithm must flip random coins and sample independent error matrices $\mathbf{E}_{i,j} \leftarrow \chi^{n \times n}$ for each execution $S(\mathbf{S}_i, \mathbf{A}_j)$ for $i, j \in [\ell]$. Otherwise, if the error matrices are derived from an additional input to the synthesizer, then the error matrices for each evaluation of the synthesizer $S(\mathbf{S}_i, \mathbf{A}_j)$ for $i, j \in [\ell]$ must inevitably be correlated and hence, the security of the synthesizer cannot be shown from the hardness of $\text{LWE}_{n,q,\chi}$.

Banerjee, Peikert, and Rosen [11] provided a way to overcome this obstacle by introducing a way of derandomizing errors in LWE samples. The idea is quite simple and elegant: instead of adding small random error terms $e \leftarrow \chi$ to each inner product $\langle \mathbf{a}, \mathbf{s} \rangle \in \mathbb{Z}_q$, one can deterministically round it to one of $p < q$ partitions or "buckets" in \mathbb{Z}_q . Concretely, the idea can be implemented by applying the *modular rounding* operation $\lfloor \cdot \rfloor_p : \mathbb{Z}_q \rightarrow \mathbb{Z}_p$ to the inner product $\langle \mathbf{a}, \mathbf{s} \rangle$, which maps $\langle \mathbf{a}, \mathbf{s} \rangle \mapsto \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \cdot p/q \rfloor$. Intuitively, adding a small noise term $e \leftarrow \chi$ to the inner product of $\langle \mathbf{a}, \mathbf{s} \rangle$ in the $\text{LWE}_{n,q,\chi}$ problem *blinds* its low-ordered bits from a distinguisher. Therefore, applying the modular rounding operation to $\langle \mathbf{a}, \mathbf{s} \rangle$, which *removes* the low-ordered bits (albeit deterministically), should make the task of distinguishing it from random just as hard.

With this intuition, [11] introduced a new computational problem called the Learning with Rounding (LWR) problem. For parameters $n, q, p \in \mathbb{N}$ where $p < q$, the $\text{LWR}_{n,q,p}$ problem asks an adversary to distinguish the distributions:

$$\begin{aligned} (\mathbf{a}_1, \lfloor \langle \mathbf{a}_1, \mathbf{s} \rangle \rfloor_p) &\approx_c (\mathbf{a}_1, u_1) \\ (\mathbf{a}_2, \lfloor \langle \mathbf{a}_2, \mathbf{s} \rangle \rfloor_p) &\approx_c (\mathbf{a}_2, u_2) \\ &\vdots \\ (\mathbf{a}_m, \lfloor \langle \mathbf{a}_m, \mathbf{s} \rangle \rfloor_p) &\approx_c (\mathbf{a}_m, u_m), \end{aligned}$$

where $\mathbf{s} \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{a}_1, \dots, \mathbf{a}_m \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$, and $u_1, \dots, u_m \stackrel{R}{\leftarrow} \mathbb{Z}_p$. The hardness of the LWR problem then induces a deterministic pseudorandom synthesizer $S : \mathbb{Z}_q^{n \times n} \times \mathbb{Z}_q^{n \times n} \rightarrow \mathbb{Z}_p^{n \times n}$ that is defined as follows:

$$S(\mathbf{S}, \mathbf{A}) = \lfloor \mathbf{S} \cdot \mathbf{A} \rfloor_p,$$

where the modular rounding is done component-wise to each entry of the matrix $\mathbf{S} \cdot \mathbf{A} \in \mathbb{Z}_q^{n \times n}$.¹

Reducing LWE to LWR. Now, the remaining question is whether the LWR problem can formally be shown to be as hard as the LWE problem. The work of [11] gave a positive answer to this question. They showed that for any B -bounded distribution χ and moduli q and p for which $q = 2Bp\lambda^{\omega(1)}$, the $\text{LWR}_{n,q,p}$ problem is as hard as the $\text{LWE}_{n,q,\chi}$ problem.² Given an adversary for the $\text{LWR}_{n,q,p}$ problem \mathcal{A} , one can construct a simple algorithm \mathcal{B} that uses \mathcal{A} to solve $\text{LWE}_{n,q,\chi}$. Specifically, on input an $\text{LWE}_{n,q,\chi}$ challenge $(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_m, b_m) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, algorithm \mathcal{B} simply provides $(\mathbf{a}_1, \lfloor b_1 \rfloor_p), \dots, (\mathbf{a}_m, \lfloor b_m \rfloor_p)$ to \mathcal{A} .

- If the values $b_1, \dots, b_m \in \mathbb{Z}_q$ are noisy inner products $b_1 = \langle \mathbf{a}_1, \mathbf{s} \rangle + e_1, \dots, b_m = \langle \mathbf{a}_m, \mathbf{s} \rangle + e_m$, then we have

$$\lfloor b_i \rfloor_p = \lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \rfloor_p = \lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle \rfloor_p$$

for all $i \in [m]$ except with negligible probability over $\mathbf{a}_1, \dots, \mathbf{a}_m \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$.

- If the values b_1, \dots, b_m are uniformly random in \mathbb{Z}_q , then the values $\lfloor b_1 \rfloor_p, \dots, \lfloor b_m \rfloor_p$ are also uniform in \mathbb{Z}_p .

Hence, the algorithm \mathcal{B} statistically simulates the correct distribution of an $\text{LWR}_{n,q,p}$ challenge for \mathcal{A} and therefore, can solve the $\text{LWE}_{n,q,\chi}$ problem with essentially the same advantage as \mathcal{A} .

The apparent limitation of this reduction is the need for the modulus q to be super-polynomial in the security parameter. If q is only polynomial, then with noticeable probability, the inner product $\langle \mathbf{a}_i, \mathbf{s} \rangle \in \mathbb{Z}_q$ for any $i \in [m]$ lands on a rounding “borderline” set

$$\begin{aligned} \text{Borderline}_B &= [-B, B] + q/p \cdot (\mathbb{Z} + 1/2) \\ &= \{ v \in \mathbb{Z}_q \mid \exists e \in [-B, B] \text{ such that } \lfloor v \rfloor_p \neq \lfloor v + e \rfloor_p \}, \end{aligned}$$

¹Note that the synthesizer maps matrices in $\mathbb{Z}_q^{n \times n}$ to matrices in $\mathbb{Z}_p^{n \times n}$ for $p < q$ and hence violates the original syntax of a synthesizer. This is a minor technicality, which can be addressed in multiple ways. One option is to have a sequence of rounding moduli $p_1, \dots, p_{\log \ell}$ such that the synthesizer can be applied iteratively. Another option is to take a pseudorandom generator $G : \mathbb{Z}_p^{n \times n} \rightarrow \mathbb{Z}_q^{n \times n}$ and define $S(\mathbf{S}, \mathbf{A}) = G(\lfloor \mathbf{S} \cdot \mathbf{A} \rfloor_p)$. Finally, one can define the synthesizer $S(\mathbf{S}, \mathbf{A}) = \lfloor \mathbf{S} \cdot \mathbf{A} \rfloor_p$ with respect to non-square matrices $S : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^{n \times m} \rightarrow \mathbb{Z}_p^{m \times m}$ such that the sets $\mathbb{Z}_q^{m \times n}$, $\mathbb{Z}_q^{n \times m}$, and $\mathbb{Z}_p^{m \times m}$ have the same cardinality $|\mathbb{Z}_q^{m \times n}| = |\mathbb{Z}_q^{n \times m}| = |\mathbb{Z}_p^{m \times m}|$.

²For the reduction to succeed with probability $1 - 2^{-\lambda}$, we must set $q \geq 2Bp \cdot 2^\lambda$. For simplicity throughout the overview, we restrict to the “negligible vs. non-negligible” type security as opposed to 2^λ -security. See Section 1.2 for further discussions.

and hence $\lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \rfloor_p \neq \lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle \rfloor_p$. In this case, one cannot guarantee that an adversary \mathcal{A} for the $\text{LWR}_{n,q,p}$ problem correctly distinguishes the samples $(\mathbf{a}_1, \lfloor \langle \mathbf{a}_1, \mathbf{s} \rangle + e_1 \rfloor_p), \dots, (\mathbf{a}_m, \lfloor \langle \mathbf{a}_m, \mathbf{s} \rangle + e_m \rfloor_p)$ from purely random samples.

2.2 Learning with Rounding and Errors.

Chaining LWE Samples. We get around the limitation of the reduction above by using what we call the *chaining method*. To demonstrate the idea, let us consider a challenge oracle $\mathcal{O}_{\tau, \mathbf{S}}$ that chains multiple $\text{LWE}_{n,q,\chi}$ samples together. The oracle $\mathcal{O}_{\tau, \mathbf{S}}$ is parameterized by a chaining parameter $\tau \in \mathbb{N}$, a set of secret vectors $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_\tau) \in \mathbb{Z}_q^{n \times \tau}$, and is defined with respect to a sampling algorithm $D_\chi : \{0, 1\}^{\lceil \log p \rceil} \rightarrow \mathbb{Z}$, which takes in as input $\lceil \log p \rceil$ random coins and samples an error value $e \in \mathbb{Z}$ according to the B -bounded error distribution χ .

- $\mathcal{O}_{\tau, \mathbf{S}}$: On its invocation, the oracle samples a public vector $\mathbf{a} \xleftarrow{R} \mathbb{Z}_q^n$ and an error term $e_1 \leftarrow \chi$. Then, for $1 \leq i < \tau$, it iteratively computes:
 - $r_i \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_i \rangle + e_i \rfloor_p$.
 - $e_{i+1} \leftarrow D_\chi(r_i)$.
It then returns $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p)$.

In words, the oracle $\mathcal{O}_{\tau, \mathbf{S}}$ generates (the rounding of) an $\text{LWE}_{n,q,\chi}$ sample $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s}_1 \rangle + e_1 \rfloor_p)$ and uses $r_1 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_1 \rangle + e_1 \rfloor_p$ as the random coins to sample $e_2 \leftarrow D_\chi(r_1)$. It then computes $r_2 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_2 \rangle + e_2 \rfloor_p$ and uses r_2 to sample the next error term $e_3 \leftarrow D_\chi(r_2)$ for the next iteration. The oracle iterates this procedure for τ steps and finally returns $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p)$.

Now, suppose that p divides q . Then a hybrid argument shows that assuming the hardness of $\text{LWE}_{n,q,\chi}$, a sample $(\mathbf{a}, b) \leftarrow \mathcal{O}_{\tau, \mathbf{S}}$ is computationally indistinguishable from uniform in $\mathbb{Z}_q^n \times \mathbb{Z}_p$. Specifically, we can argue that the first term (random coins) $r_1 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_1 \rangle + e_1 \rfloor_p$ is computationally indistinguishable from uniform in $\{0, 1\}^{\lceil \log p \rceil}$ by the hardness of $\text{LWE}_{n,q,\chi}$. Then, since r_1 is uniform, the error term $e_2 \leftarrow D_\chi(r_1)$ is correctly distributed according to χ , which implies that $r_2 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_2 \rangle + e_2 \rfloor_p$ is also computationally indistinguishable from uniform. Continuing this argument for τ iterations, we can prove that the final output $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p)$ is computationally indistinguishable from uniform in $\mathbb{Z}_q^n \times \mathbb{Z}_p$.

Chaining LWR and LWE Samples. So far, it seems as if we have not made much progress. Although the oracle $\mathcal{O}_{\tau, \mathbf{S}}$ returns an “LWR looking” sample $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$, it must still randomly sample the initial noise term $e_1 \leftarrow \chi$, which makes it useless for constructing a deterministic pseudorandom synthesizer. Our key observation, however, is that when the chaining parameter τ is big enough, then the initial error term e_1 does not affect the final output $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p)$ with overwhelming probability. In other words, e_1 can always be set to be 0 without negatively impacting the pseudorandomness of $\mathcal{O}_{\tau, \mathbf{S}}$.

To see this, consider the following modification of the oracle $\mathcal{O}_{\tau, \mathbf{S}}$:

- $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{l wre})}$: On its invocation, the oracle samples a public vector $\mathbf{a} \xleftarrow{R} \mathbb{Z}_q^n$ and initializes $e_1 = 0$. Then, for $1 \leq i < \tau$, it iteratively computes:

- $r_i \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_i \rangle + e_i \rfloor_p$.
- $e_{i+1} \leftarrow D_\chi(r_i)$.

It returns $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p)$.

In contrast to $\mathcal{O}_{\tau, \mathbf{S}}$, the oracle $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ derives the first set of random coins r_1 from an errorless $\text{LWR}_{n, q, p}$ sample $r_1 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_1 \rangle \rfloor_p$. It then uses r_1 to sample the error term $e_2 \leftarrow D_\chi(r_1)$ for the next $\text{LWE}_{n, q, \chi}$ sample to derive $r_2 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_2 \rangle + e_2 \rfloor_p$, and it continues this procedure for τ iterations. As the oracle $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ is, in effect, chaining $\text{LWR}_{n, q, p}$ and $\text{LWE}_{n, q, \chi}$ samples together, we refer to $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ as the *Learning with (both) Rounding and Errors* (LWRE) oracle.

We claim that even when q is small, as long as the chaining parameter τ is big enough, the samples that are output by the oracles $\mathcal{O}_{\tau, \mathbf{S}}$ and $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ are identical except with negligible probability. For simplicity, let us fix the modulus to be $q = 4Bp$ such that for $u \xleftarrow{\mathbb{R}} \mathbb{Z}_q$, $e_1, e_2 \leftarrow \chi$, we have

$$\Pr[\lfloor u + e_1 \rfloor_p \neq \lfloor u + e_2 \rfloor_p] \leq \frac{1}{2}. \quad (2.1)$$

Now, consider a transcript of an execution of the oracles $\mathcal{O}_{\tau, \mathbf{S}}$ and $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ for $\mathbf{S} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{n \times \tau}$, $\mathbf{a} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$, and any fixed error value $e_1 \in [-B, B]$:

$$\begin{array}{ll} \mathcal{O}_{\tau, \mathbf{S}} : & \mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})} : \\ r_1 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_1 \rangle + e_1 \rfloor_p & \tilde{r}_1 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_1 \rangle + 0 \rfloor_p \\ r_2 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_2 \rangle + e_2 \rfloor_p & \tilde{r}_2 \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_2 \rangle + \tilde{e}_2 \rfloor_p \\ \vdots & \vdots \\ r_\tau \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p & \tilde{r}_\tau \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + \tilde{e}_\tau \rfloor_p \end{array}$$

We make the following observations:

1. Since the sampler D_χ is deterministic, if there exists an index $1 \leq i^* \leq \tau$ for which $r_{i^*} = \tilde{r}_{i^*}$, then this implies that $r_i = \tilde{r}_i$ for all $i^* \leq i \leq \tau$.
2. Since the vectors $\mathbf{s}_1, \dots, \mathbf{s}_\tau$ are sampled uniformly at random from \mathbb{Z}_q^n , the inner products $\langle \mathbf{a}, \mathbf{s}_i \rangle$ for any $1 \leq i \leq \tau$ are distributed statistically close to uniform in \mathbb{Z}_q .³ Therefore, using (2.1), we have

$$\Pr[\lfloor \langle \mathbf{a}, \mathbf{s}_i \rangle + e_i \rfloor_p \neq \lfloor \langle \mathbf{a}, \mathbf{s}_i \rangle + \tilde{e}_i \rfloor_p] \leq \frac{1}{2} + \text{negl},$$

for any $1 \leq i \leq \tau$.

These observations imply that unless all of the inner products $\langle \mathbf{a}, \mathbf{s}_1 \rangle, \dots, \langle \mathbf{a}, \mathbf{s}_\tau \rangle$ land on the “borderline” set of \mathbb{Z}_q , the samples of $\mathcal{O}_{\tau, \mathbf{S}}$ and $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ coincide for

³When the modulus q is prime, then the inner product $\langle \mathbf{a}, \mathbf{s}_i \rangle$ is certainly uniform in \mathbb{Z}_q . Even when q is composite (i.e. q is divisible by p), under mild requirements on q , the inner product $\langle \mathbf{a}, \mathbf{s}_i \rangle$ is statistically close to uniform in \mathbb{Z}_q .

$\mathbf{a} \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$. Furthermore, such bad event occurs with probability at most $\approx 1/2^\tau$. Hence, even for very small values of the chaining parameter $\tau = \omega(\log \lambda)$, with overwhelming probability over the matrix $\mathbf{S} \stackrel{R}{\leftarrow} \mathbb{Z}_q^{n \times \tau}$, no information about the initial error term e_1 is revealed from a single sample of $\mathcal{O}_{\tau, \mathbf{S}}$ or $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$. This can be extended to argue that no information about e_1 is leaked from any polynomial number of samples via the union bound. Hence, for $\tau = \omega(\log \lambda)$, any polynomial number of samples from $\mathcal{O}_{\tau, \mathbf{S}}$ or $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ are statistically indistinguishable.

In the discussion above, we set the modulus $q = 4Bp$ purely for simplicity. If we set q to be slightly greater than $2Bp$ (by a polynomial factor), the chaining parameter τ can be set to be any super-constant function $\tau = \omega(1)$ to guarantee that the output of the oracles $\mathcal{O}_{\tau, \mathbf{S}}$ and $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ are statistically indistinguishable.

Synthesizer from LWRE. The oracle $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ naturally induces a computational problem. Namely, for a set of parameters n, q, p, χ , and τ , we define the $\text{LWRE}_{n, q, p, \chi, \tau}$ problem that asks an adversary to distinguish the samples $(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_m, b_m) \leftarrow \mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ from uniformly random samples $(\mathbf{a}_1, \hat{b}_1), \dots, (\mathbf{a}_m, \hat{b}_m) \stackrel{R}{\leftarrow} \mathbb{Z}_q^n \times \mathbb{Z}_p$. Using the ideas highlighted above, we can show that for small values of q and τ , the $\text{LWRE}_{n, q, p, \chi, \tau}$ is at least as hard as the $\text{LWE}_{n, m, q, \chi}$ problem. Specifically, we can first show that the oracle $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ is statistically indistinguishable from $\mathcal{O}_{\tau, \mathbf{S}}$. Then, via a hybrid argument, we can show that the oracle $\mathcal{O}_{\tau, \mathbf{S}}$ is computationally indistinguishable from a uniform sampler over $\mathbb{Z}_q^n \times \mathbb{Z}_p$ by the hardness of $\text{LWE}_{n, q, \chi}$.

The $\text{LWRE}_{n, q, p, \chi, \tau}$ problem naturally induces a pseudorandom synthesizer. One can first define an “almost” pseudorandom synthesizer $\mathcal{G} : \mathbb{Z}_q^n \times \mathbb{Z}_q^{n \times \tau} \rightarrow \mathbb{Z}_p$ that emulates the $\text{LWRE}_{n, q, p, \chi, \tau}$ oracle as follows:

- $\mathcal{G}(\mathbf{a}, \mathbf{S})$: On input $\mathbf{a} \in \mathbb{Z}_q^n$ and $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_\tau) \in \mathbb{Z}_q^{n \times \tau}$, the LWRE function sets $e_1 = 0$ and computes for $i = 1, \dots, \tau - 1$:
 1. $r_i \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_i \rangle + e_i \rfloor_p$,
 2. $e_{i+1} \leftarrow D_\chi(r_i)$.
 It then sets $b = \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p$ and returns $b \in \mathbb{Z}_p$.

It is easy to see that as long as the $\text{LWRE}_{n, q, p, \chi, \tau}$ problem is hard, then for any $\ell = \text{poly}(\lambda)$, $\mathbf{a}_1, \dots, \mathbf{a}_\ell \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{S}_1, \dots, \mathbf{S}_\ell \leftarrow \mathbb{Z}_q^{n \times \tau}$, and $u_{1,1}, \dots, u_{\ell, \ell} \stackrel{R}{\leftarrow} \mathbb{Z}_p$, we have

$$\left\{ \mathcal{G}(\mathbf{a}_i, \mathbf{S}_j) \right\}_{i, j \in [\ell]} \approx_c \left\{ u_{i, j} \right\}_{i, j \in [\ell]} \in \mathbb{Z}_p^{\ell^2}.$$

Furthermore, this indistinguishability holds even for small values of the chaining parameter $\tau = \omega(1)$ and hence, the function \mathcal{G} can be computed by shallow circuits.

The only reason why $\mathcal{G} : \mathbb{Z}_q^n \times \mathbb{Z}_q^{n \times \tau} \rightarrow \mathbb{Z}_p$ is not a pseudorandom synthesizer is that the cardinality of the sets \mathbb{Z}_q^n , $\mathbb{Z}_q^{n \times \tau}$, and \mathbb{Z}_p are different. However, this can easily be fixed by defining a synthesizer $S : (\mathbb{Z}_q^n)^{\ell_1} \times (\mathbb{Z}_q^{n \times \tau})^{\ell_2} \rightarrow (\mathbb{Z}_p)^{\ell_1 \times \ell_2}$ that takes in a set of ℓ_1 vectors $(\mathbf{a}_1, \dots, \mathbf{a}_{\ell_1}) \in (\mathbb{Z}_q^n)^{\ell_1}$, and ℓ_2 matrices $(\mathbf{S}_1, \dots, \mathbf{S}_{\ell_2}) \in$

$(\mathbb{Z}_q^{n \times \tau})^{\ell_2}$, and then returns $\{\mathcal{G}(\mathbf{a}_i, \mathbf{S}_j)\}_{i,j \in [\ell]}$. The parameters ℓ_1 and ℓ_2 can be set to be any positive integers such that

$$|(\mathbb{Z}_q^n)^{\ell_1}| = |(\mathbb{Z}_q^{n \times \tau})^{\ell_2}| = |\mathbb{Z}_p^{\ell_1 \times \ell_2}|,$$

which makes S to be a two-to-one function over a fixed domain. The PRF that is induced by the synthesizer $S : (\mathbb{Z}_q^n)^{\ell_1} \times (\mathbb{Z}_q^{n \times \tau})^{\ell_2} \rightarrow (\mathbb{Z}_p)^{\ell_1 \times \ell_2}$ corresponds to our first PRF construction.

We note that for practical implementation of the synthesizer, the large PRF key can be derived from a λ -bit PRG seed. Furthermore, the discrete Gaussian sampler D_χ can always be replaced by a suitable look-up table with pre-computed Gaussian noise as the modulus p is small. Therefore, the synthesizer can be implemented quite efficiently as it simply consists of τ inner products of two vectors modulo a small integer and their rounding. We refer to Section 4.3 for a discussion on the parameters and implementations.

2.3 Chaining Key-Homomorphic PRFs

The method of chaining multiple LWR/LWE samples can also be applied directly to existing lattice-based PRF constructions to improve their parameters. Furthermore, when applied to existing key-homomorphic PRFs, the resulting PRF is also key-homomorphic. Here, we demonstrate the idea with the PRF construction of [17] as it can be described quite compactly. In the technical section (Section 5), we show how to chain the PRF construction of [10] as it is a generalization of the previous PRF constructions of [11, 17] and it also allows us to set the parameters such that the underlying modulus q is only polynomial in the security parameter.

Background on BLMR [17]. Recall that the BLMR PRF is defined with respect to two public binary matrices $\mathbf{A}_0, \mathbf{A}_1 \in \{0, 1\}^{n \times n}$ for a suitable choice of $n = \text{poly}(\lambda)$. A PRF key is set to be a vector $\mathbf{s} \in \mathbb{Z}_q^n$, and the PRF evaluation for an input $x \in \{0, 1\}^\ell$ is defined to be the rounded matrix product

$$F^{(\text{BLMR})}(\mathbf{s}, x) = \left\lfloor \mathbf{s}^\top \prod_{i=1}^{\ell} \mathbf{A}_{x_i} \right\rfloor_p.$$

We can reason about the security of the BLMR PRF by considering its “noisy” variant that is defined as follows:

$F^{(\text{noise})}(\mathbf{s}, x)$:

1. Sample error vectors $\mathbf{e}_1, \dots, \mathbf{e}_\ell \leftarrow \chi^n$,
2. Return the vector

$$\begin{aligned} & \left\lfloor \left(\left((\mathbf{s}^\top \mathbf{A}_{x_1} + \mathbf{e}_1^\top) \cdot \mathbf{A}_{x_2} + \mathbf{e}_2^\top \right) \cdots \right) \cdot \mathbf{A}_{x_\ell} + \mathbf{e}_\ell^\top \right\rfloor_p \\ &= \left\lfloor \mathbf{s}^\top \prod_{i=1}^{\ell} \mathbf{A}_{x_i} + \underbrace{\sum_{i=1}^{\ell-1} \mathbf{e}_i^\top \prod_{j=i+1}^{\ell-1} \mathbf{A}_{x_j}}_{\mathbf{e}^*} + \mathbf{e}_\ell^\top \right\rfloor_p. \end{aligned}$$

Since the error vectors $\mathbf{e}_1, \dots, \mathbf{e}_\ell \leftarrow \chi^n$ has small norm and the matrices $\mathbf{A}_0, \mathbf{A}_1 \in \{0, 1\}^{n \times n}$ are binary, the error term \mathbf{e}^* is also small. Therefore, if the modulus q is sufficiently big, then with overwhelming probability, the error vector \mathbf{e}^* is “rounded away” and does not contribute to the final output of the PRF. This shows that when the modulus q is big, the evaluations of the functions $F(\mathbf{s}, \cdot)$ and $F^{(\text{noise})}(\mathbf{s}, \cdot)$ are statistically indistinguishable.

Now, it is easy to show that $F^{(\text{noise})}(\mathbf{s}, \cdot)$ is computationally indistinguishable from a truly random function using the hardness of the $\text{LWE}_{n,q,\chi}$ problem.⁴ We can first argue that the vector $\mathbf{s}_1^\top \mathbf{A}_{1,x_1} + \mathbf{e}_1^\top$ is computationally indistinguishable from a uniformly random vector $\mathbf{s}_2 \xleftarrow{R} \mathbb{Z}_q^n$. This implies that the vector $\mathbf{s}_2^\top \mathbf{A}_{2,x_2} + \mathbf{e}_2^\top$ is computationally indistinguishable from a random vector $\mathbf{s}_3 \xleftarrow{R} \mathbb{Z}_q^n$. We can repeat the argument for ℓ steps to prove that the final output of the PRF is computationally indistinguishable from a uniformly random output.

Chaining BLMR. For the security argument of the BLMR PRF to be valid, it is crucial that the modulus q is large enough such that the error term \mathbf{e}^* rounds away with the modular rounding operation. Specifically, the modulus q must be set to be greater than the maximum possible norm of the error term $\|\mathbf{e}^*\|$ by a super-polynomial factor in the security parameter.

To prevent this blow-up in the size of q , we can chain multiple instances of the functions $F^{(\text{BLMR})}$ and $F^{(\text{noise})}$ together. Consider the following chained PRF $F^{(\text{chain})} : \mathbb{Z}_q^{n \times \tau} \times \{0, 1\}^\ell \rightarrow \mathbb{Z}_p^n$:

- $F^{(\text{chain})}(\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_\tau), x)$:
1. Evaluate $r_2 \leftarrow F^{(\text{BLMR})}(\mathbf{s}_1, x)$.
 2. For $i = 2, \dots, \tau - 1$, compute:
 - $r_{i+1} \leftarrow F^{(\text{noise})}(\mathbf{s}_i, x; r_i)$.
 3. Return $F^{(\text{noise})}(\mathbf{s}_\tau, x; r_\tau)$.

In words, the chained PRF $F^{(\text{chain})}(\mathbf{s}, x)$ evaluates the “random coins” that are needed to evaluate the randomized PRF $F^{(\text{noise})}(\mathbf{s}_i, x)$ from the previous execution of $F^{(\text{noise})}(\mathbf{s}_{i-1}, x)$. The initial random coins are derived from the errorless BLMR PRF $F^{(\text{BLMR})}(\mathbf{s}_1, x)$.

We can prove that the chained PRF $F^{(\text{chain})}$ is secure using the same argument that was used to show the hardness of the LWRE problem. Namely, we first argue that even if the modulus q is greater than the maximum possible norm of the error term $\|\mathbf{e}^*\|$ only by a polynomial factor, for any two random coins r_i, r'_i , we have $F^{(\text{noise})}(\mathbf{s}_i, x; r_i) = F^{(\text{noise})}(\mathbf{s}_i, x; r'_i)$ with noticeable probability. Therefore, if we set τ to be sufficiently big, then we can replace $F^{(\text{BLMR})}(\mathbf{s}_1, \cdot)$ with the randomized function $F^{(\text{noise})}(\mathbf{s}_1, \cdot)$ without changing the output of the PRF. Now, we can use the fact that $F^{(\text{noise})}(\mathbf{s}_1, \cdot)$ is computationally indistinguishable from a

⁴Technically, the reduction uses the hardness of the *non-uniform* variant of the LWE problem [17] where the challenger samples the public vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ uniformly at random from $\{0, 1\}^n$ as opposed to sampling them from \mathbb{Z}_q^n . The work of [17] shows that this variant of the LWE problem is as hard as the traditional version of LWE for suitable choices of parameters.

truly random function to argue that the function $F^{(\text{chain})}(\mathbf{S}, \cdot)$ is computationally indistinguishable from a truly random function.

The parameters for the chained PRF $F^{(\text{chain})}$ provide a trade-off between the depth of the evaluation circuit and the size of the modulus q (and therefore, the quality of the LWE assumption). If we set $\tau = 1$, then we recover the original BLMR PRF, which requires very large values of the modulus q . As we increase τ , the modulus q can be set to be arbitrarily close to the maximum possible value of the error vector $\|\mathbf{e}^*\|$ for $F^{(\text{noise})}(\mathbf{s}, x)$. For the Banerjee-Peikert PRF [10], the maximum possible value of the error vector $\|\mathbf{e}^*\|$ can be made to be only polynomial in the security parameter, thereby allowing us to set q to be a polynomial function of the security parameter.

Key-homomorphism. The BLMR PRF is key-homomorphic because the modular rounding operation is an almost linear operation. Namely, for any input $x \in \{0, 1\}^\ell$ and any two keys $\mathbf{s}, \tilde{\mathbf{s}} \in \mathbb{Z}_q^n$, we have

$$\begin{aligned} F^{(\text{BLMR})}(\mathbf{s}, x) + F^{(\text{BLMR})}(\tilde{\mathbf{s}}, x) &= \left[\mathbf{s}^\top \prod_{i=1}^{\ell} \mathbf{A}_{x_i} \right]_p + \left[\tilde{\mathbf{s}}^\top \prod_{i=1}^{\ell} \mathbf{A}_{x_i} \right]_p \\ &\approx \left[(\mathbf{s} + \tilde{\mathbf{s}})^\top \prod_{i=1}^{\ell} \mathbf{A}_{x_i} \right]_p \\ &= F^{(\text{BLMR})}(\mathbf{s} + \tilde{\mathbf{s}}, x). \end{aligned}$$

Due to this algebraic structure, the “noisy” variant of the BLMR PRF is also key-homomorphic. Namely, for any input $x \in \{0, 1\}^\ell$, any two keys $\mathbf{s}, \tilde{\mathbf{s}} \in \mathbb{Z}_q^n$, and any *random coins* r, \tilde{r}, r' that is used to sample the noise, we have

$$\begin{aligned} F^{(\text{noise})}(\mathbf{s}, x; r) + F^{(\text{noise})}(\tilde{\mathbf{s}}, x; \tilde{r}) &= \left[(\mathbf{s} + \tilde{\mathbf{s}})^\top \prod_{i=1}^{\ell} \mathbf{A}_{x_i} + (\mathbf{e}^* + \tilde{\mathbf{e}}^*) \right]_p \\ &\approx F^{(\text{noise})}(\mathbf{s} + \tilde{\mathbf{s}}, x; r'), \end{aligned}$$

Now, note that the final output of the chained PRF $F^{(\text{chain})}$ on an input $x \in \{0, 1\}^\ell$ and a key $(\mathbf{s}_1, \dots, \mathbf{s}_\tau) \in \mathbb{Z}_q^{n \times \tau}$ with chaining parameter τ is simply the output of the noisy PRF $F^{(\text{noise})}(\mathbf{s}_\tau, x; r_\tau)$ where r_τ is the randomness that is derived from the previous execution of $r_\tau \leftarrow F^{(\text{noise})}(\mathbf{s}_{\tau-1}, x; r_{\tau-1})$. Therefore, for any input $x \in \{0, 1\}^\ell$, and two keys $\mathbf{S}, \tilde{\mathbf{S}} \in \mathbb{Z}_q^{n \times \tau}$, we can show that

$$\begin{aligned} F^{(\text{chain})}(\mathbf{S}, x) + F^{(\text{chain})}(\tilde{\mathbf{S}}, x) &= F^{(\text{noise})}(\mathbf{s}_\tau, x; r_\tau) + F^{(\text{noise})}(\tilde{\mathbf{s}}_\tau, x; \tilde{r}_\tau) \\ &\approx F^{(\text{noise})}(\mathbf{s}_\tau + \tilde{\mathbf{s}}_\tau, x) \\ &\approx F^{(\text{chain})}(\mathbf{S} + \tilde{\mathbf{S}}, x) \end{aligned}$$

Specifically, we can show that for a suitable choice of the modulus q , the resulting PRF is *2-almost* key-homomorphic in that for any input $x \in \{0, 1\}^\ell$ and two keys $\mathbf{S}, \tilde{\mathbf{S}} \in \mathbb{Z}_q^{n \times \tau}$, there exists an error vector $\boldsymbol{\eta} \in [0, 2]^n$ such that

$$F^{(\text{chain})}(\mathbf{S}, x) + F^{(\text{chain})}(\tilde{\mathbf{S}}, x) = F^{(\text{chain})}(\mathbf{S} + \tilde{\mathbf{S}}, x) + \boldsymbol{\eta}.$$

The exact argument to show that the chained BLMR PRF is 2-almost key-homomorphic can be used to show that the chained BP PRF [10] is also 2-almost key-homomorphic. We provide the formal details in Section 5.

3 Preliminaries

Basic notations. Unless specified otherwise, we use λ to denote the security parameter. We say a function $f(\lambda)$ is negligible in λ , denoted by $\text{negl}(\lambda)$, if $f(\lambda) = o(1/\lambda^c)$ for all $c \in \mathbb{N}$. We say that a function $f(\lambda)$ is noticeable in λ if $f(\lambda) = \Omega(1/\lambda^c)$ for some $c \in \mathbb{N}$. We say that an event happens with overwhelming probability if its complement happens with negligible probability. We say that an algorithm is efficient if it runs in probabilistic polynomial time in the length of its input. We use $\text{poly}(\lambda)$ to denote a quantity whose value is bounded by a fixed polynomial in λ .

For an integer $n \geq 1$, we write $[n]$ to denote the set of integers $\{1, \dots, n\}$. For a distribution \mathcal{D} , we write $x \leftarrow \mathcal{D}$ to denote that x is sampled from \mathcal{D} ; for a finite set S , we write $x \stackrel{\mathcal{R}}{\leftarrow} S$ to denote that x is sampled uniformly from S . For a positive integer B , we say that a distribution \mathcal{D} over \mathbb{Z} is B -bounded if $\Pr[x \leftarrow \mathcal{D} \wedge |x| > B]$ is negligible. Finally, we write $\text{Funcs}[\mathcal{X}, \mathcal{Y}]$ to denote the set of all functions mapping from a domain \mathcal{X} to a range \mathcal{Y} .

Vectors and matrices. We use bold lowercase letters (*e.g.*, \mathbf{v}, \mathbf{w}) to denote vectors and bold uppercase letters (*e.g.*, \mathbf{A}, \mathbf{B}) to denote matrices. Throughout this work, we always use the infinity norm for vectors and matrices. Therefore, for a vector \mathbf{x} , we write $\|\mathbf{x}\|$ to denote $\max_i |x_i|$. Similarly, for a matrix \mathbf{A} , we write $\|\mathbf{A}\|$ to denote $\max_{i,j} |A_{i,j}|$. If $\mathbf{x} \in \mathbb{Z}^n$ and $\mathbf{A} \in \mathbb{Z}^{n \times m}$, then $\|\mathbf{x}\mathbf{A}\| \leq n \cdot \|\mathbf{x}\| \cdot \|\mathbf{A}\|$.

Modular rounding. For an integer $p \leq q$, we define the modular “rounding” function

$$\lfloor \cdot \rfloor_p : \mathbb{Z}_q \rightarrow \mathbb{Z}_p \text{ that maps } x \rightarrow \lfloor (p/q) \cdot x \rfloor$$

and extend it coordinate-wise to matrices and vectors over \mathbb{Z}_q . Here, the operation $\lfloor \cdot \rfloor$ is the integer rounding operation over the real numbers. It can be readily checked that for any two values $x, y \in \mathbb{Z}_q$, there exists some $\eta \in \{0, 1\}$ such that $\lfloor x \rfloor_p + \lfloor y \rfloor_p = \lfloor x + y \rfloor_p + \eta$.

Bit-decomposition. Let n and q be positive integers. Then we define the “gadget matrix” $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}_n \in \mathbb{Z}_q^{n \times n \cdot \lceil \log q \rceil}$ where $\mathbf{g} = (1, 2, 4, \dots, 2^{\lceil \log q \rceil - 1})$. We define the inverse bit-decomposition function $\mathbf{G}^{-1} : \mathbb{Z}_q^{n \times m} \rightarrow \mathbb{Z}_q^{n \cdot \lceil \log q \rceil \times m}$ which expands each entry $x \in \mathbb{Z}_q$ in the input matrix into a column of size $\lceil \log q \rceil$ that consists of the bits of the binary representation of x .

3.1 Learning with Errors

In this section, we define the Learning with Errors (LWE) problem [48].

Learning with Errors. We define the LWE problem with respect to the real and ideal oracles.

Definition 3.1 (Learning with Errors). Let $\lambda \in \mathbb{N}$ be the security parameter. Then the learning with errors (LWE) problem is parameterized by a dimension $n = n(\lambda)$, modulus $q = q(\lambda)$, and error distribution $\chi = \chi(\lambda)$. It is defined with respect to the real and ideal oracles $\mathcal{O}_s^{(\text{lwe})}$ and $\mathcal{O}^{(\text{ideal})}$ that are defined as follows:

- $\mathcal{O}_s^{(\text{lwe})}$: The real oracle is parameterized by a vector $\mathbf{s} \in \mathbb{Z}_q^n$. On its invocation, the oracle samples a random vector $\mathbf{a} \xleftarrow{\text{R}} \mathbb{Z}_q^n$, and an error term $e \leftarrow \chi$. It sets $b = \langle \mathbf{a}, \mathbf{s} \rangle + e$, and returns $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.
- $\mathcal{O}^{(\text{ideal})}$: On its invocation, the ideal oracle samples a random vector $\mathbf{a} \xleftarrow{\text{R}} \mathbb{Z}_q^n$, random element $u \xleftarrow{\text{R}} \mathbb{Z}_q$, and returns $(\mathbf{a}, u) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.

The $\text{LWE}_{n,q,\chi}$ problem is to distinguish the oracles $\mathcal{O}_s^{(\text{lwe})}$ and $\mathcal{O}^{(\text{ideal})}$. More precisely, we define an adversary \mathcal{A} 's distinguishing advantage $\text{Adv}_{\text{LWE}}(n, q, \chi, \mathcal{A})$ as the probability

$$\text{Adv}_{\text{LWE}}(n, q, \chi, \mathcal{A}) = \left| \Pr [\mathcal{A}^{\mathcal{O}_s^{(\text{lwe})}}(1^\lambda) = 1] - \Pr [\mathcal{A}^{\mathcal{O}^{(\text{ideal})}}(1^\lambda) = 1] \right|,$$

where $\mathbf{s} \xleftarrow{\text{R}} \mathbb{Z}_q^n$. The $\text{LWE}_{n,q,\chi}$ assumption states that for any efficient adversary \mathcal{A} , its distinguishing advantage is negligible $\text{Adv}_{\text{LWE}}(n, q, \chi, \mathcal{A}) = \text{negl}(\lambda)$.

Compactly, the $\text{LWE}_{n,q,\chi}$ assumption states that for any $m = \text{poly}(\lambda)$, $\mathbf{s} \xleftarrow{\text{R}} \mathbb{Z}_q^n$, $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, the noisy vector-matrix product $(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)$ is computationally indistinguishable from $(\mathbf{A}, \mathbf{u}^\top)$. It follows from a standard hybrid argument that for any $m, \ell = \text{poly}(\lambda)$, $\mathbf{S} \xleftarrow{\text{R}} \mathbb{Z}_q^{\ell \times n}$, $\mathbf{A} \xleftarrow{\text{R}} \mathbb{Z}_q^{n \times m}$, $\mathbf{E} \leftarrow \chi^{\ell \times m}$, and $\mathbf{U} \xleftarrow{\text{R}} \mathbb{Z}_q^{\ell \times m}$, the noisy matrix product $(\mathbf{A}, \mathbf{S} \cdot \mathbf{A} + \mathbf{E})$ is computationally indistinguishable from (\mathbf{A}, \mathbf{U}) by the $\text{LWE}_{n,q,\chi}$ assumption.

Let $n = \text{poly}(\lambda)$ and χ be a B -bounded discrete Gaussian distribution. Then the $\text{LWE}_{n,q,\chi}$ assumption is true assuming that various worst-case lattice problems such as GapSVP and SIVP on n -dimensional lattices are hard to approximate to within a factor of $\tilde{O}(n \cdot q/B)$ by a quantum algorithm [48]. Similar reductions of $\text{LWE}_{n,q,\chi}$ to the classical hardness of approximating worst-case lattice problems are also known [46, 9, 41, 42, 20].

3.2 Elementary Number Theory

In this section, we state and prove an elementary fact in number theory that we use for our technical sections. Specifically, we analyze the distribution of the inner product of two uniformly random vectors $\langle \mathbf{a}, \mathbf{s} \rangle$ where $\mathbf{a}, \mathbf{s} \xleftarrow{\text{R}} \mathbb{Z}_q^n$ for some positive integers n and q . When n is sufficiently big, then with overwhelming probability, one of the components of \mathbf{a} (or \mathbf{s}) will be a multiplicative unit in \mathbb{Z}_q and therefore, the inner product $\langle \mathbf{a}, \mathbf{s} \rangle$ will be uniform in \mathbb{Z}_q . We formally state and prove this elementary fact in the lemma below. We first recall a general fact of the Euler totient function.

Fact 3.2 ([39]). Let $\varphi : \mathbb{Z} \rightarrow \mathbb{Z}$ be the Euler totient function. Then, there exists a constant c such that for any $q > 264$, we have $q/\varphi(q) = c \cdot \log \log q$.

The following lemma follows immediately from Fact 3.2.

Lemma 3.3. *Let $n = n(\lambda)$ and $q = q(\lambda)$ be positive integers such that $n = \Omega(\lambda \log \log q)$. Then, for any element $d \in \mathbb{Z}_q$, we have*

$$\Pr[\langle \mathbf{a}, \mathbf{s} \rangle = d] \leq 1/q + 2^{-\lambda},$$

for $\mathbf{a}, \mathbf{s} \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q^n$.

Proof. If the vector $\mathbf{a} \in \mathbb{Z}_q^n$ has at least one component that is a multiplicative unit in \mathbb{Z}_q , then since \mathbf{s} is sampled uniformly at random from \mathbb{Z}_q^n , we have $\langle \mathbf{a}, \mathbf{s} \rangle = d$ with probability $1/q$. By Fact 3.2, the probability that all components of \mathbf{a} is *not* a unit in \mathbb{Z}_q is bounded by the probability

$$\left(1 - \frac{1}{c \cdot \log \log q} \right)^n$$

for some constant c . Setting $\lambda = \Omega(n/\log \log q)$, this probability is bounded by $e^{-\lambda} < 2^{-\lambda}$. The lemma follows.

3.3 Pseudorandom Functions and Key-Homomorphic PRFs

In this section, we formally define pseudorandom functions [32] and key-homomorphic PRFs [17].

Definition 3.4 (Pseudorandom Functions [32]). *A pseudorandom function for a key space \mathcal{K} , domain \mathcal{X} , and range \mathcal{Y} is an efficiently computable deterministic function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that for any efficient adversary \mathcal{A} , we have*

$$\left| \Pr \left[k \stackrel{\text{R}}{\leftarrow} \mathcal{K} : \mathcal{A}^{F(k, \cdot)}(1^\lambda) = 1 \right] - \Pr \left[f \stackrel{\text{R}}{\leftarrow} \text{Funs}[\mathcal{X}, \mathcal{Y}] : \mathcal{A}^{f(\cdot)}(1^\lambda) = 1 \right] \right| = \text{negl}(\lambda),$$

We generally refer to the experiment where the adversary \mathcal{A} is given oracle access to the real PRF $F(k, \cdot)$ as the *real* PRF experiment. Analogously, we refer to the experiment where the adversary \mathcal{A} is given oracle access to a truly random function $f(\cdot)$ as the *ideal* PRF experiment.

Key-homomorphic PRFs are special family of pseudorandom function that satisfy an additional algebraic property. Specifically, for a key-homomorphic PRF, the key space \mathcal{K} and the range \mathcal{Y} of the PRF exhibit certain group structures such that its evaluation on any fixed input $x \in \mathcal{X}$ is homomorphic with respect to these group structures. Formally, we define a key-homomorphic PRF as follows.

Definition 3.5 (Key-Homomorphic PRFs [44, 17]). *Let (\mathcal{K}, \oplus) , (\mathcal{Y}, \otimes) be groups. Then, an efficiently computable deterministic function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ is a key-homomorphic PRF if*

- F is a secure PRF (Definition 3.4).
- For every key $k_1, k_2 \in \mathcal{K}$ and every $x \in \mathcal{X}$, we have $F(k_1, x) \otimes F(k_2, x) = F(k_1 \oplus k_2, x)$.

In this work, we will work with a slight relaxation of the notion of key-homomorphic PRFs. Namely, instead of requiring that the PRF outputs are perfectly homomorphic with respect to the PRF keys, we require that they are “almost” homomorphic in that $F(k_1, x) \otimes F(k_2, x) \approx F(k_1 \oplus k_2, x)$. Precisely, we define an almost key-homomorphic PRF as follows.

Definition 3.6 (Almost Key-Homomorphic PRFs [17]). *Let (\mathcal{K}, \oplus) , (\mathcal{Y}, \otimes) be groups and let m and p be positive integers. Then, an efficiently computable deterministic function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathbb{Z}_p^m$ is a γ -almost key-homomorphic PRF if*

- F is a secure PRF (Definition 3.4).
- For every key $k_1, k_2 \in \mathcal{K}$ and every $x \in \mathcal{X}$, there exists a vector $\mathbf{e} \in [0, \gamma]^m$ such that

$$F(k_1, x) + F(k_2, x) = F(k_1 \oplus k_2, x) + \mathbf{e} \pmod{p}.$$

Naor et al. [44] and Boneh et al. [17] gave a number of applications of (almost) key-homomorphic PRFs including distributed PRFs, symmetric-key proxy re-encryption, updatable encryption, and PRFs secure against related-key attacks.

4 Learning with Rounding and Errors

In this section, we present our new lattice-based synthesizer construction. We first define the Learning with Rounding and Errors (LWRE) problem in Section 4.1. We then show how to use the LWRE problem to construct a synthesizer in Section 4.2 and discuss its parameters in Section 4.3. We show that the LWRE problem is as hard as the standard LWE problem for suitable choices of parameters in the full version of this work.

4.1 Learning with Rounding and Errors

Definition 4.1 (Learning with Rounding and Errors). *Let $\lambda \in \mathbb{N}$ be the security parameter. The learning with rounding and errors (LWRE) problem is defined with respect to the parameters*

- LWE parameters (n, q, χ) ,
- Rounding modulus $p \in \mathbb{N}$ such that $p < q$,
- Chaining parameter $\tau \in \mathbb{N}$,

that are defined as functions of λ . Additionally, let $D_\chi : \{0, 1\}^{\lceil \log p \rceil} \rightarrow \mathbb{Z}$ be a sampling algorithm for the error distribution χ . Then, we define the $\text{LWRE}_{n, q, p, \chi, \tau}$ real and ideal oracles $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{l wre})}$ and $\mathcal{O}^{(\text{ideal})}$ as follows:

- $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{l wre})}$: *The real oracle is defined with respect to a chaining parameter $\tau \in \mathbb{N}$, and a secret matrix $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_\tau) \in \mathbb{Z}_q^{n \times \tau}$. On its invocation, the real oracle samples a vector $\mathbf{a} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^n$ and initializes $e_1 = 0$. Then, for $1 \leq i < \tau$, it iteratively computes:*

1. $r_i \leftarrow \lfloor \langle \mathbf{a}, \mathbf{s}_i \rangle + e_i \rfloor_p$.
 2. $e_{i+1} \leftarrow D_\chi(r_i)$.
- It then sets $b = \lfloor \langle \mathbf{a}, \mathbf{s}_\tau \rangle + e_\tau \rfloor_p$, and returns $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$.
- $\mathcal{O}^{(\text{ideal})}$: On its invocation, the ideal oracle samples a random vector $\mathbf{a} \xleftarrow{\text{R}} \mathbb{Z}_q^n$, a random element $u \xleftarrow{\text{R}} \mathbb{Z}_p$, and returns $(\mathbf{a}, u) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$.

The $\text{LWRE}_{n,q,p,\chi,\tau}$ problem is to distinguish the oracles $\mathcal{O}_{\mathbf{S}}^{(\text{lwre})}$ and $\mathcal{O}^{(\text{ideal})}$ for $\mathbf{S} \xleftarrow{\text{R}} \mathbb{Z}_q^{n \times \tau}$. More precisely, we define an adversary \mathcal{A} 's distinguishing advantage $\text{Adv}_{\text{LWRE}}(n, q, p, \chi, \tau, \mathcal{A})$ as the probability

$$\text{Adv}_{\text{LWRE}}(n, q, p, \chi, \tau, \mathcal{A}) = \left| \Pr [\mathcal{A}^{\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}}(1^\lambda) = 1] - \Pr [\mathcal{A}^{\mathcal{O}^{(\text{ideal})}}(1^\lambda) = 1] \right|,$$

for $\mathbf{S} \xleftarrow{\text{R}} \mathbb{Z}_q^{n \times \tau}$.

It is easy to see that when $\tau = 1$, the $\text{LWRE}_{n,q,p,\chi,\tau}$ problem is identical to the standard Learning with Rounding problem [11]. Hence, the reduction in [11] immediately shows that for $\tau = 1$, if the modulus q is sufficiently large such that $q = 2Bpn^{\omega(1)}$, then the $\text{LWRE}_{n,q,p,\chi,\tau}$ problem is as hard as the $\text{LWE}_{n,q,\chi}$ problem. We show that when τ is set to be larger, then the modulus q can be set to be significantly smaller.

Theorem 4.2. *Let λ be the security parameter and n, q, p, χ, τ be a set of parameters for the $\text{LWRE}_{n,q,p,\chi,\tau}$ problem such that p divides q . Then, for any efficient adversary \mathcal{A} making at most Q number of oracle calls, we have*

$$\text{Adv}_{\text{LWRE}}(n, q, p, \chi, \tau, \mathcal{A}) \leq Q(2Bp/q + 1/2^\lambda)^\tau + \tau \cdot \text{Adv}_{\text{LWE}}(n, q, \chi, \mathcal{A}).$$

In particular, if $Q(2Bp/q + 1/2^\lambda)^\tau = \text{negl}(\lambda)$, then the $\text{LWRE}_{n,q,p,\chi,\tau}$ problem is as hard as the $\text{LWE}_{n,q,\chi}$ problem.

We provide the proof of the theorem in the full version of this work. We provide the high-level ideas of the proof in Section 2.

Remark 4.3. The $\text{LWRE}_{n,q,p,\chi,\tau}$ problem is well-defined only when there exists a concrete sampler $D_\chi : \{0, 1\}^{\lfloor \log p \rfloor} \rightarrow \mathbb{Z}$, which uses at most $\lfloor \log p \rfloor$ random bits to sample from χ . For the discrete Gaussian distribution (over \mathbb{Z}) with Gaussian parameter $\sigma > \sqrt{n}$, there exist Gaussian samplers (i.e., [31]) that require $O(\log \lambda)$ random bits. Therefore, one can always set $p = \text{poly}(\lambda)$ to be big enough such that $\text{LWRE}_{n,q,p,\chi,\tau}$ is well defined for the discrete Gaussian distribution.

To set p to be even smaller, one can alternatively use a pseudorandom generator (PRG) to stretch the random coins that are needed by the sampler. A single element in \mathbb{Z}_p for $p = \text{poly}(\lambda)$ is not large enough to serve as a seed for a PRG. However, one can modify the oracle $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ such that it samples multiples vectors $\mathbf{a}_1, \dots, \mathbf{a}_\ell \xleftarrow{\text{R}} \mathbb{Z}_q^n$ for some $\ell = \text{poly}(\lambda)$ and then derives a vector of elements $\mathbf{r}_i = (r_{i,1}, \dots, r_{i,\ell}) \in \mathbb{Z}_p^\ell$ that can serve as a seed for any (lattice-based) pseudorandom generator.

Remark 4.4. We note that in Theorem 4.2, we impose the requirement that p perfectly divides q . This requirement is needed purely to guarantee that the rounding $\lfloor r \rfloor_p$ of a uniformly random element $r \stackrel{R}{\leftarrow} \mathbb{Z}_q$ results in a uniformly random element in \mathbb{Z}_p . However, even when q is not perfectly divisible by p (i.e. q is prime), the rounding $\lfloor r \rfloor_p$ of $r \stackrel{R}{\leftarrow} \mathbb{Z}_q$ still results in a highly unpredictable element in \mathbb{Z}_p . Therefore, by modifying the oracle $\mathcal{O}_{\tau, \mathbf{S}}^{(\text{lwre})}$ such that it applies a randomness extractor after each of the modular rounding operation, one can remove the requirement on the structure of q with respect to p .

4.2 Pseudorandom Synthesizers from LWRE

In this section, we construct our new pseudorandom synthesizer from the hardness of the $\text{LWRE}_{n,q,p,\chi,\tau}$ problem. A pseudorandom synthesizer over a domain \mathcal{D} is a function $S : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$ satisfying a specific pseudorandomness property as formulated below. We first recall the formal definition as presented in [45]. As observed in [11], one can also relax the traditional definition of a pseudorandom synthesizer by allowing the synthesizer function $S : \mathcal{D}_1 \times \mathcal{D}_1 \rightarrow \mathcal{D}_2$ to have differing domain \mathcal{D}_1 and range \mathcal{D}_2 . A synthesizer satisfying this relaxed definition still induces a PRF as long as the function can be applied iteratively. For this work, we restrict to the original definition for a simpler presentation.

Definition 4.5 (Pseudorandom Synthesizer [45]). *Let \mathcal{D} be a finite set. An efficiently computable function $S : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$ is a secure pseudorandom synthesizer if for any polynomial $\ell = \ell(\lambda)$, the following distributions are computational indistinguishable*

$$\{S(a_i, b_j)\}_{i,j \in [\ell]} \approx_c \{u_{i,j}\}_{i,j \in [\ell]},$$

where $(a_1, \dots, a_\ell) \stackrel{R}{\leftarrow} \mathcal{D}^\ell$, $(b_1, \dots, b_\ell) \stackrel{R}{\leftarrow} \mathcal{D}^\ell$, and $(u_{i,j})_{i,j \in [\ell]} \leftarrow \mathcal{D}^{\ell \times \ell}$.

Naor and Reingold [45] showed that a secure pseudorandom synthesizer induces a secure pseudorandom function. Furthermore, if the synthesizer can be computed by a low-depth circuit, then the final PRF can also be evaluated by a low-depth circuit. We formally state their result in the following theorem.

Theorem 4.6 ([45, 11]). *Suppose that there exists a pseudorandom synthesizer $S : \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{D}$ over a finite set \mathcal{D} that is computable by a circuit of size s and depth d . Then, for any $\ell = \text{poly}(\lambda)$, there exists a pseudorandom function $F : \mathcal{D}^{2\ell} \times \{0, 1\}^\ell \rightarrow \mathcal{D}$ with key space $\mathcal{D}^{2\ell}$, domain $\{0, 1\}^\ell$, and range \mathcal{D} that is computable by a circuit of size $O(s\ell)$ and depth $O(d \log \ell)$.*

As was shown in [11], the hardness of the Learning with Rounding (LWR) problem (Definition 4.1 for $\tau = 1$) naturally induces a secure pseudorandom synthesizer $S : \mathbb{Z}_q^{n \times n} \rightarrow \mathbb{Z}_q^{n \times n} \rightarrow \mathbb{Z}_p^{n \times n}$ that can be compactly defined as $S(\mathbf{S}, \mathbf{A}) = \lfloor \mathbf{S} \cdot \mathbf{A} \rfloor_p$.⁵ One can naturally extend this construction to $\text{LWRE}_{n,q,p,\chi,\tau}$

⁵The LWR synthesizer satisfies the more general definition of a pseudorandom synthesizer where the domain and range of the synthesizer can differ.

for $\tau > 1$. Unfortunately, as the $\text{LWRE}_{n,q,p,\chi,\tau}$ requires the chaining of many samples, the synthesizer does not exhibit a compact description like the LWR synthesizer. We describe the LWRE synthesizer in two steps. We first define an *LWRE function*, which satisfies the security requirement of a pseudorandom synthesizer, but does not satisfy the strict restriction on the domain and range of a synthesizer. Then, we show how to modify the LWRE function to achieve a synthesizer that satisfies Definition 4.5.

Definition 4.7 (LWRE Function). *Let λ be the security parameter and let n, q, p, χ, τ be a set of LWRE parameters. Then, we define the $\text{LWRE}_{n,q,p,\chi,\tau}$ function $\mathcal{G}_{n,q,p,\chi,\tau} : \mathbb{Z}_q^n \times \mathbb{Z}_q^{n \times \tau} \rightarrow \mathbb{Z}_p$ as follows:*

- $\mathcal{G}_{n,q,p,\chi,\tau}(\mathbf{a}, \mathbf{S})$: On input $\mathbf{a} \in \mathbb{Z}_q^n$ and $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_\tau) \in \mathbb{Z}_q^{n \times \tau}$, the LWRE function sets $e_1 = 0$ and computes for $1 \leq i < \tau$:
 1. $r_i \leftarrow \lfloor (\mathbf{a}, \mathbf{s}_i) + e_i \rfloor_p$,
 2. $e_{i+1} \leftarrow D(r_i)$.
 It then sets $b = \lfloor (\mathbf{a}, \mathbf{s}_\tau) + e_\tau \rfloor_p$ and returns $b \in \mathbb{Z}_p$.

For any $\ell = \text{poly}(\lambda)$, we can use a standard hybrid argument to show that for $\mathbf{a}_1, \dots, \mathbf{a}_\ell \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q^n$ and $\mathbf{S}_1, \dots, \mathbf{S}_\ell \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q^{n \times \tau}$, the set of elements $\{\mathcal{G}_{n,q,p,\chi,\tau}(\mathbf{a}_i, \mathbf{S}_j)\}_{i,j \in [\ell]}$ are computationally indistinguishable from ℓ^2 uniformly random elements in \mathbb{Z}_p . It readily follows that for any $\ell_1, \ell_2 = \text{poly}(\lambda)$, the function $S : \mathbb{Z}_q^{n \times \ell_1} \times (\mathbb{Z}_q^{n \times \tau})^{\ell_2} \rightarrow \mathbb{Z}_p^{\ell_1 \times \ell_2}$ that takes in as input $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_{\ell_1}) \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_q^{n \times \ell_1}$, $(\mathbf{S}_1, \dots, \mathbf{S}_{\ell_2}) \stackrel{\text{R}}{\leftarrow} (\mathbb{Z}_q^{n \times \tau})^{\ell_2}$, and returns the pairwise application of the LWRE function

$$\{\mathcal{G}_{n,q,p,\chi,\tau}(\mathbf{a}_i, \mathbf{S}_j)\}_{i \in [\ell_1], j \in [\ell_2]}$$

satisfies the security requirements for a synthesizer. Therefore, as long as ℓ_1 and ℓ_2 are set such that the cardinality of the sets $\mathbb{Z}_q^{n \times \ell_1}$, $(\mathbb{Z}_q^{n \times \tau})^{\ell_2}$, and $\mathbb{Z}_p^{\ell_1 \times \ell_2}$ have the same cardinality, the function $S : \mathbb{Z}_q^{n \times \ell_1} \times (\mathbb{Z}_q^{n \times \tau})^{\ell_2} \rightarrow \mathbb{Z}_p^{\ell_1 \times \ell_2}$ satisfies Definition 4.5. We can naturally set the parameters ℓ_1 and ℓ_2 as

$$\ell_1 = n\tau \left\lceil \frac{\log q}{\log p} \right\rceil, \quad \ell_2 = n \left\lceil \frac{\log q}{\log p} \right\rceil.$$

Formally, we define our LWRE synthesizer as follows.

Construction 4.8 (LWRE Synthesizer). Let λ be the security parameter, let n, q, p, χ, τ be a set of LWRE parameters, and let $\ell = n\tau \lceil \log q / \log p \rceil$. We define the LWRE synthesizer $S : \mathbb{Z}_q^{n \times \ell} \times \mathbb{Z}_q^{n \times \ell} \rightarrow \mathbb{Z}_q^{n \times \ell}$ as follows:

- $S(\mathbf{A}, \mathbf{S})$: On input $\mathbf{A}, \mathbf{S} \in \mathbb{Z}_q^{n \times \ell}$, the synthesizer parses the matrices
 - $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_{\ell_1})$ where $\mathbf{a}_i \in \mathbb{Z}_q^n$ and $\ell_1 = n\tau \lceil \log q / \log p \rceil$,
 - $\mathbf{S} = (\mathbf{S}_1, \dots, \mathbf{S}_{\ell_2})$ where $\mathbf{S}_j \in \mathbb{Z}_q^{n \times \tau}$ and $\ell_2 = n \lceil \log q / \log p \rceil$.

Then, the synthesizer computes the $\text{LWRE}_{n,q,p,\chi,\tau}$ function $b_{i,j} \leftarrow \mathcal{G}_{n,q,p,\chi,\tau}(\mathbf{a}_i, \mathbf{S}_j)$ for all $i \in [\ell_1]$ and $j \in [\ell_2]$. It translates the bits of $\{b_{i,j}\}_{j \in [\ell_2], i \in [\ell_1]} \in \mathbb{Z}_p^{\ell_1 \times \ell_2}$ as the representation of a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times \ell}$. It returns \mathbf{B} .

We now state the formal security statement for the LWRE synthesizer in Construction 4.8. The proof follows immediately from Definitions 4.1, 4.7, and Theorem 4.2.

Theorem 4.9 (Security). *Let λ be the security parameter and let n, q, p, χ, τ be a set of LWRE parameters. Then, assuming that the $\text{LWRE}_{n,q,p,\chi,\tau}$ problem is hard, the LWRE synthesizer in Construction 4.8 is a secure pseudorandom synthesizer (Definition 4.5).*

By combining Theorems 4.2, 4.6, and 4.9, we get the following corollary.

Corollary 4.10. *Let λ be a security parameter and $n = \text{poly}(\lambda)$ be a positive integer. Then, there exists a pseudorandom function $F : \mathcal{K} \times \{0, 1\}^{\text{poly}(\lambda)} \rightarrow \mathcal{Y}$ that can be computed by a circuit in NC^3 , and whose security can be reduced from the worst-case hardness of approximating GapSVP and SIVP to a polynomial approximation factor on n -dimensional lattices.*

We additionally discuss the parameter choices for our PRF in Section 4.3.

4.3 Parameter Instantiations

In this section, we discuss the various ways of instantiating the parameters for the LWRE synthesizer from Construction 4.8. As discussed in Section 1.2, we consider two parameter settings that provide different level of security against an adversary. For *theoretical* security, we require that an efficient adversary’s distinguishing advantage of the LWRE synthesizer to degrade super-polynomially in λ . For 2^λ -security, we require that an efficient adversary’s advantage degrades exponentially in λ . By Theorem 4.2, the main factor that we consider in determining the security of the synthesizer is the term $(2Bp/q + 1/2^\lambda)^\tau$.

Theoretical security. For theoretical security, we must let $(2Bp/q + 1/2^\lambda)^\tau = \text{negl}(\lambda)$. In one extreme, we can set $\tau = 1$ and q to be greater than $2Bp$ by a super-polynomial factor in λ , which reproduces the result of [11].⁶ As both vector-matrix multiplication and the rounding operation can be implemented in NC^1 , the resulting PRF can be implemented in NC^2 for input space $\{0, 1\}^\ell$ where $\ell = \text{poly}(\lambda)$.

If we set $\tau = \omega(1)$, then we can set q to be any function that is greater than $2Bp$ by a polynomial factor in λ . In this case, when considering the depth of the evaluation circuit, we must take into account the depth of the sampling algorithm D_χ for the error distribution χ . To base security on approximating worst-case lattice problems such as GapSVP or SIVP, we must set χ to be the discrete Gaussian distribution over \mathbb{Z} with Gaussian parameter $\sigma > \sqrt{n}$. In this case, we can either use the rejection sampling algorithm of [31] or pre-compute the samples for each possible seed for the sampler and use a look-up table. In

⁶We note that the term $1/2^\lambda$ is needed for Lemma 3.3. If q is set to be prime, then this factor can be ignored.

both cases, we can guarantee that a random element in \mathbb{Z}_p provides enough entropy to the Gaussian sampler by setting $p = \omega(\lambda^2)$.

Since the Gaussian function can be computed by an arithmetic circuit with depth $O(\log p)$, the rejection sampling algorithm can be implemented by a circuit of depth $\omega(\log \lambda \cdot \log \log \lambda)$. Therefore, the synthesizer can be evaluated by a circuit in $\text{NC}^{2+\varepsilon}$ for any constant $\varepsilon > 0$ and the final PRF can be evaluated by a circuit in $\text{NC}^{3+\varepsilon}$ for input space $\{0, 1\}^\ell$ where $\ell = \text{poly}(\lambda)$. When using a look-up table, the synthesizer can be evaluated by a circuit in $\text{NC}^{1+\varepsilon}$ for any constant $\varepsilon > 0$, and the final PRF can be evaluated in $\text{NC}^{2+\varepsilon}$.

2^λ -security. For 2^λ -security, we must let $(2Bp/q)^\tau = 1/2^{\Omega(\lambda)}$ or equivalently, $\tau \cdot \log q = \Omega(\lambda)$ for q prime. This provides a trade-off between the size of q and the chaining parameter τ , which dictates the depth needed to evaluate the PRF. In one extreme, we can set $\tau = 1$ and require the modulus to be greater than $2Bp$ by an exponential factor in λ . In the other extreme, we can let τ be linear in λ and decrease the modulus to be only a constant factor greater than $2Bp$. For practical implementations, a natural choice of parameters would be to set both τ and $\log q$ to be $\Omega(\sqrt{\lambda})$. For practical implementations, one can derive the secret keys from a single λ -bit seed using a pseudorandom generator.

Concrete instantiations The concrete parameters for our PRF can be instantiated quite flexibly depending on the applications. The modulus p can first be set to determine the output of the PRF. For instance, as the range of the PRF is \mathbb{Z}_p , the modulus p can be set to be 2^8 or 2^{16} such that the output of the PRF is byte-aligned. Then the PRF can be run multiple times (in parallel) to produce a 128-bit output.

Once p is set, the modulus q and τ can be set such that $\tau \cdot \log q \approx \lambda \cdot \log p$. For instance, to provide 2^λ -bit security, one can reasonable set q to be a 20-bit prime number and $\tau = 12$. Finally, after q is set, the LWE parameter n and noise distribution χ can be set such that the resulting LWE problem provides λ -bit level of security. Following the analysis in [5], we can set n to be around 600 (classical security) or 800 (quantum security), and χ to be either a uniform distribution over $[-24, 24]$ or an analogous Gaussian distribution with similar bits of entropy.

5 Key-Homomorphic PRFs

In this section, we show how to use the chaining method to construct key-homomorphic PRFs directly from the Learning with Errors assumption with a polynomial modulus q . Our construction is the modification of the Banerjee-Peikert (BP) PRF of [10], which generalizes the algebraic structure of previous LWE-based PRFs [11, 17].

To be precise, the BP PRF is not a single PRF family, but rather multiple PRF families that are each defined with respect to a full binary tree. The LWE parameters that govern the security of a BP PRF are determined by the structure of the corresponding binary tree. In order to construct a key-homomorphic PRF

that relies on the hardness of LWE with a polynomial modulus q , we must use the BP PRF that is defined with respect to the “right-spine” tree. However, as our modification works over any BP PRF family, we present our construction with respect to any general BP PRF. For general BP PRFs, our modification is not enough to bring the size of the modulus q to be polynomial in λ , but it still reduces its size by superpolynomial factors.

We provide our main construction in Section 5.1 and discuss its parameters in Section 5.2. We provide the proof of security and key-homomorphism in the full version.

5.1 Construction

The BP PRF construction is defined with respect to full (but not necessarily complete) binary trees. Formally, a full binary tree T is a binary tree for which every non-leaf node has two children. The shape of the full binary tree T that is used to define the PRF determines various trade-offs in the parameters and evaluation depth. As we only consider full binary trees in this work, we will implicitly refer to any tree T as a full binary tree.

Throughout the construction description and analysis, we let $|T|$ denote the number of its leaves. For any tree with $|T| \geq 1$, we let $T.\ell$ and $T.r$ denote the left and the right subtrees of T respectively (which may be empty trees). Finally, for a full binary tree T , we define its expansion factor $e(T)$ recursively as follows:

$$e(T) = \begin{cases} 0 & \text{if } |T| = 1 \\ \max\{e(T.\ell) + 1, e(T.r)\} & \text{otherwise.} \end{cases}$$

This is simply the “left-depth” of the tree, i.e., the maximum length of a root-to-leaf path, counting edges from parents to their left children.

With these notations, we define our PRF construction in a sequence of steps. We first define an input-to-matrix mapping $\mathbf{A}_T : \{0, 1\}^{|T|} \rightarrow \mathbb{Z}_q^{n \times m}$ as follows.

Definition 5.1. *Let n, q, χ be a set of LWE parameters, let $p < q$ be a rounding modulus, and let $m = n \lceil \log q \rceil$. Then, for a full binary tree T , and matrices $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$, define the function $\mathbf{A}_T : \{0, 1\}^{|T|} \rightarrow \mathbb{Z}_q^{n \times m}$ recursively:*

$$\mathbf{A}_T(x) = \begin{cases} \mathbf{A}_x & \text{if } |T| = 1 \\ \mathbf{A}_{T.\ell}(x_\ell) \cdot \mathbf{G}^{-1}(\mathbf{A}_{T.r}(x_r)) & \text{otherwise,} \end{cases},$$

where $x = x_\ell \| x_r$ for $|x_\ell| = |T.\ell|$, $|x_r| = |T.r|$.

Then, for a binary tree T , and a PRF key $\mathbf{s} \in \mathbb{Z}_q^n$, the BP PRF $F^{(\text{BP})} : \{0, 1\}^{|T|} \rightarrow \mathbb{Z}_p^m$ is defined as $F_s^{(\text{BP})}(x) = \lfloor \mathbf{s}^T \mathbf{A}_T(x) \rfloor_p$. To define our new PRF, we must first define its “noisy” variant $\mathcal{G}_{\mathbf{s}, \mathbf{E}, T} : \{0, 1\}^{|T|} \rightarrow \mathbb{Z}_p^m$ as follows.

Definition 5.2. *Let n, q, χ be a set of LWE parameters, let $p < q$ be a rounding modulus, and let $m = n \lceil \log q \rceil$. Then, for a full binary tree T , public matrices*

$\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$, error matrix $\mathbf{E} = (\mathbf{e}_1, \dots, \mathbf{e}_{e(T)}) \in \mathbb{Z}^{m \times e(T)}$, and a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$, we define the function $\mathcal{G}_{\mathbf{s}, \mathbf{E}, T} : \{0, 1\}^{|T|} \rightarrow \mathbb{Z}_p^m$:

$$\mathcal{G}_{\mathbf{s}, \mathbf{E}, T}(x) = \begin{cases} \mathbf{s}^T \mathbf{A}_T(x) + \mathbf{e}_1^T & \text{if } |T| = 1 \\ \mathcal{G}_{\mathbf{s}, \mathbf{E}, T, \ell}(x_\ell) \cdot \mathbf{G}^{-1}(\mathbf{A}_{T, r}(x_r)) + \mathbf{e}_{e(T)} & \text{otherwise,} \end{cases}$$

where $x = x_\ell \| x_r$ for $|x_\ell| = |T, \ell|$, $|x_r| = |T, r|$.

We note that when the error matrix \mathbf{E} is set to be an all-zero matrix $\mathbf{E} = \mathbf{0} \in \mathbb{Z}^{m \times e(T)}$, then the function $[\mathcal{G}_{\mathbf{s}, \mathbf{0}, T}(\cdot)]_p$ is precisely the BP PRF.

We define our new PRF to be the iterative chaining of the function $\mathcal{G}_{\mathbf{s}, \mathbf{E}, T}$. Specifically, our PRF is defined with respect to τ secret vectors $\mathbf{s}_1, \dots, \mathbf{s}_\tau \in \mathbb{Z}_q^n$ where $\tau \in \mathbb{N}$ is the chaining parameter. On input $x \in \{0, 1\}^{|T|}$, the PRF evaluation function computes $\mathbf{r}_1 \leftarrow [\mathcal{G}_{\mathbf{s}_1, \mathbf{0}, T}(x)]_p$ and uses $\mathbf{r}_1 \in \mathbb{Z}_p^m$ as a seed to derive the noise term \mathbf{E}_2 for the next iteration $\mathbf{r}_2 \leftarrow [\mathcal{G}_{\mathbf{s}_2, \mathbf{E}_2, T}(x)]_p$. The evaluation function repeats this procedure for $\tau - 1$ iterations and returns $\mathbf{r}_\tau \leftarrow [\mathcal{G}_{\mathbf{s}_{\tau-1}, \mathbf{E}_{\tau-1}, T}(x)]_p$ as the final PRF evaluation.

Construction 5.3. Let n, m, q , and χ be LWE parameters and $p < q$ be an additional rounding modulus. Then, our PRF construction is defined with respect to a full binary tree T , two public matrices $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$, a secret matrix $\mathbf{S} \in \mathbb{Z}_q^{m \times \tau}$, and a sampler $D_\chi : \{0, 1\}^{m \lceil \log p \rceil} \rightarrow \mathbb{Z}^{m \times e(T)}$ for the noise distribution χ . We define our PRF $F_{\mathbf{S}} : \{0, 1\}^{|T|} \rightarrow \mathbb{Z}_q^m$ as follows:

- $F_{\mathbf{S}}(x)$: On input $x \in \{0, 1\}^{|T|}$, the evaluation algorithm sets $\mathbf{E}_1 = \mathbf{0} \in \mathbb{Z}_q^{m \times e(T)}$. Then, for $i = 1, \dots, \tau - 1$, it iteratively computes
 1. $\mathbf{r}_i \leftarrow [\mathcal{G}_{\mathbf{s}_i, \mathbf{E}_i, T}(x)]_p$.
 2. $\mathbf{E}_{i+1} \leftarrow D_\chi(\mathbf{r}_i)$.
 It sets $\mathbf{y} \leftarrow [\mathcal{G}_{\mathbf{s}_\tau, \mathbf{E}_\tau, T}(x)]_p$, and returns $\mathbf{y} \in \mathbb{Z}_p^m$.

For the PRF to be well-defined, we must make sure that a seed $\mathbf{r}_i \leftarrow [\mathcal{G}_{\mathbf{s}_i, \mathbf{E}_i, T}(x)]_p \in \mathbb{Z}_p^m$ for $i \in [\tau - 1]$ provides enough entropy to derive the noise terms $\mathbf{E}_{i+1} \leftarrow \chi^{m \times e(T)}$. As in the case of LWRE (Remark 4.3), there are two main ways of ensuring this condition. One method is to set the rounding modulus p to be big enough such that there exists a sampler $D_\chi : \{0, 1\}^{m \lceil \log p \rceil} \rightarrow \mathbb{Z}^{m \times e(T)}$ for the noise distribution $\chi^{m \times e(T)}$. Alternatively, one can expand the seed $\mathbf{r}_i \in \mathbb{Z}_p^m$ using a pseudorandom generator to derive sufficiently many bits to sample from $\chi^{m \times e(T)}$. Since the issue of deriving the noise terms \mathbf{E}_{i+1} from the seeds \mathbf{r}_i is mostly orthogonal to the central ideas of our PRF construction, we assume that the rounding modulus p is set to be big enough such that the noise terms \mathbf{E}_{i+1} can be derived from the bits of \mathbf{r}_i .

We now state the main security theorem for the PRF in Construction 5.3.

Theorem 5.4. *Let T be any full binary tree, λ the security parameter, n, m, q, p, τ positive integers and χ a B -bounded distribution such that $m = n \lceil \log q \rceil > 2$, $(2Rmp/q)^{\tau-1} = \text{negl}(\lambda)$ for $R = |T|Bm^{e(T)}$ and p divides q . Then, assuming that the $\text{LWE}_{n, q, \chi}$ problem is hard, the PRF in Construction 5.3 is a 2-almost key-homomorphic PRF (Definition 3.5).*

We provide the proof of Theorem 5.4 in the full version. Except for the components on chaining, the proof inherits many of the arguments that are already used in [10]. For the main intuition behind the proof, we refer the readers to Section 2.3.

5.2 Instantiating the Parameters.

As in the original Banerjee-Peikert PRF [10], the size of the modulus q and the depth of the evaluation circuit is determined by the structure of the full binary tree T . The size of the modulus q is determined by the expansion factor $e(T)$ or equivalently, the length of the maximum root-to-leaf path to the left children of T . Namely, to satisfy Theorem 5.4, we require q to be large enough such that $(2Rmp/q)^{\tau-1} = \text{negl}(\lambda)$ for $R = |T|Bm^{e(T)}$.

The depth of the evaluation circuit is determined by the *sequentiality* factor $s(T)$ of the tree, which is formally defined by the recurrence

$$s(T) = \begin{cases} 0 & \text{if } |T| = 1 \\ \max\{s(T.l), s(T.r) + 1\} & \text{otherwise.} \end{cases}$$

Combinatorially, $s(T)$ denotes the length of the maximum root-to-leaf path to the right children of T . For a tree T , our PRF can be evaluated by depth $\tau \cdot s(T) \log |T|$.

When we restrict the chaining parameter to be $\tau = 1$, then we recover the parameters of the Banerjee-Peikert PRF where the modulus q is required to be $q = Rmp\lambda^{\omega(1)}$. However, setting τ to be a super-constant function $\omega(1)$, we can set $q = 2Rmp \cdot \lambda^c$ for any constant $c > 0$, which reduces the size of the modulus q by a super-polynomial factor. To guarantee that an adversary’s advantage in breaking the PRF degrades exponentially in λ , we can set τ to be linear in λ .

To set q to be polynomial in the security parameter, we can instantiate the construction with respect to the “right-spine” binary tree where the left child of any node in the tree is a leaf node. In this tree T , the sequentiality factor becomes linear in the size of the tree $s(T) = |T|$, but the expansion factor becomes a constant $e(T) = 1$. Therefore, when $\tau = \omega(1)$ (or linear in λ for 2^λ -security), the modulus q can be set to be $q = 2mp\lambda^c$ for any constant $c > 0$. The concrete parameters for our PRF can be set similarly to our synthesizer construction (see Section 4.3).

Acknowledgments

We thank the Eurocrypt reviewers for their helpful comments. This work was funded by NSF, DARPA, a grant from ONR, and the Simons Foundation. Opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA.

References

1. M. Abdalla, F. Benhamouda, and A. Passelègue. An algebraic framework for pseudorandom functions and applications to related-key security. In *CRYPTO*, 2015.

2. M. Abdalla, F. Benhamouda, and A. Passelègue. Multilinear and aggregate pseudorandom functions: New constructions and improved security. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 103–120. Springer, 2015.
3. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, 2010.
4. S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical ibe. In *CRYPTO*, 2010.
5. M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
6. J. Alperin-Sheriff and D. Apon. Dimension-preserving reductions from lwe to lwr. *IACR Cryptology ePrint Archive*, 2016(589), 2016.
7. J. Alperin-Sheriff and C. Peikert. Faster bootstrapping with polynomial error. In *CRYPTO*, 2014.
8. J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs. Learning with rounding, revisited. In *CRYPTO*. 2013.
9. B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, 2009.
10. A. Banerjee and C. Peikert. New and improved key-homomorphic pseudorandom functions. In *CRYPTO*, 2014.
11. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, 2012.
12. M. Bellare. New proofs for nmac and hmac: Security without collision-resistance. In *CRYPTO*, 2006.
13. M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In *CRYPTO*, 1996.
14. M. Bellare, R. Canetti, and H. Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *FOCS*, 1996.
15. A. Bogdanov, S. Guo, D. Masny, S. Richelson, and A. Rosen. On the hardness of learning with rounding over small modulus. In *TCC*, 2016.
16. D. Boneh, S. Kim, and H. Montgomery. Private puncturable PRFs from standard lattice assumptions. In *EUROCRYPT*, 2017.
17. D. Boneh, K. Lewi, H. Montgomery, and A. Raghunathan. Key homomorphic prfs and their applications. In *CRYPTO*. 2013.
18. D. Boneh, H. W. Montgomery, and A. Raghunathan. Algebraic pseudorandom functions with improved efficiency from the augmented cascade. In *CCS*, 2010.
19. C. Boyd, G. T. Davies, K. Gjøsteen, and Y. Jiang. Rise and shine: Fast and secure updatable encryption. *Cryptology ePrint Archive*, Report 2019/1457, 2019. <https://eprint.iacr.org/2019/1457>.
20. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, 2013.
21. Z. Brakerski, R. Tsabary, V. Vaikuntanathan, and H. Wee. Private constrained prfs (and more) from lwe. In *TCC*, 2017.
22. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, 2011.
23. Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from ring-lwe and security for key dependent messages. In *CRYPTO*, 2011.
24. Z. Brakerski and V. Vaikuntanathan. Lattice-based fhe as secure as pke. In *ITCS*, 2014.

25. Z. Brakerski and V. Vaikuntanathan. Constrained key-homomorphic PRFs from standard lattice assumptions - or: How to secretly embed a circuit in your PRF. In *TCC*, 2015.
26. R. Canetti and Y. Chen. Constraint-hiding constrained prfs for NC^1 from LWE. In *EUROCRYPT*, 2017.
27. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, 2010.
28. N. Döttling and D. Schröder. Efficient pseudorandom functions via on-the-fly adaptation. In *CRYPTO*, 2015.
29. A. Everspaugh, K. Paterson, T. Ristenpart, and S. Scott. Key rotation for authenticated encryption. In *CRYPTO*, 2017.
30. C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, 2009.
31. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
32. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.
33. S. Gorbunov and D. Vinayagamurthy. Riding on asymmetry: Efficient abe for branching programs. In *ASIACRYPT*, 2015.
34. J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
35. T. Jager, R. Kurek, and J. Pan. Simple and more efficient prfs with tight security from lwe and matrix-ddh. In *ASIACRYPT*, 2018.
36. S. Kim and D. J. Wu. Watermarking cryptographic functionalities from standard lattice assumptions. In *CRYPTO*, 2017.
37. S. Kim and D. J. Wu. Watermarking PRFs from lattices: Stronger security via extractable PRFs. In *CRYPTO*, 2019.
38. M. Klooß, A. Lehmann, and A. Rupp. (r) cca secure updatable encryption with integrity protection. In *EUROCRYPT*, 2019.
39. E. Landau. *Handbuch der Lehre von der Verteilung der Primzahlen*, volume 1. Ripol Classic Publishing House, 2000.
40. A. Lehmann and B. Tackmann. Updatable encryption with post-compromise security. In *EUROCRYPT*, 2018.
41. D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *CRYPTO*, 2011.
42. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, 2012.
43. H. Montgomery. More efficient lattice prfs from keyed pseudorandom synthesizers. In *INDOCRYPT*, 2018.
44. M. Naor, B. Pinkas, and O. Reingold. Distributed pseudo-random functions and kdcs. In *EUROCRYPT*, 1999.
45. M. Naor and O. Reingold. Synthesizers and their application to the parallel construction of pseudo-random functions. *Journal of Computer and System Sciences*, 58(2):336–375, 1999.
46. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, 2009.
47. W. Quach, D. Wichs, and G. Zirdelis. Watermarking prfs under standard assumptions: Public marking and security with extraction queries. In *TCC*, 2018.
48. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):34, 2009.