Adaptively Secure Distributed PRFs from LWE

Benoît Libert, Damien Stehlé, Radu Titiu

► To cite this version:

Benoît Libert, Damien Stehlé, Radu Titiu. Adaptively Secure Distributed PRFs from LWE. Journal of Cryptology, 2021, 34 (3), pp.1-46. 10.1007/s00145-021-09393-0 . hal-03381388

HAL Id: hal-03381388

<https://hal.inria.fr/hal-03381388>

Submitted on 16 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adaptively Secure Distributed PRFs from LWE^{*}

Benoît Libert^{1,2}, Damien Stehlé², and Radu Titiu^{2,3}

¹ CNRS, Laboratoire LIP, France

² ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France

³ Bitdefender, Bucharest, Romania

Abstract. In distributed pseudorandom functions (DPRFs), a PRF secret key SK is secret shared among N servers so that each server can locally compute a partial evaluation of the PRF on some input X . A combiner that collects t partial evaluations can then reconstruct the evaluation $F(SK, X)$ of the PRF under the initial secret key. So far, all non-interactive constructions in the standard model are based on lattice assumptions. One caveat is that they are only known to be secure in the static corruption setting, where the adversary chooses the servers to corrupt at the very beginning of the game, before any evaluation query. In this work, we construct the first fully non-interactive adaptively secure DPRF in the standard model. Our construction is proved secure under the LWE assumption against adversaries that may adaptively decide which servers they want to corrupt. We also extend our construction in order to achieve robustness against malicious adversaries.

Keywords. LWE, pseudorandom functions, distributed PRFs, threshold cryptography, adaptive security.

1 Introduction

Pseudorandom functions are fundamental objects in cryptography as most central tasks of symmetric cryptography (like secret-key encryption, message authentication or identification) can be efficiently realized from a secure PRF family. Beyond their use for cryptographic purposes, they can also be used to prove circuit lower bounds [61] and they are strongly connected to the hardness of certain tasks in learning theory [66].

Goldreich, Goldwasser and Micali (GGM) [37] showed how to build a PRF from any length-doubling pseudorandom generator (PRG). In turn, PRGs are known [42] to exist under the sole assumption that one-way functions exist. However, much more efficient constructions can be obtained by relying on specific number theoretic assumptions like the Decision Diffie-Hellman assumption [56] and related variants [30,48,18,22] or the hardness of factoring [56,57].

In the context of lattice-based cryptography, the noisy nature of hard-on-average problems, like Learning-With-Errors (LWE) [62], makes it challenging to design efficient PRF families. In order to design PRFs with small-depth evaluation circuits, several works [9,17,8] rely on the Learning-With-Rounding (LWR) technique [9], which is a “de-randomization” of LWE where noisy vectors are replaced by rounded vectors.

An appealing advantage of lattice-based techniques is that they enable the design of *key-homomorphic* PRF families [17,8]. Namely, assuming that their range and key space form an

^{*} This is the full version of a paper published in the proceedings of TCC 2018. It contains all the proofs that were omitted from the proceedings version.

additive group, for any input x and keys $\mathbf{s}, \mathbf{t}$, we have $F_{\mathbf{s}+\mathbf{t}}(x) \approx F_{\mathbf{s}}(x) + F_{\mathbf{t}}(x)$. In turn, key-homomorphic PRFs provide simple and non-interactive constructions of distributed pseudo-random functions [55]. In a (threshold) distributed PRF (DPRF), secret keys are broken into N shares $\mathbf{s}_1, \dots, \mathbf{s}_N$, each of which is given to a different server. Using its secret key share $\mathbf{s}_i$, the i -th server can locally compute a partial evaluation $F_{\mathbf{s}_i}(x)$ of the function. A dedicated server can then gather at least $t \leq N$ correct partial evaluations $F_{\mathbf{s}_{i_1}}(x), \dots, F_{\mathbf{s}_{i_t}}(x)$ and reconstruct the evaluation $F_{\mathbf{s}}(x)$ for the long-term key $\mathbf{s}$. As such, threshold PRFs inherit the usual benefits of threshold cryptography [27]. First, setting $t < N$ allows for fault-tolerant systems that can keep running when some server crashes. Second, the adversary is forced to break into t servers to compromise the security of the whole scheme. Ideally, servers should be able to generate their partial evaluations without interacting with one another.

Distributed PRFs were initially suggested by Micali and Sidney [52] and received a lot of attention since then [55, 56, 58, 29, 31]. They are motivated by the construction of distributed symmetric encryption schemes [5], distributed key distribution centers [55], or distributed coin tossing and asynchronous byzantine agreement protocols [19]. They also provide a distributed source of random coins that allows removing interaction from threshold decryption mechanisms, such as the one of Canetti and Goldwasser [21].

Boneh *et al.* [17] gave a generic construction of non-interactive DPRF from any almost key homomorphic PRF (where “almost” means that $F_{\mathbf{s}+\mathbf{t}}(x)$ only needs to be sufficiently “close” to $F_{\mathbf{s}}(x) + F_{\mathbf{t}}(x)$). Their construction, however, is only proved to be secure under *static* corruptions. Namely, the adversary has to choose the corrupted servers all-at-once and before making any evaluation query.

CONTRIBUTION. We consider the problem of proving security in the stronger *adaptive* corruption model, where the adversary chooses which servers it wants to corrupt based on the previously obtained information. In particular, an adaptive adversary is allowed to obtain partial evaluations before corrupting any server. We show that allowing a *single* partial evaluation query before the first corruption query already results in a definition which is strictly stronger than that of static security. Moreover, for a polynomial N and when $t \approx N/2$, proving adaptive security is considerably more challenging as a trivial complexity leveraging argument (i.e., guessing the set of corrupted servers upfront) makes the reduction super-polynomial, which is not satisfactory. In this work, we aim at security proofs based on reductions that do not amplify the adversary’s running time and advantage inverses more than polynomially.

In this stronger adversarial model, we provide the first realization of non-interactive distributed pseudorandom function with a security proof under a polynomial reduction. We prove the security of our construction in the standard model under the Learning-With-Errors (LWE) assumption [62] with super-polynomial approximation factors. Our security proof is given under a game-based definition (instead of a potentially stronger simulation-based one) of adaptive security. However, previous realizations [17] of non-interactive DPRFs in the standard model were analyzed under a game-based definition which is provably weaker than our definition as it only considers static corruptions.

In its basic version, our DPRF is only secure against passive adversaries. Robustness

against malicious adversaries can be achieved using leveled homomorphic signatures [40], as was suggested by earlier works on threshold lattice-based cryptography [15,16]. One caveat is that, in the adaptive corruption setting, the security proof of our robust DPRF relies on the assumption that servers can erase their internal random coins after having generated a proof that a partial PRF evaluation is correct. We note, however, that assuming reliable erasures is quite common in the context of adaptive corruptions (see, e.g., [20]). To our knowledge, we thus obtain the first DPRF candidate which is simultaneously: (i) secure under adaptive corruptions in the standard model under a well-studied assumption; (ii) robust against malicious adversaries; (iii) non-interactive (i.e., each server only sends one message to the combiner that reconstructs the final output of the PRF).

TECHNIQUES. The security proof of the generic DPRF construction of [17] stands in the static corruption model. In the adaptive corruption setting, the difficulty is that, by making N partial evaluation queries before corrupting any server, the adversary basically commits the challenger to all secret key shares. Hence, a reduction that only knows $t - 1 \approx N/2$ shares is unlikely to work as it would have to make up its mind on which set of $t - 1$ shares it wants to know at the outset of the game. In particular, this hinders a generic reduction from the security of an underlying key-homomorphic PRF. This suggests finding a reduction that knows all shares of the secret key, making it easier to consistently answer adaptive corruption queries.

To this end, we turn to lossy trapdoor functions [59], which are function families that contain both injective and lossy functions with computationally indistinguishable evaluation keys. We rely on the fact that the LWE function and its deterministic LWR variant [9] are both lossy trapdoor functions (as shown in [39,10,7]). Namely, the function that maps $\mathbf{s} \in \mathbb{Z}^n$ to $\lfloor \mathbf{A} \cdot \mathbf{s} \rfloor_p$ is injective when $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is a random matrix and becomes lossy when $\mathbf{A}$ is of the form $\bar{\mathbf{A}} \cdot \mathbf{C} + \mathbf{E}$, where $\bar{\mathbf{A}} \in \mathbb{Z}_q^{m \times n'}$, $\mathbf{C} \in \mathbb{Z}_q^{n' \times n}$ are uniformly random, $n' \ll n$ and $\mathbf{E} \in \mathbb{Z}^{m \times n}$ is a small-norm matrix. Our idea is to first construct a PRF which maps an input x to $\lfloor \mathbf{A}(x) \cdot \mathbf{s} \rfloor_p$, where $\mathbf{s} \in \mathbb{Z}^n$ is the secret key and $\mathbf{A}(x) \in \mathbb{Z}_q^{m \times n}$ is derived from public matrices. We thus evaluate a lossy trapdoor function on an input consisting of the secret key using a matrix that depends on the input. In the security proof, we use admissible hash functions [14] and techniques from fully homomorphic encryption [35] to “program” $\mathbf{A}(x)$ in such a way that, with non-negligible probability, it induces a lossy function in all evaluation queries⁴ and an injective function in the challenge phase.⁵ By choosing a large enough ratio q/p , we can make sure that evaluation queries always reveal the same information about the secret $\mathbf{s}$. Since $\lfloor \mathbf{A}(x^*) \cdot \mathbf{s} \rfloor_p$ is an injective function in the challenge phase, we can argue that it has high min-entropy, even conditionally on responses to evaluation queries. At this point, we can extract statistically uniform bits from $\lfloor \mathbf{A}(x^*) \cdot \mathbf{s} \rfloor_p$ using a deterministic randomness extractor: analogously to the deterministic encryption case [60], we need to handle a source that may be correlated with the seed.

⁴ We note that this use of lossy trapdoor functions is somewhat unusual since their injective mode is usually used to handle adversarial queries while the lossy mode comes into play in the challenge phase.

⁵ We use a “find-then-guess” security game where the adversary obtains correct evaluation for inputs of its choice before trying to distinguish a real function evaluation from a random element of the range.

We note that the above approach bears resemblance with key-homomorphic PRFs [17,8] which also evaluate functions of the form $\lfloor \mathbf{A}(x) \cdot \mathbf{s} \rfloor_p$. However, our proof method is very different in that it relies on the lossy mode of **LWE** and the homomorphic encryption scheme of [35]. The advantage of our approach is that the challenger knows the secret key $\mathbf{s}$ at all steps of the security proof. In the distributed setting, this makes it easier to handle adaptive adversaries because the reduction can always correctly answer corruption queries. In order to share the secret key $\mathbf{s}$ among N servers, we rely on the Linear Integer Secret Sharing (LISS) schemes of Damgård and Thorbek [26], which nicely fit the requirements of our security proof. Among other properties, they allow secret key shares to remain small with respect to the modulus, which helps us making sure that partial evaluations – as lossy functions of their share – always reveal the same information about uncorrupted shares. Moreover, they also enable small reconstruction constants: The secret $\mathbf{s}$ can be recovered as a linear combination of authorized shares with coefficients in $\{-1, 0, 1\}$, which is useful to avoid blowing up error terms when partial evaluations are combined together. A notable difference with [26] is that our DPRF uses a LISS scheme with Gaussian entries (instead of uniform ones), which makes it easier to analyze the remaining entropy of the key in the final step of the proof. As an additional benefit, our DPRF and its security proof are not limited to threshold access structures, but extend to access structures (e.g., monotone Boolean formulas) that can be expressed in terms of LISS.

RELATED WORK. As mentioned in [17], the early DPRF realizations [52] were only efficient when the threshold t was very small or very large with respect to the number of parties N . Before 2010, other solutions [55,56,58,29,31] either required random oracles [55] or multiple rounds of interaction [56,58,29,31]. Boneh, Lewi, Montgomery and Raghunathan [17] (BLMR) gave a generic construction of non-interactive DPRF from key-homomorphic PRFs. They also put forth the first key-homomorphic PRF in the standard model assuming the hardness of **LWE**. Banerjee and Peikert [8] generalized the BLMR construction and obtained more efficient constructions under weaker **LWE** assumptions. Boneh *et al.* [16,15] described another generic DPRF construction from a general “universal thresholdizer” tool, which allows distributing many cryptographic functionalities. So far, none of these solutions is known to provide security under adaptive corruptions.

In the context of threshold cryptography, adaptive security has been addressed in a large body of work [20,32,46,51,1,6]. These techniques, however, require interaction ⁶ and none of them is known to be compatible with existing non-interactive DPRFs. While lattice-based threshold protocols were studied by Bendlin *et al.* back in 2010 [12,13], they focused on distributing decryption operations or sharing lattice trapdoors and it is not clear how to apply them in our setting. Boneh *et al.* [16,15] showed how to generically compile cryptographic functionalities into threshold functionalities using distributed FHE. However, they do not consider adaptive corruptions and proceed by generically evaluating the circuit of the functionality at hand. While we follow their approach of using fully homomorphic signatures to acquire robustness, our basic PRF is a direct and more efficient construction.

To our knowledge, the approach of using lossy trapdoor functions to construct advanced

⁶ Except in some cases in which all players always correctly provide their contribution to the computation.

PRFs was never considered before. In spirit, our construction is somewhat similar to a random-oracle-based threshold signature proposed in [49], which also relies on the idea of always revealing the same information about the key in all evaluation queries. This DDH-based threshold signature can be turned into an adaptively secure DPRF in the random oracle model (like a variant of the Naor-Pinkas-Reingold DPRF [55]) but it has no standard-model counterpart.

After the publication of the proceedings version of this work, our adaptively secure DPRF served as a building block [50] for a multi-client functional encryption (MCFE) scheme [38] allowing to compute linear functions on inputs coming from distinct sources. We note that the DDH-based MCFE of Chotard *et al.* [23] can be interpreted as relying on an adaptively secure variant of the DPRF proposed by Naor *et al.* [55]. In both [23,50], adaptive corruptions in the MCFE setting are implicitly handled by relying on an adaptively secure DPRF.

The idea of using randomness extraction as part of the security proof of a PRF appears in [41, Section 6.2], where the function only needs to be secure in a model without evaluation queries. Here, we have to handle a different setting which prevents us from using the standard Leftover Hash Lemma.

ORGANIZATION. Section 2 recalls some relevant material about lattices, pseudorandom functions and integer secret sharing. A centralized version of our DPRF is presented in Section 3 as a warm-up. We describe its distributed variant in Section 4. In Section 5, we explain how the techniques of [16,15] apply to obtain robustness without using interaction nor random oracles.

2 Background

For any $q \geq 2$, we let $\mathbb{Z}_q$ denote the ring of integers with addition and multiplication modulo q . We always set q as a prime integer. For $2 \leq p < q$ and $x \in \mathbb{Z}_q$, we define $[x]_p := \lfloor (p/q) \cdot x \rfloor \in \mathbb{Z}_p$. This notation is readily extended to vectors over $\mathbb{Z}_p$. If $\mathbf{x}$ is a vector over $\mathbb{R}$, then $\|\mathbf{x}\|$ denotes its Euclidean norm. If $\mathbf{M}$ is a matrix over $\mathbb{R}$, then $\|\mathbf{M}\|$ denotes its induced norm. We let $\sigma_n(\mathbf{M})$ denote the least singular value of $\mathbf{M}$, where n is the rank of $\mathbf{M}$. For a finite set S , we let $U(S)$ denote the uniform distribution over S . If X is a random variable over a countable domain, the min-entropy of X is defined as $H_\infty(X) = \min_x (-\log_2 \Pr[X = x])$. If X and Y are distributions over the same domain, then $\Delta(X, Y)$ denotes their statistical distance. Let $\Sigma \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix, and $\mathbf{c} \in \mathbb{R}^n$. We define the Gaussian function on $\mathbb{R}^n$ by $\rho_{\Sigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x} - \mathbf{c})^\top \Sigma^{-1}(\mathbf{x} - \mathbf{c}))$ and if $\Sigma = \sigma^2 \cdot \mathbf{I}_n$ and $\mathbf{c} = \mathbf{0}$ we denote it by ρ_σ . For any discrete set $\Lambda \subset \mathbb{R}^n$ the discrete Gaussian distribution $D_{\Lambda, \Sigma, \mathbf{c}}$ has probability mass $\Pr_{X \sim D_{\Lambda, \Sigma, \mathbf{c}}}[X = \mathbf{x}] = \frac{\rho_{\Sigma, \mathbf{c}}(\mathbf{x})}{\rho_{\Sigma, \mathbf{c}}(\Lambda)}$, for any $\mathbf{x} \in \Lambda$. When $\mathbf{c} = \mathbf{0}$ and $\Sigma = \sigma^2 \cdot \mathbf{I}_n$ we denote it by $D_{\Lambda, \sigma}$.

Lemma 2.1 ([54, Lemma 4.4]). *For $\sigma = \omega(\sqrt{\log n})$ there is a negligible function $\varepsilon = \varepsilon(n)$ such that*

$$\Pr_{\mathbf{x} \sim D_{\mathbb{Z}^n, \sigma}} [\|\mathbf{x}\| > \sigma\sqrt{n}] \leq \frac{1 + \varepsilon}{1 - \varepsilon} \cdot 2^{-n}$$

Lemma 2.2 ([7, Lemma 2.7]). *Let p, q be positive integers such that $p < q$. Given $R > 0$ an integer, the probability that there exists $e \in [-R, R]$ such that $\lfloor y \rfloor_p \neq \lfloor y + e \rfloor_p$, when $y \leftarrow U(\mathbb{Z}_q)$, is smaller than $\frac{2Rp}{q}$.*

Lemma 2.3. *Let q be a prime and $\mathcal{M}$ be a distribution over $\mathbb{Z}_q^{m \times n}$ such that $\Delta(\mathcal{M}, U(\mathbb{Z}_q^{m \times n}))$ is less than ε . For any distribution V over $\mathbb{Z}_q^n$ we have $\Delta(\mathcal{M} \cdot V, U(\mathbb{Z}_q^m)) \leq \varepsilon + \alpha \cdot (1 - \frac{1}{q^m})$, where $\alpha := \Pr[V = \mathbf{0}]$.*

Proof. We first compute the statistical distance $\Delta(U(\mathbb{Z}_q^{m \times n}) \cdot V, U(\mathbb{Z}_q^m))$ which is equal to

$$\frac{1}{2} \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \left| \Pr_{\mathbf{M}, \mathbf{v}}[\mathbf{M}\mathbf{v} = \mathbf{y}] - \frac{1}{q^m} \right|$$

where the probabilities are taken over the choices of $\mathbf{M} \leftarrow U(\mathbb{Z}_q^{m \times n})$ and $\mathbf{v} \leftarrow V$. It is also

$$\frac{1}{2} \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \left| \Pr_{\mathbf{M}, \mathbf{v}}[\mathbf{M}\mathbf{v} = \mathbf{y} | \mathbf{v} \neq \mathbf{0}] \cdot \Pr_{\mathbf{v}}[\mathbf{v} \neq \mathbf{0}] + \Pr_{\mathbf{M}, \mathbf{v}}[\mathbf{M}\mathbf{v} = \mathbf{y} | \mathbf{v} = \mathbf{0}] \cdot \Pr_{\mathbf{v}}[\mathbf{v} = \mathbf{0}] - \frac{1}{q^m} \right|$$

Notice that when the columns of $\mathbf{M}$ are uniform in $\mathbb{Z}_q^m$ and $\mathbf{v} \neq \mathbf{0}$ then $\mathbf{M}\mathbf{v}$ is uniform in $\mathbb{Z}_q^m$ since at least one column of $\mathbf{M}$ gets multiplied by a non-zero scalar, and q is prime.

By considering the sum for all terms such that $\mathbf{y} \neq \mathbf{0}$ and then $\mathbf{y} = \mathbf{0}$, we have

$$2\Delta(U(\mathbb{Z}_q^{m \times n}) \cdot V, U(\mathbb{Z}_q^m)) = (q^m - 1) \left| \frac{1}{q^m}(1 - \alpha) - \frac{1}{q^m} \right| + \left| \frac{1}{q^m}(1 - \alpha) + \alpha - \frac{1}{q^m} \right|,$$

so that $\Delta(U(\mathbb{Z}_q^{m \times n}) \cdot V, U(\mathbb{Z}_q^m)) \leq \alpha(1 - \frac{1}{q^m})$. The claim follows from the triangle inequality

$$\Delta(\mathcal{M} \cdot V, U(\mathbb{Z}_q^m)) \leq \Delta(\mathcal{M} \cdot V, U(\mathbb{Z}_q^{m \times n}) \cdot V) + \Delta(U(\mathbb{Z}_q^{m \times n}) \cdot V, U(\mathbb{Z}_q^m)).$$

□

2.1 Lattices

For a lattice Λ , we define $\eta_\varepsilon(\Lambda)$ as the smallest $r > 0$ such that $\rho_{1/r}(\hat{\Lambda} \setminus \mathbf{0}) \leq \varepsilon$ with $\hat{\Lambda}$ denoting the dual of Λ , for any $\varepsilon \in (0, 1)$. In particular, we have $\eta_{2^{-n}}(\mathbb{Z}^n) \leq \sqrt{n}$. We define $\lambda_1^\infty(\Lambda) = \min(\|\mathbf{x}\|_\infty : \mathbf{x} \in \Lambda \setminus \mathbf{0})$.

For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we define the lattices $\Lambda^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \bmod q\}$ and $\Lambda(\mathbf{A}) = \mathbf{A}^\top \cdot \mathbb{Z}^n + q\mathbb{Z}^m$.

Lemma 2.4 ([34, Lemma 5.3]). *Let $m \geq 2n \cdot \log q$ and $q \geq 2$ prime and let $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$. With probability $\geq 1 - 2^{-\Omega(n)}$, we have $\lambda_1^\infty(\Lambda(\mathbf{A})) \geq q/4$.*

Lemma 2.5 (Adapted from [54, Lemma 4.4]). *For any n -dimensional lattice Λ , $\mathbf{x}', \mathbf{c} \in \mathbb{R}^n$ and symmetric positive definite $\Sigma \in \mathbb{R}^{n \times n}$ satisfying $\sigma_n(\sqrt{\Sigma}) \geq \eta_{2^{-n}}(\Lambda)$, we have*

$$\rho_{\Sigma, \mathbf{c}}(\Lambda + \mathbf{x}') \in [1 - 2^{-n}, 1 + 2^{-n}] \cdot \det(\Sigma)^{1/2} / \det(\Lambda).$$

Lemma 2.6. *Let $\varepsilon \in (0, 1)$, $c \in \mathbb{R}$ and $\sigma > 0$ such that $\sigma \geq \sqrt{\ln 2(1 + 1/\varepsilon)/\pi}$. Then*

$$H_\infty(D_{\mathbb{Z}, \sigma, c}) \geq \log(\sigma) - \log\left(1 + \frac{2\varepsilon}{1 - \varepsilon}\right).$$

In particular, for $\sigma = \Omega(\sqrt{n})$, we get $H_\infty(D_{\mathbb{Z}, \sigma, c}) \geq \log(\sigma) - 2^{-n}$.

Proof. From [54, Lemma 3.3] we know that $\eta_\varepsilon(\mathbb{Z}) \leq \sqrt{\ln 2(1 + 1/\varepsilon)/\pi}$. So $\sigma \geq \eta_\varepsilon(\mathbb{Z})$. By [53, Lemma 2.5], this implies that $\frac{1-\varepsilon}{1+\varepsilon} \cdot \rho_\sigma(\mathbb{Z}) \leq \rho_{\sigma, c}(\mathbb{Z})$, which translates into

$$H_\infty(D_{\mathbb{Z}, \sigma, c}) \geq H_\infty(D_{\mathbb{Z}, \sigma}) - \log\left(\frac{1 + \varepsilon}{1 - \varepsilon}\right)$$

By Poisson's summation formula, we have $\rho_\sigma(\mathbb{Z}) \geq \sigma$, so $H_\infty(D_{\mathbb{Z}, \sigma}) \geq \log \sigma$. $\square$

Definition 2.7 (LWE). *Let $m \geq n \geq 1$, $q \geq 2$ and $\alpha \in (0, 1)$ be functions of a security parameter $\lambda \in \mathbb{N}$. The LWE problem consists in distinguishing between the distributions $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{e})$ and $U(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m)$, where $\mathbf{A} \sim U(\mathbb{Z}_q^{n \times m})$, $\mathbf{s} \sim U(\mathbb{Z}_q^n)$ and $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$. For an algorithm $\mathcal{A} : \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m \rightarrow \{0, 1\}$, we define*

$$\text{Adv}_{q, m, n, \alpha}^{\text{LWE}}(\mathcal{A}) = \left| \Pr[\mathcal{A}(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{e}) = 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{u}) = 1] \right|,$$

where the probabilities are over $\mathbf{A} \sim U(\mathbb{Z}_q^{n \times m})$, $\mathbf{s} \sim U(\mathbb{Z}_q^n)$, $\mathbf{u} \sim U(\mathbb{Z}_q^m)$ and $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$ and the internal randomness of $\mathcal{A}$. We say that $\text{LWE}_{q, m, n, \alpha}$ is hard if for all ppt algorithm $\mathcal{A}$, the advantage $\text{Adv}_{q, m, n, \alpha}^{\text{LWE}}(\mathcal{A})$ is negligible.

Micciancio and Peikert [53] described a trapdoor mechanism for LWE. Their technique uses a “gadget” matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ for which anyone can publicly sample short vectors $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{G} \cdot \mathbf{x} = \mathbf{0}$. As in [53], we call a short matrix $\mathbf{R} \in \mathbb{Z}^{m \times m}$ a $\mathbf{G}$ -trapdoor for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times 2m}$ if $\mathbf{A} \cdot [\mathbf{R}^\top \mid \mathbf{I}_m]^\top = \mathbf{H} \cdot \mathbf{G}$ for some invertible matrix $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ which is referred to as the trapdoor tag. If $\mathbf{H} = \mathbf{0}$, then $\mathbf{R}$ is called a “punctured” trapdoor for $\mathbf{A}$.

Lemma 2.8 ([53, Section 5]). *Assume that $m \geq 2n \log q$. There exists a ppt algorithm GenTrap that takes as inputs matrices $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ and outputs matrices $\mathbf{R} \in \{-1, 1\}^{m \times m}$ and*

$$\mathbf{A} = [\bar{\mathbf{A}} \mid -\bar{\mathbf{A}}\mathbf{R} + \mathbf{H}\mathbf{G}] \in \mathbb{Z}_q^{n \times 2m}$$

such that if $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ is invertible, then $\mathbf{R}$ is a $\mathbf{G}$ -trapdoor for $\mathbf{A}$ with tag $\mathbf{H}$; and if $\mathbf{H} = \mathbf{0}$, then $\mathbf{R}$ is a punctured trapdoor.

Further, in case of a $\mathbf{G}$ -trapdoor, one can efficiently compute from $\mathbf{A}, \mathbf{R}$ and $\mathbf{H}$ a basis $(\mathbf{b}_i)_{i \leq 2m}$ of $\Lambda^\perp(\mathbf{A})$ such that $\max_i \|\mathbf{b}_i\| \leq O(m^{3/2})$.

Micciancio and Peikert also showed that a $\mathbf{G}$ -trapdoor for $\mathbf{A} \in \mathbb{Z}_q^{n \times 2m}$ can be used to invert the LWE function $(\mathbf{s}, \mathbf{e}) \mapsto \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}$, for any $\mathbf{s} \in \mathbb{Z}_q^n$ and any error $\mathbf{e} \in \mathbb{Z}^{2m}$ such that $\|\mathbf{e}\| < q/O(m^{3/2})$.

2.2 Admissible Hash Functions

Admissible hash functions were introduced by Boneh and Boyen [14] as a combinatorial tool for partitioning-based security proofs for which Freire *et al.* [33] gave a simplified definition. Jager [45] considered the following generalization in order to simplify the analysis of reductions under decisional assumptions.

Definition 2.9 ([45]). *Let $\ell(\lambda), L(\lambda) \in \mathbb{N}$ be functions of a security parameter $\lambda \in \mathbb{N}$. Let $\text{AHF} : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$ be an efficiently computable function. For every $K \in \{0, 1, \perp\}^L$, let the partitioning function $P_K : \{0, 1\}^\ell \rightarrow \{0, 1\}$ be defined as*

$$P_K(X) := \begin{cases} 0 & \text{if } \forall i \in [L] \quad (\text{AHF}(X)_i = K_i) \vee (K_i = \perp) \\ 1 & \text{otherwise} \end{cases}$$

*We say that AHF is a **balanced admissible hash function** if there exists an efficient algorithm $\text{AdmSmp}(1^\lambda, Q, \delta)$ that takes as input $Q \in \text{poly}(\lambda)$ and a non-negligible $\delta(\lambda) \in (0, 1]$ and outputs a key $K \in \{0, 1, \perp\}^L$ such that, for all $X^{(1)}, \dots, X^{(Q)}, X^* \in \{0, 1\}^\ell$ such that $X^* \notin \{X^{(1)}, \dots, X^{(Q)}\}$, we have*

$$\gamma_{\max}(\lambda) \geq \Pr_K [P_K(X^{(1)}) = \dots = P_K(X^{(Q)}) = 1 \wedge P_K(X^*) = 0] \geq \gamma_{\min}(\lambda),$$

where $\gamma_{\max}(\lambda)$ and $\gamma_{\min}(\lambda)$ are functions such that

$$\tau(\lambda) = \gamma_{\min}(\lambda) \cdot \delta(\lambda) - \frac{\gamma_{\max}(\lambda) - \gamma_{\min}(\lambda)}{2}$$

is a non-negligible function of λ .

Intuitively, the condition that $\tau(\lambda)$ be non-negligible requires $\gamma_{\min}(\lambda)$ to be non-negligible and the difference of $\gamma_{\max}(\lambda) - \gamma_{\min}(\lambda)$ to be small.

It is known [45] that balanced admissible hash functions exist for $\ell, L = \Theta(\lambda)$.

Theorem 2.10 ([45, Theorem 1]). *Let $(C_\ell)_{\ell \in \mathbb{N}}$ be a family of codes $C_\ell : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$ with minimal distance $c \cdot L$ for a constant $c \in (0, 1/2)$. Then, $(C_\ell)_{\ell \in \mathbb{N}}$ is a family of balanced admissible hash functions. Moreover, $\text{AdmSmp}(1^\lambda, Q, \delta)$ outputs a key $K \in \{0, 1, \perp\}^L$ for which $\eta = \lfloor \frac{\ln(2Q+Q/\delta)}{-\ln((1-c))} \rfloor$ components are not $\perp$ and*

$$\gamma_{\max} = 2^{-\eta}, \quad \gamma_{\min} = (1 - Q(1 - c))^\eta \cdot 2^{-\eta},$$

and $\tau = \gamma_{\min} \cdot \delta - \frac{\gamma_{\max} - \gamma_{\min}}{2}$ such that $\tau \geq \delta/2^{\eta+1}$.

Lemma 2.11 ([47, Lemma 8], [2, Lemma 28]). *For an input space $\mathcal{X}$, consider a mapping γ that maps a $(Q+1)$ -tuple of elements $(X^*, X_1, \dots, X_Q)$ in $\mathcal{X}$ to a probability value in $[0, 1]$. We consider the following experiment where we first execute the PRF security game, in which the adversary eventually outputs a guess $\hat{b} \in \{0, 1\}$ of the challenger's bit $b \in \{0, 1\}$ and wins with advantage ε . We denote by $X^* \in \mathcal{X}$ the challenge input and $X_1, \dots, X_Q \in \mathcal{X}$ the*

evaluation queries. At the end of the game, we flip a fair random coin $b'' \leftarrow U(\{0,1\})$. With probability $\gamma = \gamma(X^*, X_1, \dots, X_Q)$, we define $b' = \hat{b}$ and, with probability, $1 - \gamma$, we define $b' = b''$. Then, we have

$$|\Pr[b' = b] - 1/2| \geq \gamma_{\min} \cdot \varepsilon - \frac{\gamma_{\max} - \gamma_{\min}}{2},$$

where $\gamma_{\min}$ and $\gamma_{\max}$ are the maximum and minimum of $\gamma(\mathbb{X})$ for any $\mathbb{X} \in \mathcal{X}^{Q+1}$.

In our security proof, our use of admissible hash functions is inspired by the technique applied by Freire *et al.* [33] to construct programmable hash functions [43] in the multi-linear setting. Our PRF computes the LWR function for an input-dependent matrix $\mathbf{A}(x)$ obtained as a product of GSW ciphertexts [35]. Using admissible hash functions, we program the public parameters in such a way that, with noticeable probability, the gadget matrix $\mathbf{G}$ vanishes from the expression of $\mathbf{A}(x)$ in all evaluation queries, but not in the challenge phase.

2.3 (Deterministic) Randomness Extractors

A consequence of the Leftover Hash Lemma was used by Agrawal *et al.* [2] to re-randomize matrices over $\mathbb{Z}_q$ by multiplying them with small-norm matrices. We also rely on the following generalization of [2, Lemma 13].

Lemma 2.12. *Let integers m, n, ℓ such that $m > 2(n + \ell) \cdot \log q$, for some prime $q > 2$. Let $\mathbf{B}, \tilde{\mathbf{B}} \leftarrow U(\mathbb{Z}_q^{m \times \ell})$ and $\mathbf{R} \leftarrow U(\{-1, 1\}^{m \times m})$. For any matrix $\mathbf{F} \in \mathbb{Z}_q^{m \times n}$, the distributions $(\mathbf{B}, \mathbf{R} \cdot \mathbf{B}, \mathbf{R} \cdot \mathbf{F})$ and $(\mathbf{B}, \tilde{\mathbf{B}}, \mathbf{R} \cdot \mathbf{F})$ are within $2^{-\Omega(n)}$ statistical distance.*

In our security proof, we will need to extract statistically uniform bits from a high-entropy source. Here, we cannot just apply the Leftover Hash Lemma since the source may not be independent of the seed. For this reason, we will apply techniques from deterministic extraction [64, 28] and seeded extractors with seed-dependent sources [60]. In particular, we will apply a result of Dodis [28] which extends techniques due to Trevisan and Vadhan [64] to show that, for a sufficiently large $\xi > 0$, a fixed ξ -wise-independent function can be used to deterministically extract statistically uniform bits.

Lemma 2.13 ([28, Corollary 3]). *Fix any integers $\bar{n}, m, M$, any real $\varepsilon < 1$ and any collection $\mathcal{X}$ of M distributions over $\{0, 1\}^{\bar{m}}$ of min-entropy $\bar{n}$ each. Define*

$$\xi = \bar{n} + \log M, \quad \bar{k} = \bar{n} - \left(2 \log \frac{1}{\varepsilon} + \log \log M + \log \bar{n} + O(1) \right),$$

and let $\mathcal{F}$ be any family of ξ -wise independent functions from $\bar{m}$ bits to $\bar{k}$ bits. With probability at least $(1 - 1/M)$, a random function $f \leftarrow U(\mathcal{F})$ is a good deterministic extractor for the collection $\mathcal{X}$. Namely, the distribution $f(X)$ is ε -close to $U(\{0, 1\}^{\bar{k}})$ for any distribution $X \in \mathcal{X}$.

It is well-known that ξ -wise independent functions can be obtained by choosing random polynomials of degree $\xi - 1$ over $GF(2^{\bar{m}})$ (which cost $O(\xi \bar{m})$ bits to describe) and truncating their evaluations to their first $\bar{k}$ bits.

2.4 Linear Integer Secret Sharing

This section recalls the concept of linear integer secret sharing (LISS), as defined by Damgård and Thorbek [26]. The definitions below are taken from [63] where the secret to be shared lives in an interval $[-2^l, 2^l]$ centered in 0, for some $l \in \mathbb{N}$.

Definition 2.14. A **monotone access structure** on $[N]$ is a non-empty collection $\mathbb{A}$ of sets $A \subseteq [N]$ such that $\emptyset \notin \mathbb{A}$ and, for all $A \in \mathbb{A}$ and all sets B such that $A \subseteq B \subseteq [N]$, we have $B \in \mathbb{A}$. For an integer $t \in [N]$, the **threshold- t access structure** $T_{t,N}$ is the collection of sets $A \subseteq [N]$ such that $|A| \geq t$.

Let $P = [N]$ be a set of shareholders. In a LISS scheme, a dealer D wants to share a secret s in a publicly known interval $[-2^l, 2^l]$. To this end, D uses a share generating matrix $\mathbf{M} \in \mathbb{Z}^{d \times e}$ that is publicly known and a random vector $\boldsymbol{\rho} = (s, \rho_2, \dots, \rho_e)^\top$, where s is the secret to be shared $\{\rho_i\}_{i=2}^e$ are chosen uniformly in $[-2^{l_0+\lambda}, 2^{l_0+\lambda}]^e$, for a large enough $l_0 \in \mathbb{N}$. The dealer D computes a vector $\mathbf{s} = (s_1, \dots, s_d)^\top$ of share units as

$$\mathbf{s} = (s_1, \dots, s_d)^\top = \mathbf{M} \cdot \boldsymbol{\rho} \in \mathbb{Z}^d.$$

Each party in $P = \{1, \dots, N\}$ is assigned a set of share units. Letting $\psi : \{1, \dots, d\} \rightarrow P$ be a surjective function, the i -th share unit s_i is assigned to the shareholder $\psi(i) \in P$, in which case player $\psi(i)$ is said to own the i -th row of $\mathbf{M}$. If $A \subseteq P$ is a set of shareholders, $M_A \in \mathbb{Z}^{d_A \times e}$ denotes the set of rows jointly owned by A . Likewise, $\mathbf{s}_A \in \mathbb{Z}^{d_A}$ denotes the restriction of $\mathbf{s} \in \mathbb{Z}^d$ to the coordinates jointly owned by the parties in A . The j -th shareholder's share consists of $\mathbf{s}_{\psi^{-1}(j)} \in \mathbb{Z}^{d_j}$, so that it receives $d_j = |\psi^{-1}(j)|$ out of the $d = \sum_{j=1}^n d_j$ share units. The *expansion rate* $\mu = d/N$ is defined to be the average number of share units per player. Sets $A \in \mathbb{A}$ are called *qualified* and $A \notin \mathbb{A}$ are called *forbidden*.

Definition 2.15. A LISS scheme is *private* if, for any two secrets s, s' , any independent random coins $\boldsymbol{\rho} = (s, \rho_2, \dots, \rho_e)$, $\boldsymbol{\rho}' = (s', \rho'_2, \dots, \rho'_e)$ and any forbidden set A of shareholders, the distributions $\{s_i(s, \boldsymbol{\rho}) = M_i \cdot \boldsymbol{\rho} \mid i \in A\}$ and $\{s_i(s', \boldsymbol{\rho}') = M_i \cdot \boldsymbol{\rho}' \mid i \in A\}$ are $2^{-\Omega(\lambda)}$ apart in terms of statistical distance.

Damgård and Thorbek [26] showed how to construct LISS schemes from integer span programs [25].

Definition 2.16 ([25]). An *integer span program (ISP)* is a tuple $\mathcal{M} = (\mathbf{M}, \psi, \boldsymbol{\varepsilon})$, where $\mathbf{M} \in \mathbb{Z}^{d \times e}$ is an integer matrix whose rows are labeled by a surjective function $\psi : \{1, \dots, d\} \rightarrow \{1, \dots, N\}$ and $\boldsymbol{\varepsilon} = (1, 0, \dots, 0)$ is called *target vector*. The *size* of $\mathcal{M}$ is the number of rows d in $\mathbf{M}$.

Definition 2.17. Let Γ be a monotone access structure and let $\mathcal{M} = (\mathbf{M}, \psi, \boldsymbol{\varepsilon})$ an integer span program. Then, $\mathcal{M}$ is an *ISP for Γ* if it computes Γ : namely, for all $A \subseteq \{1, \dots, N\}$, the following conditions hold:

1. If $A \in \Gamma$, there exists a reconstruction vector $\boldsymbol{\lambda} \in \mathbb{Z}^{d_A}$ such that $\boldsymbol{\lambda}^\top \cdot \mathbf{M}_A = \boldsymbol{\varepsilon}^\top$.

2. If $A \notin \Gamma$, there exists $\boldsymbol{\kappa} = (\kappa_1, \dots, \kappa_e)^\top \in \mathbb{Z}^e$ such that $\mathbf{M}_A \cdot \boldsymbol{\kappa} = \mathbf{0} \in \mathbb{Z}^d$ and $\boldsymbol{\kappa}^\top \cdot \boldsymbol{\varepsilon} = 1$ (i.e., $\kappa_1 = 1$). In this case, the vector $\boldsymbol{\kappa}$ is called a sweeping vector for A .

We also define $\kappa_{\max} = \max\{|a| \mid a \text{ is an entry in some sweeping vector}\}$.

Damgård and Thorbek showed [26] that, if we have an ISP $\mathcal{M} = (\mathbf{M}, \psi, \boldsymbol{\varepsilon})$ that computes the access structure Γ , a statistically private LISS scheme for Γ can be obtained by using $\mathbf{M}$ as the share generating matrix and setting $l_0 = l + \lceil \log_2(\kappa_{\max}(e-1)) \rceil + 1$, where l is the length of the secret. A qualified set of shareholders, $A \in \mathbb{A}$, can reconstruct the initial secret $s \in [-2^\ell, 2^\ell]$ from the shares $\mathbf{s}_A \in \mathbb{Z}^{d_A}$ that are jointly owned by A , as follows. They compute a reconstruction vector $\boldsymbol{\lambda} \in \mathbb{Z}^{d_A}$, whose existence is guaranteed by the integer span program. Then the secret is obtained as $\boldsymbol{\lambda}^\top \cdot \mathbf{s}_A = \boldsymbol{\lambda}^\top \cdot \mathbf{M}_A \cdot \boldsymbol{\rho} = \boldsymbol{\varepsilon}^\top \cdot \boldsymbol{\rho} = s$.

A LISS scheme $\mathcal{L} = (\mathcal{M} = (\mathbf{M}, \psi, \boldsymbol{\varepsilon}), \Gamma, \mathcal{R}, \mathcal{K})$ is thus specified by an ISP for the access structure Γ , a space $\mathcal{R}$ of reconstruction vectors satisfying Condition 1 of Definition 2.17, and a space $\mathcal{K}$ of sweeping vectors satisfying Condition 2.

Lemma 2.18 ([63, Lemma 3.1]). *Let $l_0 = l + \lceil \log_2(\kappa_{\max}(e-1)) \rceil + 1$. If $s \in [-2^l, 2^l]$ is the secret to be shared and $\boldsymbol{\rho}$ is randomly sampled from $[-2^{l_0+\lambda}, 2^{l_0+\lambda}]^e$ conditionally on $\langle \boldsymbol{\rho}, \boldsymbol{\varepsilon} \rangle = s$, the LISS scheme derived from $\mathcal{M}$ is private. For any arbitrary $s, s' \in [-2^l, 2^l]$ and any forbidden set of shareholders $A \subset [N]$, the two distributions $\{s_A = M_A \cdot \boldsymbol{\rho} \mid \boldsymbol{\rho} \leftarrow U([-2^{l_0+\lambda}, 2^{l_0+\lambda}]^e) \text{ s.t. } \langle \boldsymbol{\rho}, \boldsymbol{\varepsilon} \rangle = s\}$, and $\{s'_A = M_A \cdot \boldsymbol{\rho} \mid \boldsymbol{\rho} \leftarrow U([-2^{l_0+\lambda}, 2^{l_0+\lambda}]^e) \text{ s.t. } \langle \boldsymbol{\rho}, \boldsymbol{\varepsilon} \rangle = s'\}$ are within statistical distance $2^{-\lambda}$.*

In the following, we do not rely on the result of Lemma 2.18 as we will share vectors sampled from Gaussian (instead of uniform) distributions using Gaussian random coins. We also depart from Lemma 2.18 in that the random coins $(\rho_2, \dots, \rho_e)$ are not sampled from a wider distribution than the secret: the standard deviation of $(\rho_2, \dots, \rho_e)$ will be the same as that of s . While this choice does not guarantee the LISS to be private in general, we will show that it suffices in our setting because we only need the secret to have sufficient min-entropy conditionally on the shares observed by the adversary. Aside from the distribution of secrets and random coins, we rely on the technique of Damgård and Thorbek [26] for building share generating matrices.

It was shown in [26] that LISS schemes can be obtained from [11, 25]. While the Benaloh-Leichter (BL) secret sharing [11] was initially designed to work over finite groups, Damgård and Thorbek generalized it [26] so as to share integers using access structures consisting of any monotone Boolean formula. In turn, this implies a LISS scheme for any threshold access structure by applying a result of Valiant [65, 36]. Their LISS scheme built upon Benaloh-Leichter [11] comes in handy for our purposes because the reconstruction coefficients and the sweeping vectors are small: as can be observed from [26, Lemma 4], the entries of $\boldsymbol{\lambda}$ live in $\{-1, 0, 1\}$ and [26, Lemma 5] shows that $\kappa_{\max} = 1$. For a monotone Boolean f , the BL-based technique allows binary share distribution matrices $M \in \{0, 1\}^{d \times e}$ such that $d, e = O(\text{size}(f))$ and which have at most $\text{depth}(f) + 1$ non-zero entries, so that each share unit s_i has magnitude $O(2^{l_0+\lambda} \cdot \text{depth}(f))$.

Valiant's result [65] implies the existence of a monotone Boolean formula of the threshold- t function $T_{t,N}$, which has size $d = O(N^{5.3})$ and depth $O(\log N)$. Since each player receives

d/N rows of $\mathbf{M}$ on average, the average share size is thus $O(N^{4.3} \cdot (l_0 + \lambda + \log \log N))$ bits. Valiant's construction was improved by Hoory *et al.* [44] who gave a monotone formula of size $O(N^{1+\sqrt{2}})$ and depth $O(\log N)$ for the majority function.⁷ This reduces the average share size to $O(N^{\sqrt{2}} \cdot (l_0 + \lambda + \log \log N))$ bits.

2.5 (Distributed) Pseudorandom Functions

A pseudorandom function family is specified by efficient algorithms (**Keygen**, **Eval**), where **Keygen** a randomized key generation algorithm that takes in a security parameter 1^λ and outputs a random key $K \leftarrow \mathcal{K}$ from a key space $\mathcal{K}$. **Eval** is a deterministic evaluation algorithm, which takes in a key $K \in \mathcal{K}$ and an input X in a domain $\mathcal{D} = \{0, 1\}^\ell$ and evaluates a function $F(K, X)$ in a range $\mathcal{R} = \{0, 1\}^\mu$. The standard security definitions for PRFs are recalled below.

Definition 2.19 (Real-or-random security). *Let λ be a security parameter and let $\kappa = \kappa(\lambda)$. A pseudorandom function $F : \mathcal{K} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\mu$ is an efficiently computable function where the first input $K \in \mathcal{K}$ is the key. Let Ω be the set of all functions that map ℓ -bit inputs to μ -bit strings. The advantage of a PRF distinguisher $\mathcal{A}$ making Q evaluation queries is defined as*

$$\text{Adv}_Q^{\mathcal{A}, \text{prf}}(\lambda) := |\Pr[\mathcal{A}^{F(K, \cdot)}(1^\lambda) = 1 \mid K \leftarrow \mathcal{K}] - \Pr[\mathcal{A}^{R(\cdot)}(1^\lambda) = 1 \mid R \leftarrow U(\Omega)]|,$$

where the probability is taken over all coin tosses. The pseudorandom function F is called secure if no ppt adversary $\mathcal{A}$ has noticeable advantage $\text{Adv}_Q^{\mathcal{A}, \text{prf}}(\lambda)$.

In some cases, it is convenient to work with the following definition.

Definition 2.20 (Find-then-guess security). *Let λ be a security parameter. A function $F : \mathcal{K} \times \{0, 1\}^\ell \rightarrow \{0, 1\}^\mu$ is a secure PRF if no PPT adversary has noticeable advantage in the following game:*

Init. *The challenger initially chooses a random key $K \leftarrow \mathcal{K}$ and also flips a random coin $b \leftarrow U(\{0, 1\})$.*

Queries 1. *On polynomially-many occasions, the adversary $\mathcal{A}$ chooses an arbitrary input $X \in \{0, 1\}^\ell$ and the challenger returns $F(K, X)$. These queries may be adaptive, i.e. depend on responses to previous queries.*

Challenge. *The adversary chooses an input X^* that differs from all queries of the previous stage. The challenger computes $Y_1 = F(K, X^*)$ and randomly samples $Y_0 \leftarrow \{0, 1\}^\mu$. Then, it returns Y_b to the adversary.*

Queries 2. *The adversary adaptively makes another series of queries for arbitrary inputs $X \neq X^*$.*

Guess. *The adversary outputs a bit $\hat{b} \in \{0, 1\}$ and wins if $\hat{b} = b$. Its advantage is defined as*

$$\text{Adv}_A^{\text{FG}}(\lambda) := |\Pr[\hat{b} = b] - 1/2| = \frac{1}{2} \cdot |\Pr[\hat{b} = 1 \mid b = 1] - \Pr[\hat{b} = 1 \mid b = 0]|.$$

⁷ Note that a threshold- t function can be obtained from the majority function by fixing the desired number of input bits, so that we need a majority function of size $\leq 2N$ to construct a threshold function $T_{t,N}$.

A standard hybrid argument over evaluation queries shows that – up to a multiplicative gap $\Theta(Q)$ between the advantage functions, where Q is the number of evaluation queries – find-then-guess security is equivalent to real-or-random security. In the following, we will work with Definition 2.20 .

Definition 2.21 (Distributed Pseudorandom Functions). *A distributed pseudorandom function (DPRF) is a tuple of (possibly randomized) efficient algorithms (Setup, Share, PEval, Eval, Combine) with the following specification.*

Setup($1^\lambda, 1^\ell, 1^t, 1^N$): *Takes as input a security parameter 1^λ , a number of servers 1^N , a threshold 1^t and a desired input length 1^ℓ and outputs public parameters pp . It also specifies the key space $\mathcal{K}$, the domain $\mathcal{D} := \{0, 1\}^\ell$ and the range $\mathcal{R}$ of the evaluation function.*

Share(pp, SK_0): *The key sharing algorithm inputs a random master secret key $SK_0 \in \mathcal{K}$ and outputs a tuple of shares $(SK_1, \dots, SK_N) \in \mathcal{K}^N$, which form a (t, N) -threshold secret sharing of SK_0 .*

PEval(pp, SK_i, X): *The partial evaluation algorithm takes as input the public parameters, a share of the secret key and an input $X \in \mathcal{D}$. It outputs a partial evaluation $Y_i \in \mathcal{R}$.*

Eval(pp, SK_0, X): *Takes as input the public parameters, the master secret key and input $X \in \mathcal{D}$. It outputs the evaluation $Y \in \mathcal{R}$.*

Combine($\mathcal{S}, (Y_{i_1}, Y_{i_2}, \dots, Y_{i_t})$): *Takes as input a t -subset $\mathcal{S} \subset [N]$, a set $\{Y_i\}_{i \in \mathcal{S}}$ of t partial evaluations and returns an evaluation $Y \in \mathcal{R}$.*

CONSISTENCY. We say that a DPRF is consistent if, for any $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$, any master key $SK_0 \leftarrow \mathcal{K}$ shared according to $(SK_1, \dots, SK_N) \leftarrow \text{Share}(SK_0)$, any t -subset $\mathcal{S} = \{i_1, \dots, i_t\} \subset [N]$ and any input $X \in \mathcal{D}$, if $Y_{i_j} = \text{PEval}(SK_{i_j}, X)$ for each $j \in [t]$, then we have $\text{Eval}(SK_0, X) = \text{Combine}(\mathcal{S}, (Y_{i_1}, \dots, Y_{i_t}))$ with overwhelming probability over the random coins of Setup and Share.

We say that a DPRF provides adaptive security if it remains secure against an adversary that can adaptively choose which servers it wants to corrupt. In particular, the adversary can arbitrarily interleave evaluation and corruption queries as long as they do not allow it to trivially win.

Definition 2.22 (Adaptive DPRF security). *Let λ be a security parameter and let integers $t, N \in \text{poly}(\lambda)$. We say that a (t, N) -DPRF is pseudorandom under adaptive corruptions if no PPT adversary has non-negligible advantage in the following game:*

1. *The challenger generates $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$ and chooses a random key $SK_0 \leftarrow \mathcal{K}$, which is broken into N shares $(SK_1, \dots, SK_N) \leftarrow \text{Share}(SK_0)$. It also initializes an empty set $\mathcal{C} \leftarrow \emptyset$ and flips a random coin $b \leftarrow U(\{0, 1\})$.*
2. *The adversary $\mathcal{A}$ adaptively interleaves the following kinds of queries.*

Corruption: *The adversary $\mathcal{A}$ chooses an index $i \in [N] \setminus \mathcal{C}$. The challenger returns SK_i to $\mathcal{A}$ and sets $\mathcal{C} := \mathcal{C} \cup \{i\}$.*

Evaluation: *The adversary $\mathcal{A}$ chooses a pair $(i, X) \in [N] \times \mathcal{D}$ and the challenger returns $Y_i = \text{PEval}(SK_i, X)$.*

3. The adversary chooses an input X^* . At this point, the challenger randomly samples $Y_0 \leftarrow U(\{0, 1\}^\mu)$ and computes $Y_1 = \text{Eval}(SK_0, X^*)$. Then, it returns Y_b to the adversary.
4. The adversary $\mathcal{A}$ adaptively makes more queries as in Stage 2 under the restriction that, at any time, we should have $|\mathcal{C} \cup \mathcal{E}| < t$, where $\mathcal{E} \subset [N]$ denotes the set of indexes for which an evaluation query of the form (i, X^*) was made in Stage 2 or in Stage 4.
5. The adversary $\mathcal{A}$ outputs a bit $\hat{b} \in \{0, 1\}$ and wins if $\hat{b} = b$. Its advantage is defined to be

$$\text{Adv}_{\mathcal{A}}^{\text{DPRF}}(\lambda) := |\Pr[\hat{b} = b] - 1/2|.$$

Definition 2.22 is a game based definition, which may not imply security in the sense of simulation-based definitions. Still, we show it is strictly stronger than the definition of static security used in [17] (which is recalled in Definition 2.23 below). It is well-known that static security does not imply adaptive security in distributed threshold protocols (see, e.g., [24]). In the case of DPRFs, we show that allowing even a *single* evaluation query before any corruption query already gives a stronger game-based definition than Definition 2.23.

Definition 2.23 (Static DPRF security). Let λ be a security parameter and let integers $t, N \in \text{poly}(\lambda)$. A (t, N) -DPRF is pseudorandom under static corruptions if no PPT adversary has non-negligible advantage in the following game:

1. The challenger generates $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$ and chooses a random key $SK_0 \leftarrow \mathcal{K}$, which is broken into N shares $(SK_1, \dots, SK_N) \leftarrow \text{Share}(SK_0)$. It also initializes empty sets $\mathcal{C}, \mathcal{V} \leftarrow \emptyset$ and flip a random coin $b \leftarrow U(\{0, 1\})$.
2. The adversary $\mathcal{A}$ chooses a set $S^* = \{i_1, \dots, i_{t-1}\}$ and the challenger returns the secret key shares $\{SK_{i_1}, \dots, SK_{i_{t-1}}\}$.
3. The adversary $\mathcal{A}$ adaptively interleaves the following kinds of queries.

Evaluation: $\mathcal{A}$ chooses an input $X \in \mathcal{D}$. The challenger replies by returning partial evaluations $\{Y_i = \text{PEval}(SK_i, X)\}_{i \in [N] \setminus S^*}$ and updating $\mathcal{V} := \mathcal{V} \cup \{X\}$.

Challenge: $\mathcal{A}$ chooses an input $X \in \mathcal{D}$. If X previously occurred in a challenge query, the challenger returns the same output as before. Otherwise, it randomly chooses $Y_{X,0} \leftarrow U(\{0, 1\}^\mu)$ and computes $Y_{X,1} = \text{Eval}(SK_0, X)$. It returns $Y_{X,b}$ and updates $\mathcal{C} := \mathcal{C} \cup \{X\}$.

It is required that $\mathcal{C} \cap \mathcal{V} = \emptyset$ at any time.

4. The adversary $\mathcal{A}$ outputs a bit $\hat{b} \in \{0, 1\}$ and wins if $\hat{b} = b$. Its advantage is defined to be $\text{Adv}_{\mathcal{A}}^{\text{DPRF}}(\lambda) := |\Pr[\hat{b} = b] - 1/2|$.

We may assume w.l.o.g. that the adversary only makes one challenge query in the experiment of Definition 2.23. Indeed, a standard hybrid argument allows showing that security in the single-challenge sense implies security when polynomially-many queries are allowed.

Theorem 2.24. For any $t, N \in \text{poly}(\lambda)$ such that $\binom{N}{t-1}$ is super-polynomial, there exists a DPRF family which is secure in the sense of Definition 2.23 but insecure in the sense of Definition 2.22.

Proof. Let $\Pi = (\text{Setup}, \text{Share}, \text{PEval}, \text{Eval}, \text{Combine})$ be a (t, N) -DPRF family which provides static security as captured by Definition 2.23. We assume that secret keys live in a large finite field $\mathbb{F}_p$ (the key-homomorphic-based constructions implied by [17,8] can be modified to satisfy this condition). We modify the DPRF family Π into $\Pi^* = (\text{Setup}^*, \text{Share}^*, \text{PEval}^*, \text{Eval}^*, \text{Combine}^*)$ which is insecure in the experiment of Definition 2.22 but remains secure under Definition 2.23 for the same threshold t . The construction Π^* goes as follows.

Setup $^*(1^\lambda, 1^\ell, 1^t, 1^N)$: Run $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$ and output pp .

Share $^*(\text{pp}, SK_0)$: Given $SK_0 = k$, conduct the following steps.

1. Run $(k_1, \dots, k_N) \leftarrow \text{Share}(\text{pp}, k)$.
2. Generate an independent $(t-1, N)$ Shamir secret sharing of k and let $(k'_1, \dots, k'_N)$ be the resulting shares.
3. Choose a random subset $\mathcal{T} \subset [N]$ of cardinality $|\mathcal{T}| = N - t + 1$. Choose uniformly random vector $(k''_1, \dots, k''_N) \leftarrow U(\mathbb{F}_p^N)$ and define $(\tilde{k}_1, \dots, \tilde{k}_N)$ in the following way

$$\tilde{k}_i := \begin{cases} k'_i & \text{if } i \notin \mathcal{T} \\ k''_i & \text{if } i \in \mathcal{T} \end{cases}$$

For each $i \in [N]$ define the i -th secret key share as $SK_i := (k_i, \tilde{k}_i, \mathcal{T})$.

Output $(SK_1, \dots, SK_N)$.

PEval $^*(\text{pp}, SK_i, X)$: Given the secret key share $SK_i = (k_i, \tilde{k}_i, \mathcal{T})$, compute the partial evaluation $F_{k_i}(X) = \text{PEval}(\text{pp}, k_i, X)$ and output $Y_i = (F_{k_i}(X), \mathcal{T})$.

Eval $^*(\text{pp}, SK_0, X)$: Compute $Y \leftarrow \text{Eval}(\text{pp}, k, X)$ and output Y .

Combine $^*(\mathcal{S}, (Y_{i_1}, \dots, Y_{i_t}))$: For each $\kappa \in [t]$, parse Y_{i_κ} as a pair $Y_{i_\kappa} = (\tilde{Y}_{i_\kappa}, \mathcal{T})$, where $\tilde{Y}_{i_\kappa} = F_{k_{i_\kappa}}(X)$. Run $Y \leftarrow \text{Combine}(\mathcal{S}, (\tilde{Y}_{i_1}, \dots, \tilde{Y}_{i_t}))$ and return Y .

It is easy to see that Π^* is trivially insecure in the sense of Definition 2.22. Before corrupting any server, the adversary can make a single evaluation query on an arbitrary server to learn $\mathcal{T}$. By corrupting all servers $i \in [N] \setminus \mathcal{T}$, the adversary can then obtain $t-1$ shares of $SK_0 = k$ so as to reconstruct it.

The proof that Π^* remains statically secure relies on the idea that, in the game of Definition 2.23, the adversary has to choose the set $\mathcal{S}^*$ of corrupted servers *before* making any evaluation query. Hence, with overwhelming probability $1 - 1/\binom{N}{t-1}$, we have $\mathcal{S}^* \neq [N] \setminus \mathcal{T}$, in which case $\{\tilde{k}_i\}_{i \in \mathcal{S}^*}$ is information-theoretically independent of SK_0 . We give a reduction below.

Assuming that we have a DPRF adversary $\mathcal{B}$ against Π^* , we build a DPRF adversary $\mathcal{B}_0$ against Π as follows. First, $\mathcal{B}_0$ chooses a random subset $\mathcal{T} \subset [N]$ of cardinality $|\mathcal{T}| = N - t + 1$. Then, $\mathcal{B}_0$ starts interacting with $\mathcal{B}$ that chooses an arbitrary $(t-1)$ -subset $\mathcal{S}^* = \{i_1^*, \dots, i_{t-1}^*\} \in [N]$. If $\mathcal{S}^* = [N] \setminus \mathcal{T}$, $\mathcal{B}_0$ halts and considers $\mathcal{B}$ as successful. Otherwise, we have $|\mathcal{S}^* \cap \mathcal{T}| \geq 1$, which means that at least one of the $\{\tilde{k}_{i_\kappa}\}_{\kappa=1}^{t-1}$ is supposed to have been obtained by replacing a coordinate of $(k'_{i_1}, \dots, k'_{i_{t-1}})$ with a random element of $\mathbb{F}_p$.

Since $\{k'_{i_\kappa}\}_{\kappa=1}^{t-1}$ form a Shamir secret sharing over a finite field, they live on a degree- $(t-2)$ polynomial which is also a $(t-1)$ -wise independent function. Hence, replacing at least one of the shares $\{k'_{i_\kappa}\}_{\kappa=1}^{t-1}$ with a random value results in a randomly distributed $(\tilde{k}_{i_1^*}, \dots, \tilde{k}_{i_{t-1}^*})$ which is uncorrelated to k_0 .

In this case, $\mathcal{B}'$ sends $\mathcal{S}^*$ to its own DPRF challenger and obtains the key share $k_{i_1^*}, \dots, k_{i_{t-1}^*}$. It also chooses $(\tilde{k}_{i_1^*}, \dots, \tilde{k}_{i_{t-1}^*}) \leftarrow U(\mathbb{F}_p^{t-1})$ uniformly and gives $\{(k_{i_\kappa^*}, \tilde{k}_{i_\kappa^*}, \mathcal{T})\}_{\kappa=1}^{t-1}$ to $\mathcal{B}$. Then, $\mathcal{B}_0$ answers $\mathcal{B}$'s evaluation queries by invoking its own challenger on the same inputs and constructing the responses in the obvious way. At the end of the experiment, $\mathcal{B}_0$ outputs whatever $\mathcal{B}$ outputs.

If ε denotes the advantage of $\mathcal{B}$ and **Guess** denotes the event that $\mathcal{S}^* = [N] \setminus \mathcal{T}$, we know that $\Pr[\mathcal{B}_0 \text{ wins} | \neg \text{Guess}] = \frac{1}{2} + \varepsilon$. This implies

$$\begin{aligned} \Pr[\mathcal{B}_0 \text{ wins}] &= \Pr[\mathcal{B}_0 \text{ wins} \mid \text{Guess}] \cdot \Pr[\text{Guess}] + \Pr[\mathcal{B}_0 \text{ wins} \mid \neg \text{Guess}] \cdot \Pr[\neg \text{Guess}] \\ &> \Pr[\mathcal{B}_0 \text{ wins} \mid \neg \text{Guess}] \cdot \Pr[\neg \text{Guess}] \\ &= \left(\frac{1}{2} + \varepsilon\right) \cdot (1 - \Pr[\text{Guess}]) \\ &> \left(\frac{1}{2} + \varepsilon\right) \cdot \left(1 - 1/\binom{N}{t-1}\right) > \left(\frac{1}{2} + \varepsilon\right) - 1/\binom{N}{t-1}. \end{aligned}$$

Consequently, if $\varepsilon > 0$ is noticeable, so is $|\Pr[\mathcal{B}_0 \text{ wins}] - \frac{1}{2}| > \varepsilon - 2^{-\Omega(\lambda)}$. $\square$

The above separation holds for small non-constant values of t and N if we assume polynomial or slightly super-polynomial adversaries. Indeed, if $t = \Theta(\log \lambda)$ and $N = \Theta(\log^2 \lambda)$, for example, the adversary's probability to guess $\mathcal{T}$ is $1/\binom{N}{t-1} < \left(\frac{t-1}{N-t}\right)^{t-1}$, which is roughly $1/(\log \lambda)^{\log \lambda} = \lambda^{-\omega(1)}$. For any constant $c > 0$, if $t = \Theta(\lambda^{1/c})$ and $N = \Theta(\lambda^{2/c})$, the guessing probability becomes sub-exponentially small as well.

3 A Variant of the BLMR PRF

Before describing our distributed PRF, we present its centralized version which can be seen as a variant of the key-homomorphic PRFs described by Boneh *et al.* [17] and Banerjee-Peikert PRFs [8]. However, the security proof is very different in that it does not use a hybrid argument over the input bits. Instead, it applies the strategy of partitioning the input space into disjoint subspaces (analogously to proof techniques for, e.g., identity-based encryption [67]) and builds on the lossy mode of LWE [39].

In [17,8], a PRF evaluation of an input x is of the form $\mathbf{y} = \lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p \in \mathbb{Z}_p^m$, where $\mathbf{s} \in \mathbb{Z}_q^n$ is the secret key and $\mathbf{A}(X) \in \mathbb{Z}_q^{n \times m}$ is an input-dependent matrix obtained from public matrices $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$. Our variant is similar at a high level, with two differences. First, we derive $\mathbf{A}(x)$ from a set of $2L$ public matrices $\{\mathbf{A}_{i,0}, \mathbf{A}_{i,1}\}_{i=1}^L$. Second, $\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p$ is not quite our PRF evaluation. Instead, we obtain the PRF value by using $\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p$ as a source of entropy for a deterministic randomness extractor.

The security proof departs from [17,8] by exploiting the connection between the schemes

and the Gentry-Sahai-Waters FHE [35]. For each $i \in [L]$ and $b \in \{0, 1\}$, we interpret the matrix $\mathbf{A}_{i,b} \in \mathbb{Z}_q^{n \times m}$ as a GSW ciphertext $\mathbf{A}_{i,b} = \mathbf{A} \cdot \mathbf{R}_{i,b} + \mu_{i,b} \cdot \mathbf{G}$, where $\mathbf{R}_{i,b} \in \{-1, 1\}^{m \times m}$, $\mu_{i,b} \in \{0, 1\}$ and $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ is the gadget matrix of [53]. Before evaluating the PRF on an input X , we encode $X \in \{0, 1\}^\ell$ into $x \in \{0, 1\}^L$ using an admissible hash function. Then, we homomorphically derive $\mathbf{A}(x)$ as a GSW ciphertext $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x + (\prod_{i=1}^L \mu_{i,x[i]}) \cdot \mathbf{G}$, for some small-norm $\mathbf{R}_x \in \mathbb{Z}^{m \times m}$. By carefully choosing $\{\mu_{i,b}\}_{i \in [L], b \in \{0,1\}}$, the properties of admissible hash functions ensure that the product $\prod_{i=1}^L \mu_{i,x[i]}$ cancels out in all evaluation queries but evaluates to 1 on the challenge input X^* .

In the next step of the proof, we move to a modified experiment where the random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is replaced by a lossy matrix $\mathbf{A}^\top = \tilde{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{E}$, where $\tilde{\mathbf{A}} \leftarrow U(\mathbb{Z}_q^{n' \times m})$, $\mathbf{C} \leftarrow U(\mathbb{Z}_q^{n' \times n})$ and $\mathbf{E} \in \mathbb{Z}^{m \times n}$ is a short Gaussian matrix. This modification has the consequence of turning $\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p$ into a lossy function of $\mathbf{s}$ on all inputs X for which $\prod_{i=1}^L \mu_{i,x[i]} = 0$. At the same time, the function remains injective whenever $\prod_{i=1}^L \mu_{i,x[i]} = 1$. Using the properties of admissible hash functions, we still have a noticeable probability that the function be lossy in all evaluation queries and injective in the challenge phase. Moreover, by using a small-norm secret $\mathbf{s} \in \mathbb{Z}^n$ and setting the ratio q/p large enough, we can actually make sure that evaluation queries always reveal the same information (namely, the product $\mathbf{C} \cdot \mathbf{s}$) about $\mathbf{s}$. As long as we have $\prod_{i=1}^L \mu_{i,x^*[i]} = 1$ for the challenge input X^* , the rounded value $\tilde{\mathbf{z}} = \lfloor \mathbf{A}(x^*)^\top \cdot \mathbf{s} \rfloor_p = \lfloor (\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G})^\top \cdot \mathbf{s} \rfloor_p$ is guaranteed to have a lot of entropy as an injective function of an unpredictable $\mathbf{s}$. At this point, we can extract statistically uniform bits from the source $\tilde{\mathbf{z}}$. Since the latter depends on x^* (which can be correlated with the seed included in public parameters), we need an extractor that can operate on seed-dependent sources. Fortunately, deterministic extractors come in handy for this purpose.

3.1 Decomposing Random Matrices into Invertible Binary Matrices

In the following, we set $\omega = n \lceil \log q \rceil$ and $m = 2\omega$ and define

$$\mathbf{G} = [\mathbf{I}_n \otimes (1, 2, 4, \dots, 2^{\lceil \log q \rceil - 1}) \mid \mathbf{I}_n \otimes (1, 2, 4, \dots, 2^{\lceil \log q \rceil - 1})] \in \mathbb{Z}_q^{n \times m}$$

which is a variant of the gadget matrix of [53]. We also define $\mathbf{G}^{-1} : \mathbb{Z}_q^{n \times m} \rightarrow \mathbb{Z}^{m \times m}$ to be the deterministic algorithm that inputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and outputs a binary matrix $\mathbf{G}^{-1}(\mathbf{A}) \in \{0, 1\}^{m \times m}$ such that $\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{A}) = \mathbf{A}$, defined as follows: given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, parse $\mathbf{A}$ as $[\mathbf{A}_L \mid \mathbf{A}_R]$ with $\mathbf{A}_L, \mathbf{A}_R \in \mathbb{Z}_q^{n \times \omega}$. We set $\mathbf{G}^{-1}(\mathbf{A}) \in \{0, 1\}^{m \times m}$ to be the binary matrix that consists of the binary decomposition $\text{BD}(\mathbf{A}_L) \in \{0, 1\}^{\omega \times \omega}$ in its upper-left corner and $\text{BD}(\mathbf{A}_R) \in \{0, 1\}^{\omega \times \omega}$ in its lower-right corner. The next lemma shows that the binary decomposition of a matrix sampled uniformly random from $\mathbb{Z}_q^{n \times \omega}$ is $\mathbb{Z}_q$ -invertible with probability at least $\exp(-3\sqrt{\exp(1)})$.

Hence, if we sample $141 \cdot \omega > \omega / \exp(-3 \cdot \sqrt{\exp(1)})$ vectors $\mathbf{a}_i \leftarrow U(\mathbb{Z}_q^n)$ and stack up the binary decompositions of $\mathbf{a}_i^\top$, the probability that the resulting matrix contains a $\mathbb{Z}_q$ -invertible sub-matrix over $\{0, 1\}^\omega$ is at least $1 - 2^{-\Omega(\omega)}$.

We can thus sample a random matrix $\mathbf{A} = [\mathbf{A}_L \mid \mathbf{A}_R] \leftarrow U(\mathbb{Z}_q^{n \times m})$ constrained by the condition $\mathbf{G}^{-1}(\mathbf{A}) \in \{0, 1\}^{m \times m}$ is invertible mod q , by sampling $\mathbf{A}_L \leftarrow U(\mathbb{Z}_q^{n \times \omega})$ and $\mathbf{A}_R \leftarrow U(\mathbb{Z}_q^{n \times \omega})$ in parallel, such that their binary decompositions are $\mathbb{Z}_q$ -invertible.

Lemma 3.1 (Adapted from [17, Lemma A.3]). *Let $\omega = n \lceil \log q \rceil$. If $q \geq 2^{\omega/n} \cdot (1 - \frac{1}{2n})$, the probability that $\text{BD}(\mathbf{A}) \in \{0, 1\}^{\omega \times \omega}$ is $\mathbb{Z}_q$ -invertible for a uniformly random matrix $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times \omega})$ is at least $\exp(-3\sqrt{\exp(1)})$.*

Proof. As in the proof of [17, Lemma A.3], we use the observation that, for any i linearly independent vectors $\mathbf{v}_1, \dots, \mathbf{v}_i \in \mathbb{Z}_q^\omega$ over $\mathbb{Z}_q$, if $V = \text{span}_{\mathbb{Z}_q}(\mathbf{v}_1, \dots, \mathbf{v}_i)$, we have $|V \cap \{0, 1\}^\omega| \leq 2^i$.

For an index $i \in [\omega - 1]$, suppose that we have $\mathbb{Z}_q$ -independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_i \in \{0, 1\}^\omega$ and that $\mathbf{b}_{i+1} \in \{0, 1\}^\omega$ is obtained as the binary decomposition of a random $\mathbf{a}_{i+1} \leftarrow U(\mathbb{Z}_q^n)$. The probability that $\mathbf{b}_{i+1}$ is independent of $\mathbf{b}_1, \dots, \mathbf{b}_i$ is $\geq (q^n - 2^i)/q^n$. If we sample $\mathbf{a}_1, \dots, \mathbf{a}_\omega \leftarrow U(\mathbb{Z}_q^n)$, the probability that their binary decompositions are linearly independent over $\mathbb{Z}_q$ is

$$\prod_{i=0}^{\omega-1} \Pr[\mathbf{b}_{i+1} \notin \text{span}_{\mathbb{Z}_q}(\mathbf{b}_1, \dots, \mathbf{b}_i)] = \prod_{i=0}^{\omega-1} \frac{q^n - 2^i}{q^n} \quad (1)$$

Note that the factors of the right-hand-side member of (1) are all positive: indeed, we have $2^{\omega/n} \cdot (1 - \frac{1}{2n}) \leq q$ which is equivalent to $2^{\omega-1}/q^n \leq \frac{1}{2} \cdot (1 - \frac{1}{2n})^{-n} < \sqrt{\exp(1)}/2 < 1$ for all values of $n \geq 2$.

We now proceed to bound (1) as

$$\begin{aligned} \prod_{i=0}^{\omega-1} \left(1 - \frac{2^i}{q^n}\right) &= \exp\left(\sum_{i=0}^{\omega-1} \ln\left(1 - \frac{2^i}{q^n}\right)\right) \\ &\geq \exp\left(-\frac{3}{q^n} \cdot \sum_{i=0}^{\omega-1} 2^i\right) \geq \exp\left(-\frac{3}{q^n} \cdot 2^\omega\right) \geq \exp(-3 \cdot \sqrt{\exp(1)}), \end{aligned}$$

where the inequality holds because $\ln(1 - x) \geq -3x$ for all $x \in (0, \sqrt{\exp(1)}/2)$.

□

3.2 A Centralized Construction

Let λ be a security parameter and let $\ell \in \Theta(\lambda)$, $L \in \Theta(\lambda)$. We use parameters consisting of prime moduli p and q such that $q/p > 2^{L+\lambda} \cdot r$, dimensions $n, m, k \in \text{poly}(\lambda)$ such that $m \geq 2n \cdot \lceil \log q \rceil$, an integer $\beta > 0$, $\alpha > 0$ and $r = m^{L+2} \cdot n \cdot \beta \cdot \alpha q$. We rely on the following ingredients.

- A balanced admissible hash function $\text{AHF} : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$.
- A family Π_λ of ξ -wise independent hash functions $\pi_i : \mathbb{Z}_p^m \rightarrow \mathbb{Z}_p^k$ for a suitable $\xi > 0$ that will be determined later on. Let π be a random member of Π_λ . For example, the function π can be a random polynomial $\pi(Z) \in GF(p^m)[Z]$ of degree $\xi - 1$ with outputs truncated to their k first coordinates.

We also choose a Gaussian parameter $\sigma > 0$, which will specify an interval $[-\beta, \beta] = [-\sigma\sqrt{n}, \sigma\sqrt{n}]$ where the coordinates of the secret will be confined (with probability exponentially close to 1). We also need a rounding parameter $r > 0$, set as indicated above.

The pseudorandom function family assumes the availability of public parameters

$$\text{pp} := \left(q, \pi, \mathbf{A}_0, \{ \mathbf{A}_{i,0}, \mathbf{A}_{i,1} \in \mathbb{Z}_q^{n \times m} \}_{i=1}^L, \text{AHF}, r, \sigma \right),$$

where $\mathbf{A}_0 \sim U(\mathbb{Z}_q^{n \times m})$ and $\mathbf{A}_{i,0}, \mathbf{A}_{i,1} \sim U(\mathbb{Z}_q^{n \times m})$ for each $i \in [L]$, conditioned on $\mathbf{G}^{-1}(\mathbf{A}_{i,b}) \in \mathbb{Z}_q^{m \times m}$ being $\mathbb{Z}_q$ -invertible for all $i \in [L]$ and $b \in \{0, 1\}$.

Keygen(pp): Given pp , sample a vector $\mathbf{s} \leftarrow D_{\mathbb{Z}^n, \sigma}$ so that $\|\mathbf{s}\|_\infty < \beta = \sigma\sqrt{n}$ with overwhelming probability. The secret key is $SK := \mathbf{s} \in [-\beta, \beta]^n$.

Eval(pp, SK, X): Given $SK = \mathbf{s} \in \mathbb{Z}^n$ and an input $X \in \{0, 1\}^\ell$,

1. Compute $x = \text{AHF}(X) \in \{0, 1\}^L$ and parse it as $x = x_1 \dots x_L$.
2. Compute

$$\mathbf{z} = \left\lfloor \left(\mathbf{A}(x) \right)^\top \cdot \mathbf{s} \right\rfloor_p \in \mathbb{Z}_p^m, \quad (2)$$

where

$$\mathbf{A}(x) = \mathbf{A}_0 \cdot \prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i, x_i}),$$

and output $\mathbf{y} = \pi(\mathbf{z}) \in \mathbb{Z}_p^k$.

We remark that the way to compute $\mathbf{z} \in \mathbb{Z}_p^m$ in (2) is reminiscent of the key-homomorphic PRFs of [17,8]. Unlike [17,8], our security proof requires the secret $\mathbf{s}$ to have small entries. Also, our PRF is not key-homomorphic as the output is $\mathbf{y} = \pi(\mathbf{z}) \in \mathbb{Z}_p^k$ instead of $\mathbf{z} \in \mathbb{Z}_p^m$. Fortunately, losing the key-homomorphic property does not prevent us from building a DPRF since the randomness extraction step is only applied to the result of combining t partial evaluations in the **Combine** algorithm.

Theorem 3.2. *Set an entropy lower bound $\bar{n} = \lfloor n \cdot \log \sigma - n' \cdot \log q \rfloor - 1$ as $\Omega(\lambda)$. If we choose the output length $\bar{k} = k \cdot \log p$ in such a way that*

$$\xi = \bar{n} + \ell, \quad \bar{k} = \bar{n} - 2 \cdot (\lambda + \log \ell + \log \bar{n}),$$

then the construction above is a secure PRF family under the $\text{LWE}_{q,m,n',\alpha}$ assumption.

Proof. The proof considers a sequence of hybrid games. In each game, we call W_i the event that $b' = b$.

Game₀: This is the real find-then-guess PRF experiment where the adversary $\mathcal{A}$ outputs $\hat{b} \in \{0, 1\}$ and the challenger defines $b' := \hat{b}$. By definition, the adversary's advantage is $\text{Adv}(\mathcal{A}) := |\Pr[W_0] - 1/2|$, where W_0 is event that $b' = b$.

Game₁: This game is identical to Game₀ with the following changes. First, the challenger runs $K \leftarrow \text{AdmSmp}(1^\lambda, Q, \delta)$, for $\delta := \text{Adv}(\mathcal{A})$, to generate a key $K \in \{0, 1, \perp\}^L$ for a balanced admissible hash function $\text{AHF} : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$. When the adversary halts and outputs $\hat{b} \in \{0, 1\}$, the challenger checks if the conditions

$$P_K(X^{(1)}) = \dots = P_K(X^{(Q)}) = 1 \wedge P_K(X^*) = 0 \quad (3)$$

are satisfied, where X^* is the challenge input and $X^{(1)}, \dots, X^{(Q)}$ are the adversarial queries. If these conditions do not hold, the challenger ignores $\mathcal{A}$'s output $\hat{b} \in \{0, 1\}$ and overwrites it with a random bit $b'' \leftarrow \{0, 1\}$ to define $b' = b''$. If the conditions of (3) are satisfied, the challenger sets $b' = \hat{b}$. By Lemma 2.11, we have

$$\begin{aligned} |\Pr[W_1] - 1/2| &= |\Pr[b' = b] - 1/2| \\ &\geq \gamma_{\min} \cdot \text{Adv}(\mathcal{A}) - \frac{1}{2} \cdot (\gamma_{\max} - \gamma_{\min}) = \tau. \end{aligned}$$

By Theorem 2.10, we get $|\Pr[W_1] - 1/2| \geq \delta/2^{\eta+1}$, where η is chosen as in Theorem 2.10 and satisfies $2^{\eta+1} \leq 2 \cdot (2Q + Q/\delta)^a$, where $a = -\log_{1-c} 2 > 1$ with $c \in (0, 1/2)$ from Theorem 2.10. This yields $2 \cdot (3Q)^a \cdot |\Pr[W_1] - 1/2| \geq |\Pr[W_0] - 1/2|^{a+1}$. Since $Q \in \text{poly}(\lambda)$ and a is constant, if $\mathcal{A}$'s advantage in Game₁ is negligible, so is its advantage in Game₀.

Game₂: In this game, we modify the generation of $\mathbf{pp}$ in the following way. Initially, the challenger samples a uniformly random matrix $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$. Next, for each $i \in [L]$, it samples $\mathbf{R}_{i,0}, \mathbf{R}_{i,1} \leftarrow U(\{-1, 1\}^{m \times m})$ and defines $\{\mathbf{A}_{i,0}, \mathbf{A}_{i,1}\}_{i=1}^L$ as follows for all $i \in [L]$ and $j \in \{0, 1\}$:

$$\mathbf{A}_{i,j} := \begin{cases} \mathbf{A} \cdot \mathbf{R}_{i,j} & \text{if } (j \neq K_i) \wedge (K_i \neq \perp) \\ \mathbf{A} \cdot \mathbf{R}_{i,j} + \mathbf{G} & \text{if } (j = K_i) \vee (K_i = \perp) \end{cases} \quad (4)$$

It also defines $\mathbf{A}_0 = \mathbf{A} \cdot \mathbf{R}_0 + \mathbf{G}$ for a random $\mathbf{R}_0 \leftarrow U(\{-1, 1\}^{m \times m})$. Since $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ was chosen uniformly, the Leftover Hash Lemma ensures that $\{\mathbf{A}_{i,0}, \mathbf{A}_{i,1}\}_{i=1}^L$ are statistically close to the uniform distribution over $\mathbb{Z}_q^{n \times m}$. The challenger keeps sampling $\mathbf{R}_{ij}$ until $\mathbf{G}^{-1}(\mathbf{A}_{ij})$ is $\mathbb{Z}_q$ -invertible in such a way that the simulated $\mathbf{A}_{ij}$ are statistically uniform under the constraint that $\mathbf{G}^{-1}(\mathbf{A}_{ij})$ be invertible over $\mathbb{Z}_q$. By Lemma 3.1, the probability that $\mathbf{G}^{-1}(\mathbf{A}_{ij})$ is invertible in $\mathbb{Z}_q$ is roughly $(\exp(-3\sqrt{\exp(1)}))^2$, so that the public parameters can be simulated efficiently. Since the distribution of $\mathbf{pp}$ is statistically unchanged, it follows that $|\Pr[W_2] - \Pr[W_1]| \leq L \cdot 2^{-\lambda}$.

We note that, at each query X , we can view $\mathbf{A}(x)$ as a GSW encryption

$$\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x + \left(\prod_{i=1}^L \mu_i \right) \cdot \mathbf{G},$$

for some small norm $\mathbf{R}_x \in \mathbb{Z}^{m \times m}$, where

$$\mu_i := \begin{cases} 0 & \text{if } (\text{AHF}(X)_i \neq K_i) \wedge (K_i \neq \perp) \\ 1 & \text{if } (\text{AHF}(X)_i = K_i) \vee (K_i = \perp) \end{cases}$$

If the conditions of (3) are satisfied, at each query $X^{(i)}$, the admissible hash function ensures that $x^{(i)} = \text{AHF}(X^{(i)})$ satisfies

$$\mathbf{A}(x^{(i)}) = \mathbf{A} \cdot \mathbf{R}_{x^{(i)}}, \quad (5)$$

for some small norm $\mathbf{R}_{x^{(i)}} \in \mathbb{Z}^{m \times m}$. Moreover, the challenge input X^* is mapped to an L -bit string $x^* = \text{AHF}(X^*)$ such that

$$\mathbf{A}(x^*) = \mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G}. \quad (6)$$

Game₃: In this game, we modify the distribution of $\mathbf{pp}$ and replace the uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ by a lossy matrix such that

$$\mathbf{A}^\top = \bar{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{E} \in \mathbb{Z}_q^{m \times n}, \quad (7)$$

where $\bar{\mathbf{A}} \leftarrow U(\mathbb{Z}_q^{n' \times m})$, $\mathbf{C} \leftarrow U(\mathbb{Z}_q^{n' \times n})$ and $\mathbf{E} \leftarrow D_{\mathbb{Z}^{m \times n}, \alpha q}$, for $n' \ll n$. The matrix (7) is thus “close” to a matrix $\bar{\mathbf{A}}^\top \cdot \mathbf{C}$ of much lower rank than n . Under the LWE assumption in dimension n' , this change should not significantly alter $\mathcal{A}$ ’s behavior and a straightforward reduction $\mathcal{B}$ shows that $|\Pr[W_3] - \Pr[W_2]| \leq n \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{q,m,n',\alpha}}(\lambda)$, where the factor n comes from the use of an LWE assumption with n secrets.

Due to the modification introduced in Game₃, if the conditions of (3) are satisfied, each query $X^{(i)}$ receives a response of the form

$$\left\lfloor \left(\mathbf{A} \cdot \mathbf{R}_{x^{(i)}} \right)^\top \cdot \mathbf{s} \right\rfloor_p = \left\lfloor \left(\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{R}_{x^{(i)}}^\top \cdot \mathbf{E} \right) \cdot \mathbf{s} \right\rfloor_p. \quad (8)$$

Game₄: In this game, we modify the evaluation oracle and introduce a bad event. We define bad_x to be the event that for the input $x \in \{0, 1\}^L$ such that $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x$, for some small-norm $\mathbf{R}_x \in \mathbb{Z}^{m \times m}$, we have

$$\left\lfloor \left(\mathbf{A} \cdot \mathbf{R}_x \right)^\top \cdot \mathbf{s} \right\rfloor_p \neq \left\lfloor \left(\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \right) \cdot \mathbf{s} \right\rfloor_p. \quad (9)$$

We also define **BAD** as the event that the adversary $\mathcal{A}$ makes an evaluation query $x \in \{0, 1\}^L$ such that the event bad_x occurs. Note that the challenger can detect this event since it knows $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n' \times m}$, $\mathbf{C} \in \mathbb{Z}_q^{n' \times n}$ and $\mathbf{E} \in \mathbb{Z}^{m \times n}$ satisfying (7). If bad_x occurs, the challenger overwrites $\mathcal{A}$ ’s output $\hat{b}$ with a random bit $b'' \leftarrow \{0, 1\}$ and sets $b' = b''$ (otherwise, it sets $b' = \hat{b}$ as before).

Claim. If $q/p > 2^{L+\lambda} \cdot r$, where $r = m^{L+2} \cdot n \cdot \beta \cdot \alpha q$, we have the inequality

$$|\Pr[W_4] - \Pr[W_3]| \leq \Pr[\text{BAD}] \leq 2^{-\Omega(\lambda)}.$$

Proof. We use the fact that, for each query $x = x_1 \dots x_L$, the matrix $\mathbf{R}_x \in \mathbb{Z}^{m \times m}$ is of the form

$$\begin{aligned}
\mathbf{R}_x &= \mathbf{R}_0 \cdot \prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad + \mathbf{R}_{1,x_1} \cdot \prod_{i=2}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad + \mu_{1,x_1} \cdot \mathbf{R}_{2,x_2} \cdot \prod_{i=3}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad + \mu_{1,x_1} \mu_{2,x_2} \cdot \mathbf{R}_{3,x_3} \cdot \prod_{i=4}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad \vdots \\
&\quad + \mu_{1,x_1} \mu_{2,x_2} \dots \mu_{L-2,x_{L-2}} \cdot \mathbf{R}_{L-1,x_{L-1}} \cdot \mathbf{G}^{-1}(\mathbf{A}_{L,x_L}) \\
&\quad + \mu_{1,x_1} \mu_{2,x_2} \dots \mu_{L-1,x_{L-1}} \cdot \mathbf{R}_{L,x_L}.
\end{aligned}$$

In order to bound $\Pr[\text{bad}_x]$, we first observe that, as argued in [7, Proof of Theorem 7.3], $\mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top$ is statistically uniform over $\mathbb{Z}_q^{m \times n'}$, even conditionally on $\mathbf{R}_0^\top \cdot \mathbf{E} \in \mathbb{Z}^{m \times n}$. Hence, for a fixed choice of x , the product $\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \mathbf{s})$ contains a term

$$\underbrace{\left(\prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \right)^\top}_{\mathbf{G}_x^\top} \cdot \mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \mathbf{s}),$$

for which we know that $\mathbf{G}_x^\top \cdot \mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top$ is statistically uniform over $\mathbb{Z}_q^{m \times n'}$ since $\mathbf{G}_x^\top$ is a $\mathbb{Z}_q$ -invertible matrix by construction.

Lemma 2.3 implies that $\mathbf{G}_x^\top \cdot \mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \mathbf{s})$ is statistically close to the uniform distribution over $\mathbb{Z}_q^m$. By Lemma 2.2, we thus have

$$\Pr[\text{bad}_x] = \Pr \left[\left| \left(\mathbf{A} \cdot \mathbf{R}_x \right)^\top \cdot \mathbf{s} \right|_p \neq \left| \left(\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \right) \cdot \mathbf{s} \right|_p \right] \leq m \cdot \frac{2rp}{q}. \quad (10)$$

where $r = m^{L+2} \cdot n \cdot \beta \cdot \alpha q$. Indeed, we have $\|\mathbf{R}_x^\top\|_\infty \leq \frac{m(m^{L+1}-1)}{m-1}$, $\|\mathbf{s}\|_\infty \leq \beta$ and

$$\|\mathbf{E}\|_\infty = \max_{i \in [m]} \left(\sum_{j=1}^n |e_{i,j}| \right) \leq \sqrt{n} \max_{i \in [m]} \|\mathbf{e}_i\| \leq n \cdot \alpha q.$$

The conditions of Lemma 2.2 are met since

$$\|\mathbf{R}_x^\top \cdot \mathbf{E} \cdot \mathbf{s}\|_\infty \leq \|\mathbf{R}_x^\top\|_\infty \cdot \|\mathbf{E}\|_\infty \cdot \|\mathbf{s}\|_\infty \leq \frac{m(m^{L+1}-1)}{m-1} \cdot n \cdot \alpha q \cdot \beta \leq r,$$

which yields inequality (10).

By taking the union bound over all possible $x \in \{0, 1\}^L$, we can bound the probability $\Pr[\text{BAD}] \leq 2^L \cdot m \cdot \frac{2rp}{q}$.

Notice that Game 3 and Game 4 are identical if **BAD** does not occur. This implies the inequality

$$|\Pr[W_4] - \Pr[W_3]| \leq \Pr[\text{BAD}] \leq 2^L \cdot m \cdot \frac{2rp}{q} = 2m \cdot 2^{-\lambda}$$

which is negligible in λ . $\square$

We recall that, if **BAD** does not occur, we have

$$\left[(\mathbf{A} \cdot \mathbf{R}_{x^{(i)}})^\top \cdot \mathbf{s} \right]_p = \left[(\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C}) \cdot \mathbf{s} \right]_p. \quad (11)$$

at each query $x^{(i)}$. We note that the right-hand-side member of (11) is completely determined by $\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top$ and the product $\mathbf{C} \cdot \mathbf{s} \in \mathbb{Z}_q^{n'}$. This means that $\mathcal{A}$'s queries always reveal the same $n' \cdot \log q$ bits of information about $\mathbf{s} \in \mathbb{Z}^n$, meaning that $\mathbf{s}$ retains a lot of entropy even after all queries.

Game₅: In this game, we modify the challenge value for which, if $b = 1$, the adversary is given a random $\mathbf{y}^* \leftarrow U(\mathbb{Z}_p^k)$. Clearly, we have $\Pr[W_5] = 1/2$ since the distribution of the challenge value does not depend on $b \in \{0, 1\}$. In Game₄, if $b = 1$, the output of the PRF is of the form

$$\mathbf{y}^* = \pi \left(\left[(\mathbf{A}(x^*))^\top \cdot \mathbf{s} \right]_p \right) = \pi \left(\left[(\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G})^\top \cdot \mathbf{s} \right]_p \right). \quad (12)$$

We argue that the right-hand-side member of (12) is statistically uniform over $\mathbb{Z}_p^k$. Indeed, the source $\mathbf{z}^* = \left[(\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G})^\top \cdot \mathbf{s} \right]_p$ depends on a term $\mathbf{G}^\top \cdot \mathbf{s} \in \mathbb{Z}_q^m$, which is an injective function of $\mathbf{s}$. More precisely, if **BAD** does not occur, we know that $\mathcal{A}$ has at most $n' \cdot \log q$ bits of information about $\mathbf{s}$.

We also note that $\mathbf{z}^*$ can be written

$$\mathbf{z}^* = \left[(\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G})^\top \cdot \mathbf{s} \right]_p = \left[\mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{s} \right]_p + \left[\mathbf{G}^\top \cdot \mathbf{s} \right]_p + \mathbf{e}_{s, x^*}, \quad (13)$$

for some $\mathbf{e}_{s, x^*} \in \{0, 1\}^m$. Moreover, the proof of the claim implies that

$$\left[\mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{s} \right]_p = \left[\mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \cdot \mathbf{s} \right]_p$$

since it considers a union bound to show that equality (9) holds for all x 's with overwhelming probability. This implies $H_\infty(\left[\mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{s} \right]_p \mid \mathbf{C} \cdot \mathbf{s}) = 0$ with high probability. In the right-hand-side member of (13), we also observe that $\left[\mathbf{G}^\top \cdot \mathbf{s} \right]_p + \mathbf{e}_{s, x^*}$ is an injective function of $\mathbf{s}$. Indeed, we have

$$\left[\mathbf{G}^\top \cdot \mathbf{s} \right]_p + \mathbf{e}_{s, x^*} = (p/q) \cdot \mathbf{G}^\top \cdot \mathbf{s} - \mathbf{t}_{s, x^*} + \mathbf{e}_{s, x^*}$$

for some $\mathbf{t}_{s,x^*} \in (0, 1)^m$, so that

$$(q/p) \cdot (\lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_p + \mathbf{e}_{s,x^*}) = \mathbf{G}^\top \cdot \mathbf{s} + \mathbf{e}'_{s,x^*} \quad (14)$$

for some $\mathbf{e}'_{s,x^*} \in (-q/p, 2 \cdot q/p)^m$. The vector $\mathbf{s}$ is thus uniquely determined by (14) using the public trapdoor of $\mathbf{G}$ as long as $q/p \ll q$.

When we consider the entropy of $\mathbf{z}^*$ conditionally on $\mathcal{A}$'s view, we have

$$\begin{aligned} H_\infty(\mathbf{z}^* \mid \mathbf{C} \cdot \mathbf{s}) &= H_\infty(\lfloor \mathbf{R}_{x^*}^\top \cdot \mathbf{A}^\top \cdot \mathbf{s} \rfloor_p + \lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_p + \mathbf{e}_{s,x^*} \mid \mathbf{C} \cdot \mathbf{s}) \\ &= H_\infty(\lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_p + \mathbf{e}_{s,x^*} \mid \mathbf{C} \cdot \mathbf{s}) \\ &= H_\infty(\mathbf{s} \mid \mathbf{C} \cdot \mathbf{s}) \geq n \cdot \log \sigma - n' \cdot \log q - 1. \end{aligned}$$

Here, the second equality follows from the fact that, for any random variables X, Y, Z defined over an additive group, we have $H_\infty(Y + Z \mid X) = H_\infty(Z \mid X)$ if $H_\infty(Y \mid X) = 0$. The last inequality follows from Lemma 2.5: the distribution of $\mathbf{s} \in \mathbb{Z}_q^n$ conditionally on $\mathbf{C} \cdot \mathbf{s}$ is $\mathbf{s}_0 + D_{\Lambda^\perp(\mathbf{C}), \sigma, -\mathbf{s}_0}$, where $\mathbf{s}_0 \in \mathbb{Z}_q^n$ is an arbitrary solution of $\mathbf{C} \cdot \mathbf{s}_0 = \mathbf{C} \cdot \mathbf{s}$. Then, Lemma 2.5 implies that the point with highest probability in $D_{\Lambda^\perp(\mathbf{C}), \sigma, -\mathbf{s}_0}$ occurs with probability $\leq 2 \det(\Lambda^\perp(\mathbf{C}))/\sigma^n$. Since $\det(\Lambda^\perp(\mathbf{C})) = q^{n'}$ with probability $1 - 2^{-\Omega(\lambda)}$, we have $H_\infty(\mathbf{s} \mid \mathbf{C} \cdot \mathbf{s}) \geq n \cdot \log \sigma - n' \cdot \log q - 1$.

To extract statistically uniform bits from $\tilde{\mathbf{z}}^*$, we need to account for its possible dependency on x^* which may depend on $\mathbf{pp}$. Assuming that $P_K(X^*) = 0$, the source $\tilde{\mathbf{z}}^*$ is taken from a distribution determined by X^* within a collection of less than 2^ℓ distributions (i.e., those for which $P_K(X^*) = 0$), which all have min-entropy $\bar{n} = n \cdot \log \sigma - n' \cdot \log q - 1$. By applying Lemma 2.13 with $\varepsilon = 2^{-\lambda}$ for a collection $\mathcal{X}$ of at most $M = 2^\ell$ distributions, we obtain that the distribution of $\pi(\tilde{\mathbf{z}}^*)$ is $2^{-\lambda}$ -close to the uniform distribution over $\mathbb{Z}_p^k$. $\square$

4 The DPRF Construction

We build a distributed PRF by using a LISS inside the PRF construction of Section 3. As mentioned earlier, the latter is well-suited to our purposes because, in the security proof, the secret key is known to the challenger at any time. When the secret key $\mathbf{s}$ is shared using a LISS, the challenger is always able to consistently answer corruption queries because it has all shares at disposal.

In the construction, we rely on the specific LISS construction of Damgård and Thorbek [26], which is based on the Benaloh-Leichter secret sharing [11]. This particular LISS scheme is well-suited to our needs for several reasons. First, it has binary share generating matrices, which allows obtaining relatively short shares of $\mathbf{s} \in \mathbb{Z}^n$: in the security proof, this is necessary to ensure that the adversary always obtains the same information about uncorrupted shares in partial evaluation queries. Another advantage of the Benaloh-Leichter-based LISS is that its reconstruction constants live in $\{-1, 0, 1\}$, which avoids blowing up the homomorphism errors when partial evaluations are combined together. Finally, its sweeping vectors also have their coordinates in $\{-1, 0, 1\}$ (whereas they may be exponentially large in

the number N of servers in the construction due to Cramer and Fehr [25]) and we precisely need sweeping vectors κ to be small in the proof of our Lemma 4.4.

For simplicity, this section presents a DPRF which is only secure against passive adversaries. In Section 5, we show that our scheme can be made robust against malicious adversaries (Definition 5.1) using homomorphic signatures by applying techniques suggested in [15,16].

4.1 Description

Setup($1^\lambda, 1^\ell, 1^t, 1^N$): On input a security parameter λ , a number of servers N , a threshold $t \in [1, N]$ and an input length $\ell \in \Theta(\lambda)$, set $d, e = O(N^{1+\sqrt{2}})$. Choose a real $\alpha > 0$, a Gaussian parameter $\sigma = \Omega(\sqrt{e \cdot n})$, which will specify an interval $[-\beta, \beta] = [-\sigma\sqrt{n}, \sigma\sqrt{n}]$ where the coordinates of the secret will live (with probability exponentially close to 1). Next, do the following.

1. Choose prime moduli p, q and u such that $p/u > d \cdot 2^{L+\lambda}$ and $q/p > 2^{L+\lambda} \cdot r$, where the dimensions $n, m, k \in \text{poly}(\lambda)$ are such that $m \geq 2n \cdot \lceil \log q \rceil$, and $r = m^{L+2} \cdot n \cdot \beta^* \cdot \alpha q$ with $\beta^* = O(\beta \cdot \log N)$.
2. Choose a balanced admissible hash function $\text{AHF} : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$, for a suitable $L \in \Theta(\lambda)$. Choose a family Π_λ of ξ -wise independent hash functions $\pi : \mathbb{Z}_u^m \rightarrow \mathbb{Z}_u^k$, for a suitable integer $\xi > 0$, with $\pi \leftarrow U(\Pi_\lambda)$.
3. Choose random matrices $\mathbf{A}_0 \leftarrow U(\mathbb{Z}_q^{n \times m})$ and $\mathbf{A}_{i,b} \leftarrow U(\mathbb{Z}_q^{n \times m})$, for each $i \in [L]$, $b \in \{0, 1\}$, subject to the constraint that $\mathbf{G}^{-1}(\mathbf{A}_{i,b}) \in \mathbb{Z}^{m \times m}$ be $\mathbb{Z}_q$ -invertible for all $i \in [L]$ and $b \in \{0, 1\}$.

Output

$$\text{pp} := \left(q, p, u, \pi, \mathbf{A}_0, \{\mathbf{A}_{i,0}, \mathbf{A}_{i,1} \in \mathbb{Z}_q^{n \times m}\}_{i=1}^L, \text{AHF} \right),$$

Share(pp, SK_0): Given pp and a key $SK_0 = \mathbf{s}$ consisting of an integer vector $\mathbf{s}$ sampled from the Gaussian distribution $D_{\mathbb{Z}^n, \sigma}$, return $\perp$ if $\mathbf{s} \notin [-\beta, \beta]^n$, where $\beta = \sigma\sqrt{n}$. Otherwise, generate a LISS of $\mathbf{s}$ as follows.

1. Using the BL-based LISS scheme, construct the matrix $\mathbf{M} \in \{0, 1\}^{d \times e}$ that computes the Boolean formula associated with the $T_{t,N}$ threshold function. By using [44], we obtain a matrix $\mathbf{M} \in \{0, 1\}^{d \times e}$, so that each row of $\mathbf{M}$ contains $O(\log N)$ non-zero entries.
2. For each $k \in [n]$, generate a LISS of the k -th coordinate s_k of $\mathbf{s} \in \mathbb{Z}^n$. To this end, define a vector $\boldsymbol{\rho}_k = (s_k, \rho_{k,2}, \dots, \rho_{k,e})^\top$, with Gaussian entries $\rho_{k,2}, \dots, \rho_{k,e} \leftarrow D_{\mathbb{Z}, \sigma}$, and compute

$$\mathbf{s}_k = (s_{k,1}, \dots, s_{k,d})^\top = \mathbf{M} \cdot \boldsymbol{\rho}_k \in \mathbb{Z}^d,$$

whose entries are smaller than $\|\mathbf{s}_k\|_\infty \leq \beta^* = O(\beta \cdot \log N)$.

3. Define the matrix $\mathbf{S} = [\mathbf{s}_1 \mid \dots \mid \mathbf{s}_n] \in \mathbb{Z}^{d \times n}$. For each $j \in [N]$, define the share of P_j to be the sub-matrix $\mathbf{S}_{I_j} = \mathbf{M}_{I_j} \cdot [\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n] \in \mathbb{Z}^{d_j \times n}$, where $I_j = \psi^{-1}(j) \subset \{1, \dots, d\}$ is the set of indexes such that P_j owns the sub-matrix $\mathbf{M}_{I_j} \in \{0, 1\}^{d_j \times e}$.

For each $j \in [N]$, the share $SK_j = \mathbf{S}_{I_j} \in \mathbb{Z}^{d_j \times n}$ is privately sent to P_j .

PEval(pp, SK_j, X): Given $SK_j = \mathbf{S}_{I_j} \in \mathbb{Z}^{d_j \times n}$ and an input $X \in \{0, 1\}^\ell$,

1. Compute $x = \text{AHF}(X) \in \{0, 1\}^L$ and parse it as $x = x_1 \dots x_L$.
2. Parse $\mathbf{S}_{I_j}^\top = [\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n]^\top \cdot \mathbf{M}_{I_j}^\top \in \mathbb{Z}^{n \times d_j}$ as $[\bar{\mathbf{s}}_{j,1} \mid \dots \mid \bar{\mathbf{s}}_{j,d_j}]$. For each $\theta \in \{1, \dots, d_j\}$, compute

$$\mathbf{z}_{j,\theta} = \left[\left(\mathbf{A}(x) \right)^\top \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p \in \mathbb{Z}_p^m, \quad (15)$$

where

$$\mathbf{A}(x) = \mathbf{A}_0 \cdot \prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}),$$

and output the partial evaluation $\mathbf{Y}_j = [\mathbf{z}_{j,1} \mid \dots \mid \mathbf{z}_{j,d_j}] \in \mathbb{Z}_p^{m \times d_j}$.

Eval(pp, SK_0, X): Given $SK_0 = \mathbf{s} \in \mathbb{Z}^n$ and an input $X \in \{0, 1\}^\ell$,

1. Compute $x = \text{AHF}(X) \in \{0, 1\}^L$ and write it as $x = x_1 \dots x_L$.
2. Compute

$$\tilde{\mathbf{z}} = \left[\left(\mathbf{A}(x) \right)^\top \cdot \mathbf{s} \right]_p \in \mathbb{Z}_p^m,$$

where $\mathbf{A}(x) = \mathbf{A}_0 \cdot \prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i})$, and output $\mathbf{y} = \pi([\tilde{\mathbf{z}}]_u) \in \mathbb{Z}_u^k$.

Combine($\mathcal{S}, (\mathbf{Y}_{j_1}, \dots, \mathbf{Y}_{j_t})$): Write the set $\mathcal{S}$ as $\mathcal{S} = \{j_1, \dots, j_t\}$ and parse each partial evaluation $\mathbf{Y}_{j_\kappa} \in \mathbb{Z}_p^{m \times d_{j_\kappa}}$ as $[\mathbf{z}_{j_\kappa,1} \mid \dots \mid \mathbf{z}_{j_\kappa,d_{j_\kappa}}]$ for all $\kappa \in [t]$.

1. Determine the vector $\boldsymbol{\lambda}_{\mathcal{S}} \in \{-1, 0, 1\}^{d_{\mathcal{S}}}$ such that $\boldsymbol{\lambda}_{\mathcal{S}}^\top \cdot \mathbf{M}_{\mathcal{S}} = (1, 0, \dots, 0)^\top$, where $\mathbf{M}_{\mathcal{S}} \in \{0, 1\}^{d_{\mathcal{S}} \times e}$ is the sub-matrix of $\mathbf{M}$ owned by the parties in $\mathcal{S}$ and $d_{\mathcal{S}} = \sum_{\kappa=1}^t d_{j_\kappa}$ with $d_{j_\kappa} = |\psi^{-1}(j_\kappa)|$ for all $\kappa \in [t]$. Then, parse $\boldsymbol{\lambda}_{\mathcal{S}}$ as $[\boldsymbol{\lambda}_{j_1}^\top \mid \dots \mid \boldsymbol{\lambda}_{j_t}^\top]^\top$, where $\boldsymbol{\lambda}_{j_\kappa} \in \{-1, 0, 1\}^{d_{j_\kappa}}$ for all $\kappa \in [t]$.
2. Compute $\tilde{\mathbf{z}} = \sum_{\kappa=1}^t \mathbf{Y}_{j_\kappa} \cdot \boldsymbol{\lambda}_{j_\kappa} \in \mathbb{Z}_p^m$, which equals

$$\tilde{\mathbf{z}} = \left[\left(\mathbf{A}(x) \right)^\top \cdot \mathbf{s} \right]_p + \mathbf{e}_z \in \mathbb{Z}_p^m,$$

for some $\mathbf{e}_z \in \{-2d_{\mathcal{S}}, \dots, 2d_{\mathcal{S}}\}^m$.

3. Compute $\mathbf{z} = [\tilde{\mathbf{z}}]_u \in \mathbb{Z}_u^m$, which equals

$$\mathbf{z} = \left[\left[\left(\mathbf{A}(x) \right)^\top \cdot \mathbf{s} \right]_p \right]_u \in \mathbb{Z}_u^m$$

with overwhelming probability. Finally, output $\mathbf{y} = \pi(\mathbf{z}) \in \mathbb{Z}_u^k$.

By setting $\sigma = \sqrt{e \cdot n} = O(N^{\frac{1+\sqrt{2}}{2}}) \cdot \sqrt{n}$ as allowed by [44], we have share units of magnitude $\beta^* = \Theta(\sigma \sqrt{n} \log N) = O\left(N^{\frac{1+\sqrt{2}}{2}} \log N\right) \cdot n$. Since $d = O(N^{1+\sqrt{2}})$, the average share size amounts to $\frac{d \cdot n \cdot \log \beta^*}{N} = n \cdot N^{\sqrt{2}} \cdot (\log n + O(\log N))$ bits.

Regarding the parameters, Theorem 4.2 allows us to rely on the presumed hardness of $\text{LWE}_{q,m,n',\alpha}$ for n' which may be set as $\Theta(n(\log Nn)/(\log q))$ if $n(\log Nn) = \Omega(\lambda)$. To make sure that the best known attacks on LWE require 2^λ bit operations, it suffices to have $\alpha q = \Omega(\sqrt{n'})$ and $n' \log q / \log^2 \alpha = \Omega(\lambda / \log \lambda)$. We may set $n = \text{poly}(\lambda)$ (for a small degree polynomial) and $q = 2^{\Omega(\lambda \log \lambda)}$ since r contains a term $m^L = \text{poly}(\lambda)^{\Theta(\lambda)} = 2^{O(\lambda \log \lambda)}$.

We remark that our modulus q is exponential in the input length L , but not in the number of servers N . In contrast, the DPRF of [17] requires an exponential modulus in N incurred by the use of Shamir's secret sharing and the technique of clearing out the denominators [3].

4.2 Security and Correctness

We now show that the construction provides statistical consistency.

Lemma 4.1. *Let $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$ and let a secret key $SK_0 = s \leftarrow D_{\mathbb{Z}^n, \sigma}$, which is shared as $(SK_1, \dots, SK_N) \leftarrow \text{Share}(\text{pp}, SK_0)$. For any t -subset $\mathcal{S} = \{j_1, \dots, j_t\} \subset [N]$ and input $X \in \{0, 1\}^\ell$, if $Y_{j_k} = \text{PEval}(\text{pp}, SK_{j_k}, X)$ for all $\kappa \in [t]$, we have*

$$\text{Combine}(\mathcal{S}, (Y_{j_1}, \dots, Y_{j_t})) = \text{Eval}(\text{pp}, SK_0, X)$$

with probability exponentially close to 1.

Proof. Let $\boldsymbol{\lambda}_\mathcal{S} \in \{-1, 0, 1\}^{d_\mathcal{S}}$ such that $\boldsymbol{\lambda}_\mathcal{S}^\top \cdot \mathbf{M}_\mathcal{S} = (1, 0, \dots, 0) \in \mathbb{Z}^e$. If we parse $\boldsymbol{\lambda}_\mathcal{S}$ as $[\boldsymbol{\lambda}_{j_1}^\top \mid \dots \mid \boldsymbol{\lambda}_{j_t}^\top]^\top$ we have

$$\mathbf{s} = [\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n]^\top \cdot \mathbf{M}_\mathcal{S}^\top \cdot \boldsymbol{\lambda}_\mathcal{S} = \sum_{k=1}^t \mathbf{S}_{I_{j_k}}^\top \cdot \boldsymbol{\lambda}_{j_k}.$$

In turn, this implies

$$\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p = \left\lfloor \sum_{k=1}^t \mathbf{A}(x)^\top \cdot \mathbf{S}_{I_{j_k}}^\top \cdot \boldsymbol{\lambda}_{j_k} \right\rfloor_p = \sum_{k=1}^t \left\lfloor \mathbf{A}(x)^\top \cdot \mathbf{S}_{I_{j_k}}^\top \right\rfloor_p \cdot \boldsymbol{\lambda}_{j_k} + \mathbf{e}, \quad (16)$$

where the last equality of (16) stems from fact that, for any two vectors $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{Z}_q^m$, we have $\lfloor \mathbf{v}_1 + \mathbf{v}_2 \rfloor_p = \lfloor \mathbf{v}_1 \rfloor_p + \lfloor \mathbf{v}_2 \rfloor_p + \mathbf{e}^+$, for some $\mathbf{e}^+ \in \{0, 1\}^m$, and $\lfloor \mathbf{v}_1 - \mathbf{v}_2 \rfloor_p = \lfloor \mathbf{v}_1 \rfloor_p - \lfloor \mathbf{v}_2 \rfloor_p + \mathbf{e}^-$, where $\mathbf{e}^- \in \{-1, 0\}^m$. The error vector $\mathbf{e}$ of (16) thus lives in $\{-d_\mathcal{S}, \dots, d_\mathcal{S}\}^m$. By the definition of $\mathbf{Y}_{j_k} = \lfloor \mathbf{A}(x)^\top \cdot \mathbf{S}_{I_{j_k}}^\top \rfloor_p$, if we define $\tilde{\mathbf{z}} := \sum_{k=1}^t \mathbf{Y}_{j_k} \cdot \boldsymbol{\lambda}_{j_k}$ and $\mathbf{e}_z := -\mathbf{e} \in \{-d_\mathcal{S}, \dots, d_\mathcal{S}\}^m$, we have

$$\tilde{\mathbf{z}} = \left\lfloor \left(\mathbf{A}(x) \right)^\top \cdot \mathbf{s} \right\rfloor_p + \mathbf{e}_z \in \mathbb{Z}_p^m.$$

Then, we observe that $\mathbf{A}(x)^\top \cdot \mathbf{s}$ is of the form $\mathbf{T}_q \cdot \mathbf{A}_0^\top \cdot \mathbf{s}$, for some matrix $\mathbf{T}_q = (\prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}))^\top$ which is a product of $\mathbb{Z}_q$ -invertible matrices. By Lemma 2.3, the vector $\mathbf{A}_0^\top \cdot \mathbf{s}$ is statistically close to the uniform distribution $U(\mathbb{Z}_q^m)$. Since $\mathbf{T}_q$ is invertible, the vector $\mathbf{A}(x)^\top \cdot \mathbf{s}$ is itself statistically close to $U(\mathbb{Z}_q^m)$. Hence, the vector $\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p$ is statistically close to $U(\mathbb{Z}_p^m)$ since the statistical distance between $\lfloor U(\mathbb{Z}_q^m) \rfloor_p$ and $U(\mathbb{Z}_p^m)$ is at most $m \cdot (p/q)$. Therefore we can apply Lemma 2.2, which implies that

$$\left\lfloor \lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p \right\rfloor_u = \left\lfloor \lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p + \mathbf{e}_z \right\rfloor_u$$

except with probability $2^L \cdot m \cdot \frac{4d \cdot u}{p} \leq 2^L \cdot m \cdot \frac{4d \cdot u}{p} \leq m \cdot 2^{-\lambda}$.

This shows that the equality $\lfloor \tilde{\mathbf{z}} \rfloor_u = \left\lfloor \lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p \right\rfloor_u$ holds with overwhelming probability if the vector $\tilde{\mathbf{z}} := \sum_{k=1}^t \mathbf{Y}_{j_k} \cdot \boldsymbol{\lambda}_{j_k}$ in the left-hand-side member is computed by **Combine** and the right-hand-side member is the $\lfloor \tilde{\mathbf{z}} \rfloor_u$ computed by **Eval**. $\square$

Theorem 4.2. *Assume an entropy lower bound $\bar{n} = \lfloor n \cdot \log \sigma - \frac{n}{2} \cdot \log e - n' \cdot \log q \rfloor - 1$ is $\Omega(\lambda)$. If we set the output length $\bar{k} = k \cdot \log u$ so as to have*

$$\xi = \bar{n} + \ell, \quad \bar{k} = \bar{n} - 2 \cdot (\lambda + \log \ell + \log \bar{n}),$$

then the construction above is an adaptively secure DPRF family under the $\text{LWE}_{q,m,n',\alpha}$ assumption.

Proof. The proof considers a sequence of hybrid games. In each game, we call W_i the event that $b' = b$.

Game₀: This is the experiment described in Definition 2.22. Namely, the challenger initially samples a secret Gaussian vector $SK_0 = \mathbf{s} \leftarrow D_{\mathbb{Z}^n, \sigma}$, which is shared by computing

$$\mathbf{S}_{I_j} = \mathbf{M}_{I_j} \cdot [\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n] = M_{I_j} \cdot \boldsymbol{\Gamma} \in \mathbb{Z}^{d_j \times n} \quad \forall j \in [N],$$

where

$$\boldsymbol{\Gamma} = [\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n] = \begin{bmatrix} \mathbf{s}^\top \\ \rho_{1,2} \cdots \rho_{n,2} \\ \vdots \quad \ddots \quad \vdots \\ \rho_{1,e} \cdots \rho_{n,e} \end{bmatrix} \in \mathbb{Z}^{e \times n},$$

with $\rho_{k,\nu} \leftarrow D_{\mathbb{Z}, \sigma}$ for all $(k, \nu) \in [1, n] \times [2, e]$. At each partial evaluation query $(j, X^{(i)}) \in [N] \times \{0, 1\}^\ell$, the adversary $\mathcal{A}$ obtains

$$\mathbf{Y}_j = \left\lfloor (\mathbf{A}(x))^\top \cdot \mathbf{S}_{I_j}^\top \right\rfloor_p \in \mathbb{Z}_p^{m \times d_j}. \quad (17)$$

In the challenge phase, the adversary chooses an input $X^* \in \{0, 1\}^\ell$. It obtains a random vector $\mathbf{y}^* \leftarrow U(\mathbb{Z}_u^k)$ if the challenger's bit is $b = 0$. If $b = 1$, it obtains the real evaluation $\mathbf{y}^* = \pi(\lfloor \tilde{\mathbf{z}}^* \rfloor_u) \in \mathbb{Z}_u^k$, where

$$\tilde{\mathbf{z}}^* = \left\lfloor \left(\mathbf{A}(x^*) \right)^\top \cdot \mathbf{s} \right\rfloor_p \in \mathbb{Z}_p^m,$$

with $\mathbf{A}(x^*) = \mathbf{A}_0 \cdot \prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i, x_i^*})$ and $x^* = \text{AHF}(X^*) \in \{0, 1\}^L$. At the end of the game, we denote $\mathcal{C}^* \subseteq [N]$ the set of servers that were corrupted by $\mathcal{A}$ or such that an evaluation query of the form (i, X^*) was made. By hypothesis, we have $|\mathcal{C}^*| < t$. When the adversary halts, it outputs $\hat{b} \in \{0, 1\}$ and the challenger defines $b' := \hat{b}$. The adversary's advantage is $\mathbf{Adv}(\mathcal{A}) := |\Pr[W_0] - 1/2|$, where W_0 is event that $b' = b$.

Game₁: This game is identical to Game₀ with the following changes. First, the challenger runs $K \leftarrow \text{AdmSmp}(1^\lambda, Q, \delta)$ to generate a key $K \in \{0, 1, \perp\}^L$ for a balanced admissible hash function $\text{AHF} : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$, with $\delta := \mathbf{Adv}(\mathcal{A})$ and Q is an upper bound on the number of queries that the adversary makes. When the adversary halts and outputs $\hat{b} \in \{0, 1\}$, the challenger checks if the conditions

$$P_K(X^{(1)}) = \dots = P_K(X^{(Q)}) = 1 \quad \wedge \quad P_K(X^*) = 0 \quad (18)$$

are satisfied, where X^* is the challenge input and $X^{(1)}, \dots, X^{(Q)}$ are the adversarial queries. If these conditions do not hold, the challenger ignores $\mathcal{A}$'s output $\hat{b} \in \{0, 1\}$ and overwrites it with a random bit $b'' \leftarrow \{0, 1\}$ to define $b' = b''$. If conditions (18) are satisfied, the challenger sets $b' = \hat{b}$. By Lemma 2.11, we have

$$\begin{aligned} |\Pr[W_1] - 1/2| &= |\Pr[b' = b] - 1/2| \\ &\geq \gamma_{\min} \cdot \mathbf{Adv}(\mathcal{A}) - \frac{1}{2} \cdot (\gamma_{\max} - \gamma_{\min}) = \tau. \end{aligned}$$

By Theorem 2.10, we have $|\Pr[W_1] - 1/2| \geq \delta/2^{\eta+1}$, where η is chosen as in Theorem 2.10 and satisfies $2^{\eta+1} \leq 2 \cdot (2Q + Q/\delta)^a$, where $a = -\log_{1-c} 2 > 1$ with $c \in (0, 1/2)$. This implies $2 \cdot (3Q)^a \cdot |\Pr[W_1] - 1/2| \geq |\Pr[W_0] - 1/2|^{a+1}$. For polynomially many queries Q and constant a , if $\mathcal{A}$'s advantage is negligible in Game₁, so is it in Game₀.

Game₂: In this game, we modify the generation of $\mathbf{pp}$ in the following way. Initially, the challenger samples a uniformly random matrix $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$. Next, for each $i \in [L]$, it samples $\mathbf{R}_{i,0}, \mathbf{R}_{i,1} \leftarrow U(\{-1, 1\}^{m \times m})$ and defines $\{\mathbf{A}_{i,0}, \mathbf{A}_{i,1}\}_{i=1}^L$ as follows: for all $i \in [L]$ and $j \in \{0, 1\}$:

$$\mathbf{A}_{i,j} := \begin{cases} \mathbf{A} \cdot \mathbf{R}_{i,j} & \text{if } (j \neq K_i) \wedge (K_i \neq \perp) \\ \mathbf{A} \cdot \mathbf{R}_{i,j} + \mathbf{G} & \text{if } (j = K_i) \vee (K_i = \perp) \end{cases} \quad (19)$$

It also defines $\mathbf{A}_0 = \mathbf{A} \cdot \mathbf{R}_0 + \mathbf{G}$ for a random $\mathbf{R}_0 \leftarrow U(\{-1, 1\}^{m \times m})$. Since $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ was chosen uniformly, the Leftover Hash Lemma ensures that $\{\mathbf{A}_{i,0}, \mathbf{A}_{i,1}\}_{i=1}^L$ are statistically independent and uniformly distributed over $\mathbb{Z}_q^{n \times m}$. The challenger keeps sampling the

matrix $\mathbf{R}_{ij}$ until $\mathbf{G}^{-1}(\mathbf{A}_{ij})$ is $\mathbb{Z}_q$ -invertible, so that the simulated matrices $\mathbf{A}_{ij}$ are statistically uniform conditionally on $\mathbf{G}^{-1}(\mathbf{A}_{ij})$ being invertible over $\mathbb{Z}_q$. By Lemma 3.1, the probability that $\mathbf{G}^{-1}(\mathbf{A}_{ij})$ is $\mathbb{Z}_q$ -invertible is roughly the constant $\left(\exp(-3\sqrt{\exp(1)})\right)^2$, which makes it possible to simulate the public parameters efficiently. Since the distribution of $\mathbf{pp}$ is statistically unchanged, it follows that $|\Pr[W_2] - \Pr[W_1]| \leq L \cdot 2^{-\lambda}$.

We note that, at each query X , we can view $\mathbf{A}(x)$ as a GSW encryption

$$\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x + \left(\prod_{i=1}^n \mu_i\right) \cdot \mathbf{G},$$

for some small norm $\mathbf{R}_x \in \mathbb{Z}^{m \times m}$, where

$$\mu_i := \begin{cases} 0 & \text{if } (\text{AHF}(X)_i \neq K_i) \wedge (K_i \neq \perp) \\ 1 & \text{if } (\text{AHF}(X)_i = K_i) \vee (K_i = \perp) \end{cases}$$

If conditions (18) are satisfied, at each query $X^{(i)}$, the admissible hash function ensures that $x^{(i)} = \text{AHF}(X^{(i)})$ satisfies

$$\mathbf{A}(x^{(i)}) = \mathbf{A} \cdot \mathbf{R}_{x^{(i)}}, \quad (20)$$

for some small norm $\mathbf{R}_{x^{(i)}} \in \mathbb{Z}^{m \times m}$. Moreover, the admissible hash function maps the challenge input X^* to an L -bit string $x^* = \text{AHF}(X^*)$ such that

$$\mathbf{A}(x^*) = \mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G}. \quad (21)$$

Game₃: In this game, we modify the distribution of $\mathbf{pp}$ and replace the uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ by a lossy matrix such that

$$\mathbf{A}^\top = \bar{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{E} \in \mathbb{Z}_q^{m \times n}, \quad (22)$$

where $\bar{\mathbf{A}} \leftarrow U(\mathbb{Z}_q^{n' \times m})$, $\mathbf{C} \leftarrow U(\mathbb{Z}_q^{n' \times n})$ and $\mathbf{E} \leftarrow D_{\mathbb{Z}^{m \times n}, \alpha q}$, for n' significantly smaller than n . The matrix in (22) is thus “computationally close” to a matrix $\bar{\mathbf{A}}^\top \cdot \mathbf{C}$ of much lower rank than n . Under the LWE assumption with in dimension n' , this change should not significantly alter $\mathcal{A}$ ’s behavior and a straightforward reduction $\mathcal{B}$ shows that $|\Pr[W_3] - \Pr[W_2]| \leq n \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{q,m,n',\alpha}}(\lambda)$, where the factor n comes from the use of an LWE assumption with n secrets.

The modification introduced in Game₃ has the following consequence. Assuming that conditions (18) are satisfied, for each partial evaluation query $X^{(i)}$ such that $X^{(i)} \neq X^*$, the response is of the form $\mathbf{Y}_j = [\mathbf{z}_{j,1} \mid \dots \mid \mathbf{z}_{j,d_j}] \in \mathbb{Z}_p^{m \times d_j}$, where

$$\begin{aligned} \mathbf{z}_{j,\theta} &= \left[\left(\mathbf{A} \cdot \mathbf{R}_{x^{(i)}} \right)^\top \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p \\ &= \left[\left(\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{R}_{x^{(i)}}^\top \cdot \mathbf{E} \right) \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p \quad \forall \theta \in [d_j]. \end{aligned}$$

Game₄: In this game, we modify the evaluation oracle and introduce a bad event. We define BAD to be the event that the adversary makes a partial evaluation query (j, X) such that the AHF-encoded input $x = \text{AHF}(X) \in \{0, 1\}^L$ corresponds to a matrix $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x$, for some small-norm $\mathbf{R}_x \in \mathbb{Z}^{m \times m}$, such that we have

$$\mathbf{z}_{j,\theta} = \left[\left(\mathbf{A} \cdot \mathbf{R}_x \right)^\top \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p \neq \left[\left(\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \right) \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p. \quad (23)$$

for some $\theta \in [d_j]$. Note that the challenger can detect this event since it knows $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n' \times m}$, $\mathbf{C} \in \mathbb{Z}_q^{n' \times n}$ and $\mathbf{E} \in \mathbb{Z}^{m \times n}$ satisfying (22). If BAD occurs, the challenger overwrites $\mathcal{A}$'s output $\hat{b}$ with a random bit $b'' \leftarrow \{0, 1\}$ and sets $b' = b''$ (otherwise, it sets $b' = \hat{b}$ as before). Lemma 4.3 shows that we have $|\Pr[W_4] - \Pr[W_3]| \leq \Pr[\text{BAD}] \leq 2^{-\Omega(\lambda)}$.

We note that, if BAD does not occur, we have

$$\left[\left(\mathbf{A} \cdot \mathbf{R}_{x^{(i)}} \right)^\top \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p = \left[\left(\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \right) \cdot \bar{\mathbf{s}}_{j,\theta} \right]_p \quad \forall (j, \theta) \in [N] \times [d_j] \quad (24)$$

at each query $(j, X^{(i)})$ for which $X^{(i)} \neq X^*$. We note that the right-hand-side member of (24) is fully determined by $\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top$ and the product $\mathbf{C} \cdot \bar{\mathbf{s}}_{j,\theta} \in \mathbb{Z}_q^{n'}$. This means that partial evaluation queries $(j, X^{(i)})$ such that $X^{(i)} \neq X^*$ always reveal the same information (namely, $\mathbf{C} \cdot \bar{\mathbf{s}}_{j,\theta} \in \mathbb{Z}_q^{n'}$) about $\bar{\mathbf{s}}_{j,\theta} \in \mathbb{Z}^n$.

Conversely, the right-hand-side member of (24) uniquely determines $\mathbf{C} \cdot \bar{\mathbf{s}}_{j,\theta}$ with high probability: observe that $\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top$ is statistically uniform over $\mathbb{Z}_q^{m \times n'}$, so by Lemma 2.4, the quantity $[\mathbf{R}_{x^{(i)}}^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \mathbf{s})]_p$ is an injective function of $\mathbf{C} \cdot \mathbf{s} \bmod q$. It comes that partial evaluation queries information-theoretically reveal $\mathbf{C} \cdot \mathbf{s} \bmod q$, but we will show that $\mathbf{s}$ still retains high entropy in $\mathcal{A}$'s view.

Game₅: We modify the challenge value for which, if $b = 1$, the adversary is given a random $\mathbf{y}^* \leftarrow U(\mathbb{Z}_u^k)$. Clearly, we have $\Pr[W_5] = 1/2$ since the distribution of the challenge value does not depend on $b \in \{0, 1\}$. Moreover, we will show that $|\Pr[W_5] - \Pr[W_4]| \leq 2^{-\Omega(\lambda)}$.

We claim that, conditionally on $\mathcal{A}$'s view, the vector $\mathbf{y}^*$ is already statistically uniform over $\mathbb{Z}_u^k$ in Game₄. Recall that the source $[\tilde{\mathbf{z}}^*]_u$ depends on an injective function $\mathbf{G}^\top \cdot \mathbf{s} \in \mathbb{Z}_q^m$ of the vector $\mathbf{s}$. In Lemma 4.4, we show that this vector has high min-entropy if BAD does not occur.

We observe that the source $[\tilde{\mathbf{z}}^*]_u$ can be written

$$[\tilde{\mathbf{z}}^*]_u = \left[\left[\left(\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G} \right)^\top \cdot \mathbf{s} \right]_p \right]_u \quad (25)$$

$$\begin{aligned} &= \left[\left(\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G} \right)^\top \cdot \mathbf{s} \right]_u + \mathbf{e}_{s,x,u} && \text{with } \mathbf{e}_{s,x,u} \in \{-1, 0\}^m \\ &= [\mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{s}]_u + [\mathbf{G}^\top \cdot \mathbf{s}]_u + \mathbf{e}_{s,x,u} + \mathbf{e}_{s,x}, && \text{with } \mathbf{e}_{s,x} \in \{0, 1\}^m \\ &= [\mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{s}]_u + [\mathbf{G}^\top \cdot \mathbf{s}]_u + \mathbf{e}'_{s,x}, && \text{with } \mathbf{e}'_{s,x} \in \{-1, 0, 1\}^m. \end{aligned} \quad (26)$$

The proof of Lemma 4.3 implies that $\lfloor \mathbf{R}_{x^*}^\top \cdot \mathbf{A}^\top \cdot \mathbf{s} \rfloor_p = \lfloor \mathbf{R}_{x^*}^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \cdot \mathbf{s} \rfloor_p$ with overwhelming probability. In turn, this implies that $H_\infty(\lfloor \mathbf{R}_{x^*}^\top \cdot \mathbf{A}^\top \cdot \mathbf{s} \rfloor_u \mid \mathbf{C} \cdot \mathbf{s}) = 0$ with high probability. In the expression of $\tilde{\mathbf{z}}^*$ in (26), we also remark that $\lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_u + \mathbf{e}'_{s,x}$ is an injective function of $\mathbf{s} \in \mathbb{Z}^n$. To see this, observe that

$$\lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_u + \mathbf{e}'_{s,x} = (u/q) \cdot \mathbf{G}^\top \cdot \mathbf{s}' - \mathbf{t}_{s,x} + \mathbf{e}'_{s,x}$$

for some $\mathbf{t}_{s,x} \in (0, 1)^m$, so that

$$(q/u) \cdot (\lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_u + \mathbf{e}'_{s,x}) = \mathbf{G}^\top \cdot \mathbf{s} + \mathbf{e}''_{s,x} \quad (27)$$

for some $\mathbf{e}''_{s,x} \in (-q/u, 2 \cdot q/u)^m$. The vector $\mathbf{s}$ is thus uniquely determined by (27) using the public trapdoor of $\mathbf{G}$ so long as $q/u \ll q$.

Consider the entropy of $\tilde{\mathbf{z}}^*$ conditionally on $\mathcal{A}$'s view. We have

$$\begin{aligned} H_\infty(\lfloor \tilde{\mathbf{z}}^* \rfloor_u \mid \mathbf{C} \cdot \mathbf{\Gamma}^\top, \{\mathbf{S}_{I_j}\}_{j \in \mathcal{C}^*}) \\ &= H_\infty(\lfloor \mathbf{R}_{x^*}^\top \cdot \mathbf{A}^\top \cdot \mathbf{s} \rfloor_u + \lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_u + \mathbf{e}'_{s,x} \mid \mathbf{C} \cdot \mathbf{\Gamma}^\top, \{\mathbf{S}_{I_j}\}_{j \in \mathcal{C}^*}) \\ &= H_\infty(\lfloor \mathbf{G}^\top \cdot \mathbf{s} \rfloor_u + \mathbf{e}'_{s,x} \mid \mathbf{C} \cdot \mathbf{\Gamma}^\top, \{\mathbf{S}_{I_j}\}_{j \in \mathcal{C}^*}) \\ &= H_\infty(\mathbf{s} \mid \mathbf{C} \cdot \mathbf{\Gamma}^\top, \{\mathbf{S}_{I_j}\}_{j \in \mathcal{C}^*}) \geq n \cdot \log \sigma - \frac{n}{2} \cdot \log e - n' \cdot \log q - 1. \end{aligned}$$

Here, the last inequality is given by Lemma 4.4. The second equality follows from the fact that, for any random variables X, Y, Z defined over an additive group, we have the equality $H_\infty(Y + Z \mid X) = H_\infty(Z \mid X)$ if $H_\infty(Y \mid X) = 0$.

In order to extract statistically random bits from $\tilde{\mathbf{z}}^*$, we must take into account that it possibly depends on x^* which may depend on $\mathbf{pp}$. As long as $P_K(X^*) = 0$, the source $\tilde{\mathbf{z}}^*$ is taken from a distribution determined by the challenge input $X^* \in \{0, 1\}^\ell$ within a collection of less than 2^ℓ distributions (namely, those inputs X for which $P_K(X) = 0$), which all have min-entropy $\bar{n} \geq n \log \sigma - \frac{n}{2} \cdot \log e - n' \log q - 1$. By applying Lemma 2.13 with $\varepsilon = 2^{-\lambda}$ for a collection $\mathcal{X}$ of at most $M = 2^\ell$ distributions, we obtain that the distribution of $\pi(\lfloor \tilde{\mathbf{z}}^* \rfloor_u)$ is $2^{-\Omega(\lambda)}$ -close to the uniform distribution over $\mathbb{Z}_u^k$. $\square$

Lemma 4.3. *Assume that $q/p > 2^{L+\lambda} \cdot r$, where $r = m^{L+2} \cdot n \cdot \beta^* \cdot \alpha q$ with $\beta^* = O(\beta \cdot \log N)$. Then, we have the inequality $|\Pr[W_4] - \Pr[W_3]| \leq \Pr[\text{BAD}] \leq 2^{-\Omega(\lambda)}$.*

Proof. We use the fact that, for each query $x = x_1 \dots x_L$, we have

$$\begin{aligned}
\mathbf{R}_x &= \mathbf{R}_0 \cdot \prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad + \mathbf{R}_{1,x_1} \cdot \prod_{i=2}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad + \mu_{1,x_1} \cdot \mathbf{R}_{2,x_2} \cdot \prod_{i=3}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad + \mu_{1,x_1} \mu_{2,x_2} \cdot \mathbf{R}_{3,x_3} \cdot \prod_{i=4}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \\
&\quad \vdots \\
&\quad + \mu_{1,x_1} \mu_{2,x_2} \dots \mu_{L-1,x_{L-1}} \cdot \mathbf{R}_{L,x_L}.
\end{aligned} \tag{28}$$

We define $\text{bad}_{x,j,\theta}$ as the event that the inequality (23) occurs for a specific string $x \in \{0,1\}^L$ and a specific pair $(j, \theta) \in [N] \times [d_j]$. We first proceed to bound $\Pr[\text{bad}_{x,j,\theta}]$. To this end, we apply Lemma 2.12 and first observe that, as argued in [7, proof of Theorem 7.3], the matrix $\mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top$ is within statistical distance $2^{-\Omega(\lambda)}$ from the uniform distribution over $\mathbb{Z}_q^{m \times n'}$, even conditionally on $\mathbf{R}_0^\top \cdot \mathbf{E} \in \mathbb{Z}^{m \times n}$. Hence, for a fixed choice of $x \in \{0,1\}^L$ and $(j, \theta) \in [N] \times [d_j]$, the product $\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \bar{\mathbf{s}}_{j,\theta})$ contains a term

$$\underbrace{\left(\prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i}) \right)^\top}_{\mathbf{G}_x^\top} \cdot \mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \bar{\mathbf{s}}_{j,\theta}),$$

for which we know that $\mathbf{G}_x^\top \cdot \mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top$ is statistically uniform over $\mathbb{Z}_q^{m \times n'}$ since $\mathbf{G}_x^\top$ is a $\mathbb{Z}_q$ -invertible matrix by construction.

Lemma 2.3 then implies that $\mathbf{G}_x^\top \cdot \mathbf{R}_0^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \bar{\mathbf{s}}_{j,\theta})$ is statistically close to the uniform distribution over $\mathbb{Z}_q^m$. By Lemma 2.2, we thus have

$$\Pr[\text{bad}_{x,j,\theta}] = \Pr \left[\left| \left(\mathbf{A} \cdot \mathbf{R}_x \right)^\top \cdot \bar{\mathbf{s}}_{j,\theta} \right|_p \neq \left| \left(\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C} \right) \cdot \bar{\mathbf{s}}_{j,\theta} \right|_p \right] \leq m \cdot \frac{2rp}{q}. \tag{29}$$

where $r = m^{L+2} \cdot n \cdot \beta^* \cdot \alpha q$, with $\beta^* = O(\beta \cdot \log N)$.

Indeed, we have $\|\mathbf{R}_x^\top\|_\infty \leq m^{L+2}$, $\|\bar{\mathbf{s}}_{j,\theta}\|_\infty \leq \beta^* = O(\beta \cdot \log N)$ and

$$\|\mathbf{E}\|_\infty = \max_{i \in [m]} \left(\sum_{j=1}^n |e_{i,j}| \right) \leq \sqrt{n} \max_{i \in [m]} \|\mathbf{e}_i\| \leq n \cdot \alpha q.$$

The conditions of Lemma 2.2 are satisfied since

$$\|\mathbf{R}_x^\top \cdot \mathbf{E} \cdot \bar{\mathbf{s}}_{j,\theta}\|_\infty \leq \|\mathbf{R}_x^\top\|_\infty \cdot \|\mathbf{E}\|_\infty \cdot \|\bar{\mathbf{s}}_{j,\theta}\|_\infty \leq m^{L+2} \cdot n \cdot \alpha q \cdot \beta^* = r,$$

which yields inequality (29).

By taking a union bound over all possible inputs $x \in \{0,1\}^L$ and all possible pairs $(j, \theta) \in [N] \times [d_j]$, we obtain the upper bound $\Pr[\text{BAD}] \leq 2^L \cdot m \cdot N \cdot d \cdot \frac{2rp}{q}$.

Since **Game**₃ and **Game**₄ are identical if **BAD** does not occur, we finally obtain

$$|\Pr[W_4] - \Pr[W_3]| \leq \Pr[\text{BAD}] \leq 2^L \cdot m \cdot N \cdot d \cdot \frac{2rp}{q} = 2m \cdot N \cdot d \cdot 2^{-\lambda},$$

which completes the proof of the lemma. $\square$

Lemma 4.4. *In **Game**₄, the min-entropy of $\mathbf{s} \in \mathbb{Z}^n$ conditionally on $\mathcal{A}$'s view is at least $n \cdot \log \sigma - \frac{n}{2} \cdot \log e - n' \cdot \log q - \frac{n}{2^n}$.*

Proof. Let us assume that **BAD** does not occur in **Game**₄ since, if it does, the challenger replaces the adversary's output with a random bit, in which case both games have the same outcome. We show that, assuming $\neg \text{BAD}$, the shared secret vector $\mathbf{s}$ retains high min-entropy conditionally on the adversary's view.

Let us first recap what the adversary can see in **Game**₄. For each partial evaluation query (j, X^*) , the response $\lfloor (\mathbf{A} \cdot \mathbf{R}_{x^*} + \mathbf{G})^\top \cdot \mathbf{S}_{I_j}^\top \rfloor_p$ consists of non-lossy functions of $\mathbf{S}_{I_j}^\top \in \mathbb{Z}^{n \times d_j}$. We thus consider partial evaluation queries of the form (j, X^*) as if they were corruption queries and assume that they information-theoretically reveal $\mathbf{S}_{I_j}$ (we thus merge the two sets $\mathcal{C}$ and $\mathcal{E}$ of Definition 2.22 into one set $\mathcal{C}^*$). As for uncorrupted shares $\{\mathbf{S}_{I_j}\}_{j \in [N] \setminus \mathcal{C}^*}$, partial evaluation queries $(j, X^{(i)})$ for which $X^{(i)} \neq X^*$ only reveal the information $\{\mathbf{C} \cdot \mathbf{S}_{I_j}^\top\}_{j \in [N] \setminus \mathcal{C}^*}$. More precisely, those partial evaluations $\{\mathbf{Y}_j\}_{j \in [N] \setminus \mathcal{C}^*}$ can be written

$$\mathbf{Y}_j = \lfloor (\mathbf{A}(x^{(i)}))^\top \cdot \mathbf{S}_{I_j}^\top \rfloor_p = \lfloor (\mathbf{R}_{x^{(i)}} \cdot \bar{\mathbf{A}}^\top \cdot \mathbf{C}) \cdot \mathbf{S}_{I_j}^\top \rfloor_p \quad (30)$$

where

$$\mathbf{S}_{I_j}^\top = \begin{bmatrix} \boldsymbol{\rho}_1^\top \\ \vdots \\ \boldsymbol{\rho}_n^\top \end{bmatrix} \cdot \mathbf{M}_{I_j}^\top \in \mathbb{Z}^{n \times d_j}$$

is a product of $\mathbf{M}_{I_j}^\top$ with the matrix $[\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n]^\top \in \mathbb{Z}^{n \times e}$ whose first column is the secret $SK_0 = \mathbf{s} \in \mathbb{Z}^n$. Hence, the information revealed by (30) for $j \in [N] \setminus \mathcal{C}^*$ is only a lossy function $\mathbf{C} \cdot \mathbf{S}_{I_j}^\top$ of the share $\mathbf{S}_{I_j}$, namely,

$$\begin{aligned} \mathbf{C} \cdot \begin{bmatrix} \boldsymbol{\rho}_1^\top \\ \vdots \\ \boldsymbol{\rho}_n^\top \end{bmatrix} \cdot \mathbf{M}_{I_j}^\top &= \left[\mathbf{C} \cdot \mathbf{s} \mid \mathbf{C} \cdot \begin{pmatrix} \rho_{2,2} \\ \vdots \\ \rho_{n,2} \end{pmatrix} \mid \dots \mid \mathbf{C} \cdot \begin{pmatrix} \rho_{2,e} \\ \vdots \\ \rho_{n,e} \end{pmatrix} \right] \cdot \mathbf{M}_{I_j}^\top, \\ &= \mathbf{C} \cdot \boldsymbol{\Gamma}^\top \cdot \mathbf{M}_{I_j}^\top, \end{aligned} \quad (31)$$

where

$$\boldsymbol{\Gamma} = [\boldsymbol{\rho}_1 \mid \dots \mid \boldsymbol{\rho}_n] = \begin{bmatrix} \mathbf{s}^\top \\ \hline \rho_{2,2} \dots \rho_{n,2} \\ \vdots \quad \ddots \quad \vdots \\ \rho_{2,e} \dots \rho_{n,e} \end{bmatrix} \in \mathbb{Z}^{e \times n}$$

is the matrix of Gaussian entries which is used to compute secret key shares

$$\mathbf{S}_{I_j} = \mathbf{M}_{I_j} \cdot \mathbf{\Gamma} \quad \forall j \in [N].$$

The information revealed by exposed shares $\{\mathbf{S}_{I_j}\}_{j \in \mathcal{C}^*}$ can thus be written

$$\mathbf{S}_{I_j} = [\mathbf{s}_{I_j,1} \mid \dots \mid \mathbf{s}_{I_j,n}] = \mathbf{M}_{I_j} \cdot \mathbf{\Gamma} \in \mathbb{Z}^{d_j \times n} \quad \forall j \in \mathcal{C}^*. \quad (32)$$

At this stage, we see that proving the following fact on distributions is sufficient to complete the proof of the lemma.

Fact *Let $M_{\mathcal{C}^*}$ to be the sub-matrix of $\mathbf{M}$ obtained by stacking up the rows assigned to corrupted parties $j \in \mathcal{C}^*$. Conditionally on*

$$(\mathbf{C}, \mathbf{C} \cdot \mathbf{\Gamma}^\top \cdot \mathbf{M}^\top, \mathbf{M}_{\mathcal{C}^*}, \mathbf{M}_{\mathcal{C}^*} \cdot \mathbf{\Gamma}), \quad (33)$$

the vector $\mathbf{s}^\top = (1, 0, \dots, 0)^\top \cdot \mathbf{\Gamma}$ has min-entropy at least

$$n \cdot \log \sigma - \frac{n}{2} \cdot \log e - n' \cdot \log q - \frac{n}{2^n}.$$

To prove this statement, we apply arguments inspired from [4, Lemma 1]. First, we observe that conditioning on (33) is the same as conditioning on $(\mathbf{C}, \mathbf{C} \cdot \mathbf{\Gamma}^\top \cdot \mathbf{M}_{[N] \setminus \mathcal{C}^*}^\top, \mathbf{M}_{\mathcal{C}^*}, \mathbf{M}_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$ since $\mathbf{M}_{\mathcal{C}^*} \cdot \mathbf{\Gamma}$ and $\mathbf{C}$ are given. In fact, it is sufficient to prove the result when conditioning on

$$(\mathbf{C}, \mathbf{C} \cdot \mathbf{\Gamma}^\top, \mathbf{M}_{\mathcal{C}^*}, \mathbf{M}_{\mathcal{C}^*} \cdot \mathbf{\Gamma}),$$

as $\mathbf{C} \cdot \mathbf{\Gamma}^\top \cdot \mathbf{M}_{[N] \setminus \mathcal{C}^*}^\top$ is computable from $\mathbf{C} \cdot \mathbf{\Gamma}^\top$. By the definition of an Integer Span Program, we know that there exists a sweeping vector $\boldsymbol{\kappa} \in \mathbb{Z}^e$ whose first coordinate is $\kappa_1 = 1$ and such that $\mathbf{M}_{\mathcal{C}^*} \cdot \boldsymbol{\kappa} = \mathbf{0}$. The rows of $\mathbf{M}_{\mathcal{C}^*}$ thus live in the lattice $\mathcal{L}_{\mathcal{C}^*} = \{\mathbf{m} \in \mathbb{Z}^e : \langle \mathbf{m}, \boldsymbol{\kappa} \rangle = 0\}$. Hence, if we define a matrix $L_{\mathcal{C}^*} \in \mathbb{Z}^{(e-1) \times e}$ whose rows form a basis of $\mathcal{L}_{\mathcal{C}^*}$, we may prove the min-entropy lower bound conditioned on

$$(\mathbf{C}, \mathbf{C} \cdot \mathbf{\Gamma}^\top, L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma}).$$

This is because $L_{\mathcal{C}^*} \cdot \mathbf{\Gamma}$ provides at least as much information as $\mathbf{M}_{\mathcal{C}^*} \cdot \mathbf{\Gamma}$.

We first consider the distribution of $\mathbf{\Gamma}$, conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$. Since the columns of $\mathbf{\Gamma}$ are statistically independent, we may look at them individually. For each $i \in [n]$, we let $\boldsymbol{\rho}_i^* \in \mathbb{Z}^e$ be an arbitrary solution of $L_{\mathcal{C}^*} \cdot \boldsymbol{\rho}_i^* = L_{\mathcal{C}^*} \cdot \boldsymbol{\rho}_i \in \mathbb{Z}_q^{e-1}$. The distribution of $\boldsymbol{\rho}_i \in \mathbb{Z}^e$ conditionally on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \boldsymbol{\rho}_i)$ is $\boldsymbol{\rho}_i^* + D_{\Lambda, \sigma, -\boldsymbol{\rho}_i^*}$, where $\Lambda = \{\mathbf{x} \in \mathbb{Z}^e \mid L_{\mathcal{C}^*} \cdot \mathbf{x} = \mathbf{0}\}$ is the 1-dimensional lattice $\Lambda = \boldsymbol{\kappa} \cdot \mathbb{Z}$.

At this stage, we know that conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$, each row $(s_i, \rho_{i,2}, \dots, \rho_{i,e})^\top$ of $\mathbf{\Gamma}^\top$ is Gaussian over an affine line. We use this observation to show that conditioning on $(\mathbf{C}, \mathbf{C} \cdot \mathbf{\Gamma}^\top, L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$ is the same as conditioning on $(\mathbf{C}, \mathbf{C} \cdot \mathbf{s}, L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$.⁸ In fact, we

⁸ Note that conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$, the *rows* of $\mathbf{\Gamma}^\top$ are Gaussian on affine lines, but a *column* of $\mathbf{\Gamma}^\top$ is an inner product of unit vector with all these rows.

claim that, conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$, the last $e - 1$ columns of $\mathbf{\Gamma}^\top$ do not reveal any more information than its first column. Indeed, conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$, each $\boldsymbol{\rho}_i$ can be written $\boldsymbol{\rho}_i = \xi_i \cdot \boldsymbol{\kappa} + \boldsymbol{\rho}_i^*$ for some integer $\xi_i \in \mathbb{Z}$. We may assume that the shifting vector $\boldsymbol{\rho}_i^* = (\rho_{i,1}^*, \dots, \rho_{i,e}^*)^\top \in \mathbb{Z}_q^e$ is known to $\mathcal{A}$ as it can be obtained from $L_{\mathcal{C}^*} \cdot \boldsymbol{\rho}_i$ via de-randomized Gaussian elimination. Writing $\boldsymbol{\kappa} = (\kappa_1, \dots, \kappa_e)$, the j -th column $(\mathbf{\Gamma}^\top)_j$ of $\mathbf{\Gamma}^\top$ is

$$(\mathbf{\Gamma}^\top)_j = \kappa_j \cdot \begin{pmatrix} \xi_1 \\ \vdots \\ \xi_n \end{pmatrix} + \begin{pmatrix} \rho_{1,j}^* \\ \vdots \\ \rho_{n,j}^* \end{pmatrix} \quad \forall j \in [e].$$

As $\kappa_1 = 1$, we have

$$(\mathbf{\Gamma}^\top)_j = \kappa_j \cdot (\mathbf{\Gamma}^\top)_1 - \kappa_j \cdot \begin{pmatrix} \rho_{1,1}^* \\ \vdots \\ \rho_{n,1}^* \end{pmatrix} + \begin{pmatrix} \rho_{1,j}^* \\ \vdots \\ \rho_{n,j}^* \end{pmatrix} \quad \forall j \in [e].$$

In the latter, the last two terms are information-theoretically known to $\mathcal{A}$ (once we have conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$) and so is κ_j .

We now study the distribution of $\mathbf{s} = (\mathbf{\Gamma}^\top)_1$ conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$. By statistical independence, we may consider each coordinate $s_i = (1, 0, \dots, 0)^\top \cdot \boldsymbol{\rho}_i$ of $\mathbf{s}$ individually. Recall that, conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$, each $\boldsymbol{\rho}_i$ is distributed as $\boldsymbol{\rho}_i^* + D_{\boldsymbol{\kappa}\mathbb{Z}, \sigma, -\boldsymbol{\rho}_i^*}$. Let us write $\boldsymbol{\rho}_i^* = y \cdot \boldsymbol{\kappa} + (\boldsymbol{\rho}_i^*)^\perp$, with $y \in \mathbb{R}$ and $(\boldsymbol{\rho}_i^*)^\perp$ orthogonal to $\boldsymbol{\kappa}$. Then,

$$\begin{aligned} \boldsymbol{\rho}_i^* + D_{\boldsymbol{\kappa}\mathbb{Z}, \sigma, -\boldsymbol{\rho}_i^*} &= (\boldsymbol{\rho}_i^*)^\perp + y \cdot \boldsymbol{\kappa} + D_{\boldsymbol{\kappa}\mathbb{Z}, \sigma, -y \cdot \boldsymbol{\kappa} - (\boldsymbol{\rho}_i^*)^\perp} \\ &= (\boldsymbol{\rho}_i^*)^\perp + y \cdot \boldsymbol{\kappa} + \boldsymbol{\kappa} \cdot D_{\mathbb{Z}, \sigma / \|\boldsymbol{\kappa}\|, -y}. \end{aligned}$$

We now take the inner product with $(1, 0, \dots, 0)$ and use the fact that $\kappa_1 = 1$ to obtain that, conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$, the coordinate s_i is distributed as $(\boldsymbol{\rho}_i^*)^\perp_1 + y + D_{\mathbb{Z}, \sigma / \|\boldsymbol{\kappa}\|, -y}$. As $\boldsymbol{\kappa} \in \{-1, 0, 1\}^e$ with the Benaloh-Leichter-based LISS scheme of [26], and by our choice of σ , we have that $\sigma / \|\boldsymbol{\kappa}\| = \Omega(\sqrt{n})$. Using Lemma 2.6, this implies that each s_i has min-entropy $\geq \log(\sigma / \|\boldsymbol{\kappa}\|) - 2^{-n} \geq \log \sigma - \frac{1}{2} \log e - 2^{-n}$. Overall, we obtain

$$H_\infty(\mathbf{s} \mid L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma}) \geq n \cdot \log \sigma - \frac{n}{2} \cdot \log e - \frac{n}{2^n}.$$

We are now ready to conclude. By the above, to prove the fact (and hence the lemma), it suffices to obtain a lower bound on the min-entropy of $\mathbf{s}$ conditioned on $(\mathbf{C}, \mathbf{C} \cdot \mathbf{s}, L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$. We then use the above min-entropy lower bound on $\mathbf{s}$ conditioned on $(L_{\mathcal{C}^*}, L_{\mathcal{C}^*} \cdot \mathbf{\Gamma})$ and the fact that given $\mathbf{C}$, the quantity $\mathbf{C} \cdot \mathbf{s} \in \mathbb{Z}_q^{n'}$ reveals at most $n' \log q$ bits. $\square$

4.3 Reducing the Number of Matrices in Public Parameters

As described in Section 4.1, our construction requires large public parameters $\mathbf{pp}$ containing $\Theta(\lambda)$ random matrices $\{\mathbf{A}_{i,0}, \mathbf{A}_{i,1}\}_{i=1}^L$. In order to reduce the number of public matrices from

$L = \Theta(\lambda)$ to $\Theta(\log^c \lambda)$, for some constant $c > 0$, it is possible to use a technique suggested by Yamada [68] modulo slight changes in our security proof.

The proof of Theorem 4.2 uses an admissible hash function keyed by $K \in \{0, 1, \perp\}^L$ to encode the input X as $x = \text{AHF}(X) \in \{0, 1\}^L$, which is mapped to a matrix of the form $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x + P_K(x) \cdot \mathbf{G}$, where $P_K : \{0, 1\}^L \rightarrow \{0, 1\}$ is the partitioning function of Definition 2.9. Recall that the latter should evaluate to 1 in the challenge input x^* and 0 in partial evaluation queries $x \neq x^*$. In the transition from **Game**₃ to **Game**₄, the proof of Theorem 4.2 requires the inequality (23) to occur with negligible probability in partial evaluation queries. In order to prove this in Lemma 4.3, we exploit the fact that the matrix $\mathbf{R}_x$ is of the form (28) to show that, for a lossy matrix $\mathbf{A}^\top = \mathbf{A}^\top \cdot \mathbf{C} + \mathbf{E}$, the first term $\mathbf{R}_x^\top \cdot \mathbf{A}^\top \cdot (\mathbf{C} \cdot \mathbf{s}_{j,\theta})$ of $\mathbf{A}(x)^\top \cdot \mathbf{s}_{j,\theta}$ is statistically uniform because it contains a term which is a product of a statistically random vector with a product $\prod_{i=1}^L \mathbf{G}^{-1}(\mathbf{A}_{i,x_i})$ of $\mathbb{Z}_q$ -invertible binary matrices. This invertibility is enforced by choosing $\{\mathbf{A}_{i,b}\}$ in such a way that $\{\mathbf{G}^{-1}(\mathbf{A}_{i,b})\}$ are invertible. In the proof of Lemma 4.1, a similar argument invokes the invertibility of matrices $\mathbf{G}^{-1}(\mathbf{A}_{i,b})$ to prove that $\mathbf{A}^\top(x) \cdot \mathbf{s}$ is close to uniform, which is necessary to prove the correctness of the scheme by applying Lemma 2.2.

In order to use Yamada's partitioning technique [68], some changes are needed in the proof of Lemma 4.3. The constructions of [68] rely on the observation that the key K of the admissible hash function of [45] is very sparse (with only $O(\log \lambda)$ non- $\perp$ entries) and can be encoded using $O(\log \lambda)$ integers of size $O(\log \lambda)$ each. Letting $\mathbf{T}$ denote the compact representation of K , Yamada defines a modified partitioning function $P'_\mathbf{T}(\cdot)$, which is equivalent to $P_K(\cdot)$ and operates over $\mathbf{T}$. The homomorphic operations of GSW then make it possible to evaluate $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}'_x + P'_\mathbf{T}(x) \cdot \mathbf{G}$ from GSW encryptions⁹ of the bits of $\mathbf{T}$.

Lemma 4.5 ([68]). *Let $\mathcal{K}_{\text{MAH}} = \{\mathbf{T} \subseteq [2\ell] \mid |\mathbf{T}| < \eta\}$ and $\mathcal{X} = \{0, 1\}^\ell$. For any key $K \in \{0, 1, \perp\}^L$ with at most $\eta = O(\log \lambda)$ entries in $\{0, 1\}$, define $\mathbf{T} = \text{Encode}_{\text{MAH}}(K)$ as*

$$\mathbf{T} := \{2i - K_i \mid i \in [L], K_i \neq \perp\} \in \mathcal{K}_{\text{MAH}}.$$

Then, the partitioning function

$$P'_\mathbf{T}(X) := \begin{cases} 0 & \text{if } \mathbf{T} \subseteq \mathbf{S}(X) \\ 1 & \text{otherwise} \end{cases} \quad \text{where } \mathbf{S}(X) = \{2i - \text{AHF}(X)_i \mid i \in [L]\}$$

satisfies $P_K(X) = 0$ if and only if $P'_\mathbf{T}(X) = 0$. Moreover, $\eta' := |\mathbf{T}| = O(\log \lambda)$, so that $\mathbf{T}$ costs $u = O(\log^2 \lambda)$ bits to represent. Finally, there exist deterministic polynomial-time algorithms ($\text{PubEval}_{\text{MAH}}, \text{TrapEval}_{\text{MAH}}$) where $\text{PubEval}_{\text{MAH}}$ (resp. $\text{TrapEval}_{\text{MAH}}$) takes as input $X \in \{0, 1\}^\ell$ and $\{\mathbf{A}_i = \mathbf{A} \cdot \mathbf{R}_i + \kappa_i \cdot \mathbf{G}\}_{i \in [u]}$ (resp. X and $\{\mathbf{R}_i, \kappa_i\}_{i \in [u]}$), where $u = O(\log^2 \lambda)$ and $\kappa_i \in \{0, 1\}$, $\mathbf{R}_i \in \{-1, 0, 1\}^{m \times m}$ for all $i \in [u]$. It outputs a matrix $\mathbf{A}_{P,X} = \text{PubEval}_{\text{MAH}}(X, \{\mathbf{A}_i\}_{i \in [u]})$ (resp. $\mathbf{R}_{P,X} = \text{TrapEval}_{\text{MAH}}(X, \{\kappa_i, \mathbf{R}_i\}_{i \in [u]})$) such that

$$\mathbf{A}_{P,X} = \mathbf{A} \cdot \mathbf{R}_{P,X} + P'_\mathbf{T}(X) \cdot \mathbf{G} \in \mathbb{Z}_q^{n \times m} \quad (34)$$

and $\mathbf{R}_{P,X} \in \mathbb{Z}^{m \times m}$ has norm $\|\mathbf{R}_{P,X}\|_\infty \leq m^3 u (L + 1)$.

⁹ The proof of Yamada's IBE [68] requires to have $P'_\mathbf{T}(x^*) = 0$ in the challenge phase and $P'_\mathbf{T}(x) = 1$ in all adversarial queries while we need the opposite. This is not a problem here since we can just evaluate the logical NOT of his partitioning function $P'_\mathbf{T}$.

In our setting, one difficulty is that the matrix $\mathbf{R}_{P,X}$ of (34) is no longer of the form (28). When we run $\text{TrapEval}_{\text{MAH}}$, we have little control on the $\mathbb{Z}_q$ -invertibility of matrices produced by $\mathbf{G}^{-1}(\cdot)$ as the homomorphic evaluation of $P'_T(\cdot)$ proceeds via several recursive multiplication steps (see [68, Section 5.2] for details) where the decomposed matrices depend on the input x . A solution to this problem is to define $\mathbf{A}(x) = \mathbf{A}_0 + \text{PubEval}_{\text{MAH}}(X, \{\mathbf{A}_i\}_{i \in [u]})$ using an additional random matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times m}$ in the scheme. In addition, it also allows proving correctness since $\mathbf{A}_0^\top \cdot \mathbf{s}$ is already close to uniform, so that it is no longer necessary to rely on the invertibility of $\mathbf{G}^{-1}(\mathbf{A}_{i,b})$. In the security proof, we can generate $\mathbf{A}_0$ as a GSW encryption of 0 in such a way that $\mathbf{A}(x)$ can still be seen as a GSW encryption $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}_x + P'_T(x) \cdot \mathbf{G}$ of the partitioning function. Besides, the proof of Lemma 4.3 goes through when we set $\mathbf{A}_0 = \mathbf{A} \cdot \mathbf{R}_0$ for a random $\mathbf{R}_0 \leftarrow U(\{-1, 0, 1\}^{m \times m})$ since $\mathbf{R}_x = \mathbf{R}_0 + \mathbf{R}_{P,X}$ contains a term $\mathbf{R}_0$ that never gets multiplied by outputs of $\mathbf{G}^{-1}(\cdot)$, which allows applying the same arguments when we want to show that $\mathbf{R}_x^\top \cdot \bar{\mathbf{A}}^\top \cdot (\mathbf{C} \cdot \mathbf{s}_{j,\theta})$ is statistically uniform over $\mathbb{Z}_q^m$.

5 Achieving Robustness from Homomorphic Signatures

In this section, we show how to extend our DPRF so as to achieve robustness against faulty servers at the expense of assuming reliable erasures on behalf of the servers. Note that relying on erasures is common in the context of adaptive corruptions (see, e.g., [20]). To achieve robustness, we follow the approach of [16,15] which relies on homomorphic signatures to obtain a mechanism allowing to verify that servers' computations are carried out correctly.

In short, a leveled homomorphic signature makes it possible to publicly derive a signature Φ_C on a circuit evaluation $C(M)$ given a signature Φ on an initial message M . In our setting as well as in [16,15], we need a homomorphic signature which is *context-hiding*, meaning that a homomorphically evaluated signature Φ_C is statistically independent of the initial message M : more formally, there should be a simulator that creates Φ_C from scratch from the signing key sk and the value $C(M)$ and, yet, the joint distribution of $(C(M), \Phi_C)$ should be statistically close to that obtained by homomorphically deriving Φ_C from (M, Φ) . One difficulty is that the fully homomorphic signatures of [40] are not known to be simultaneously context-hiding and adaptively unforgeable (i.e., context-hiding security was only achieved for selectively unforgeable schemes). Fortunately, selective unforgeability is sufficient for our purposes (and those of [16]). We can thus instantiate our construction using the context-hiding construction of Gorbunov, Vaikuntanathan and Wichs [40], which relies on the Short-Integer-Solution (SIS) assumption. Since the LWE assumption implies the SIS assumption, we can achieve robustness without any additional assumption, analogously to [16,15].

A Robust DPRF is a DPRF endowed with a verification algorithm Robust.Verify that takes as input the public parameters $\mathbf{pp}$ of the DPRF, some public part vk_j of the shared key, along with an input X and a candidate partial evaluation Y_j on behalf of server $j \in [N]$, together with a corresponding label ϕ_j . It outputs 1 or 0 depending on whether (Y_j, ϕ_j) is deemed valid or not. For correctness, we require that $\text{Robust.Verify}(\mathbf{pp}, Y_j, \phi_j, vk_j, X) = 1$ for all $j \in [N]$, $\mathbf{pp} \leftarrow \text{Robust.Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$, $X \in \mathcal{X}$ and $(Y_j, \phi_j) \leftarrow \text{Robust.PEval}(\mathbf{pp}, SK_j, X)$, where $((SK_1, vk_1), \dots, (SK_N, vk_N)) \leftarrow \text{Robust.Share}(SK_0)$ is obtained by sharing the key

$SK_0 \leftarrow \mathcal{K}$ and $\{vk_j\}_{j \in [N]}$ are the public parts of the corresponding shares. The robustness property requires that it be infeasible for a corrupted server $j \in [N]$ holding the secret share SK_j to produce an incorrect partial evaluation Y_j , for some input X , that still satisfies $\text{Robust.Verify}(\text{pp}, Y_j, \phi_j, vk_j, X) = 1$.

Definition 5.1 (Robust DPRF). *A robust DPRF specified by algorithms $(\text{Robust.Setup}, \text{Robust.Share}, \text{Robust.PEval}, \text{Robust.Eval}, \text{Robust.Combine})$ together with a polynomial-time verification algorithm Robust.Verify such that:*

Correctness: *For any $\text{pp} \leftarrow \text{Robust.Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$, any key $SK_0 \leftarrow \mathcal{K}$, any shares $((SK_1, vk_1), \dots, (SK_N, vk_N)) \leftarrow \text{Robust.Share}(SK_0)$, any index $j \in [N]$, any $X \in \mathcal{X}$, and any $(Y_j, \phi_j) \leftarrow \text{Robust.PEval}(\text{pp}, SK_j, X)$ we have $\text{Robust.Verify}(\text{pp}, j, Y_j, \phi_j, vk_j, X) = 1$.*

Robustness: *For any $j \in [N]$ and any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\lambda)$ such that:*

$$\Pr \left[\text{Robust.Verify}(\text{pp}, j, \tilde{Y}_j, \tilde{\phi}_j, vk_j, \tilde{X}) = 1 \mid \begin{array}{c} (\tilde{Y}_j, \tilde{\phi}_j, \tilde{X}) \leftarrow \mathcal{A}(\text{pp}, \mathbf{SK}) \\ \wedge \\ \tilde{Y}_j \neq Y_j \end{array} \right] \leq \text{negl}(\lambda)$$

where the probability is taken over the random coins of $\text{pp} \leftarrow \text{Robust.Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$, $SK_0 \leftarrow \mathcal{K}$, $\mathbf{SK} := ((SK_1, vk_1), \dots, (SK_N, vk_N)) \leftarrow \text{Robust.Share}(SK_0)$, $(Y_j, \phi_j) \leftarrow \text{Robust.PEval}(SK_j, \tilde{X})$, and the coin tosses of $\mathcal{A}$.

5.1 Homomorphic Signatures

In this section we recall the definition of leveled homomorphic signature [40]. We use the simplified version presented in [16].

Definition 5.2 (Context-Hiding Homomorphic Signature). *A (leveled) context-hiding homomorphic signature scheme is a tuple of efficient algorithms $(\text{KeyGen}, \text{Sign}, \text{SignEval}, \text{Verify}, \text{Hide}, \text{HVerify})$ with the following specifications.*

KeyGen $(1^\lambda, 1^d, 1^K) \rightarrow (sk, vk)$: *On input the security parameter λ , a circuit depth bound d and a data set bound K , the algorithm outputs a signing key and a verification key.*

Sign $(sk, \mathbf{m}) \rightarrow \Phi$: *On input of the signing key sk and a message $\mathbf{m} \in \{0, 1\}^K$, the algorithm outputs a signature Φ .*

SignEval $(C, \Phi) \rightarrow \Phi^*$: *On input of a circuit $C : \{0, 1\}^K \rightarrow \{0, 1\}$ and a signature, the algorithm outputs a homomorphically evaluated signature Φ^* .*

Verify $(vk, C, y, \Phi^*) \rightarrow 0/1$: *Given a verification key vk , a circuit $C : \{0, 1\}^K \rightarrow \{0, 1\}$, an output value y and a signature Φ^* , the algorithm outputs 1 or 0.*

Hide $(vk, y, \Phi) \rightarrow \tilde{\Phi}$: *On input of a verification key vk , an output value y , and a signature Φ , this (randomized) algorithm outputs a transformed signature $\tilde{\Phi}$.*

HVerify $(vk, C, y, \tilde{\Phi}) \rightarrow 0/1$: *Given a verification key vk , a circuit $C : \{0, 1\}^K \rightarrow \{0, 1\}$, an output value y and a transformed signature $\tilde{\Phi}$, the algorithm outputs 1 or 0.*

Correctness: For all $\lambda, d, K \in \mathbb{N}$, $(sk, vk) \leftarrow \text{KeyGen}(1^\lambda, 1^d, 1^K)$, $\mathbf{m} \in \{0, 1\}^K$, $\Phi \leftarrow \text{Sign}(sk, \mathbf{m})$, $C : \{0, 1\}^K \rightarrow \{0, 1\}$ a circuit of depth at most d and $y = C(\mathbf{m})$, if $\Phi^* \leftarrow \text{SignEval}(C, \Phi)$, then $\text{Verify}(vk, C, y, \Phi^*) = 1$. Moreover, for any (C, y, Φ) such that $\text{Verify}(vk, C, y, \Phi) = 1$ and any $\tilde{\Phi} \leftarrow \text{Hide}(vk, y, \Phi)$, we have $\text{HVerify}(vk, C, y, \tilde{\Phi}) = 1$ w.h.p. over the randomness of Hide .

For our purposes, we need a homomorphic signature that satisfies two security definitions called *unforgeability* and *context-hiding* [40].

Definition 5.3 (Selective Unforgeability). *We say that a leveled context-hiding homomorphic signature scheme $\text{HS} = (\text{KeyGen}, \text{Sign}, \text{SignEval}, \text{Verify}, \text{Hide}, \text{HVerify})$ provides selective unforgeability if, for any PPT adversary $\mathcal{A}$, there is a negligible function $\text{negl}(\lambda)$ such that*

$$\text{Adv}_{\text{HS}, \mathcal{A}}^{\text{uf}}(\lambda) := \Pr[\text{Exp}_{\text{HS}, \mathcal{A}}^{\text{uf}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where $\text{Exp}_{\text{HS}, \mathcal{A}}^{\text{uf}}(\lambda)$ is defined below:

1. $(sk, vk) \leftarrow \text{KeyGen}(1^\lambda, 1^d, 1^K)$
2. $\mathbf{m}^* \leftarrow \mathcal{A}(1^\lambda)$
3. $\Phi \leftarrow \text{Sign}(sk, \mathbf{m}^*)$
4. $(C^*, y^*, \tilde{\Phi}^*) \leftarrow \mathcal{A}(\Phi, vk)$

The experiment outputs 1 if and only if the following conditions are satisfied:

- The depth of C^* is at most d
- $C^*(\mathbf{m}^*) \neq y^*$
- $\text{HVerify}(vk, C^*, y^*, \tilde{\Phi}^*) = 1$.

We need a homomorphic signature satisfying a notion of context-hiding security capturing that derived signatures hide everything but the result $C(\mathbf{m})$ of the circuit evaluation. This notion is achieved via the Hide algorithm that takes as input a derived signature Φ (which is not context hiding) and applies some sanitization procedure that makes the resulting signature context hiding. More precisely, the “hiding” signature $\tilde{\Phi}$ can be simulated given only C and $y = C(\mathbf{m})$, no matter which original signature was used to create it. Our Definition 5.4 hereunder is formulated in a slightly different way than (and is actually implied by) [40, Definition 6.1]. We chose this formulation to make it explicit that indistinguishability holds even when the original signatures are given.

Definition 5.4 (Context-hiding). *A leveled homomorphic signature $\text{HS} = (\text{KeyGen}, \text{Sign}, \text{SignEval}, \text{Verify}, \text{Hide}, \text{HVerify})$ is statistically context-hiding if there exists a simulator Sim such that no (computationally unbounded) adversary has noticeable advantage in the following game:*

1. The challenger flips a coin $b \leftarrow U(\{0, 1\})$, generates $(sk, vk) \leftarrow \text{KeyGen}(1^\lambda, 1^d, 1^K)$ and gives (sk, vk) to the adversary $\mathcal{A}$.

2. On polynomially-many occasions, $\mathcal{A}$ submits adaptive queries, each of which consists of a message y , a circuit $C : \{0, 1\}^K \rightarrow \{0, 1\}$, and a signature Φ . If $\text{HVerify}(vk, C, y, \Phi) = 0$, the challenger returns $\perp$. Otherwise, the challenger responds as follows:
 - If $b = 0$, the challenger evaluates $\tilde{\Phi} \leftarrow \text{Hide}(vk, y, \Phi)$ and returns $\tilde{\Phi}$.
 - If $b = 1$, the challenger computes $\tilde{\Phi} \leftarrow \text{Sim}(sk, C, y)$ and returns $\tilde{\Phi}$ to $\mathcal{A}$.
3. Finally, the adversary outputs a bit $b' \in \{0, 1\}$ and its advantage is measured by the distance $\text{Adv}_{\mathcal{A}}^{\text{ch}}(\lambda) := |\Pr[b' = b] - 1/2|$.

The context-hiding construction of Gorbunov *et al.* [40, Section 6] is easily seen to satisfy the above security definition.

5.2 A Robust DPRF Construction

In this section, we use homomorphic signatures to make our DPRF construction of Section 4 robust against malicious servers. For this purpose, we need a leveled homomorphic signatures [40] that is context-hiding. Following the approach of [16,15], we have the trusted dealer provide each server $j \in [N]$ with a secret key share SK_j together with a homomorphic signature Φ_j of that secret share. By leveraging the homomorphism of the signature scheme, each server j can derive (without knowing the signing key) valid signatures for circuits C_X that evaluate $Y_j = \text{PEval}(SK_j, X)$ on input of SK_j . Each server can thus produce a valid signature for the partial evaluation of the PRF on any input X using its own share SK_j and the corresponding signature. The unforgeability of homomorphic signatures makes it infeasible to come up with a valid signature on an incorrect partial evaluation of the PRF for the input X .

Let a DPRF specified by algorithms (DPRF.Setup , DPRF.Share , DPRF.PEval , DPRF.Eval , DPRF.Combine), as described in Section 4. Let a homomorphic signature scheme (HS.KeyGen , HS.Sign , HS.SignEval , HS.Verify , HS.Hide , HS.HVerify). We now describe a robust DPRF as follows.

Robust.Setup($1^\lambda, 1^\ell, 1^t, 1^N$): Run $\text{pp} \leftarrow \text{DPRF.Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$. Define the input space $\mathcal{X}$ and the key space $\mathcal{K}$ to be the same as in the underlying DPRF. Let K the number of bits to represent any element of $\mathcal{K}$ and d is the maximal depth of any Boolean circuit that computes

$$C_X(\cdot) := \text{DPRF.PEval}(\cdot, X)$$

for any input $X \in \mathcal{X}$. Return $\text{Robust.pp} = \{\text{pp}, \mathcal{K}, \mathcal{X}, 1^d\}$.

Robust.Share: In order to share $SK_0 \leftarrow \mathcal{K}$, run $(sk_j, vk_j) \leftarrow \text{HS.KeyGen}(1^\lambda, 1^d, 1^K)$ for every $j \in [N]$ and do the following.

1. Run $(SK_1, \dots, SK_N) \leftarrow \text{DPRF.Share}(SK_0)$.
2. Compute $\Phi_j = \text{HS.Sign}(sk_j, SK_j)$ for all $j \in [N]$.

For each $j \in [N]$, the j -th share comprises the secret pair (SK_j, Φ_j) and the signature verification key vk_j , which is made publicly available.

Robust.PEval $((SK_j, \Phi_j), X)$: Given its secret key share (SK_j, Φ_j) , the j -th server computes a partial evaluation on X by conducting the following steps:

1. Compute $Y_j \leftarrow \text{DPRF.PEval}(SK_j, X)$.
 2. Compute $\Phi_{j,X} \leftarrow \text{HS.SignEval}(C_X, \Phi_j)$, where C_X is a Boolean circuit such that $C_X(\cdot) = \text{DPRF.PEval}(\cdot, X)$.
 3. Compute $\tilde{\Phi}_{j,X} \leftarrow \text{HS.Hide}(vk_j, Y_j, \Phi_{j,X})$ and erase the random coins used by HS.Hide .
- Then, return $(Y_j, \tilde{\Phi}_{j,X})$.

Robust.Eval (SK_0, X) : Compute $Y \leftarrow \text{DPRF.Eval}(SK_0, X)$.

Robust.Combine $(\mathcal{S}, (Y_{j_1}, \dots, Y_{j_t}))$: On input a t -subset of $[N]$ and t partial evaluations $(Y_{j_1}, \dots, Y_{j_t})$, output $Y \leftarrow \text{DPRF.Combine}(\mathcal{S}, (Y_{j_1}, \dots, Y_{j_t}))$.

Robust.Verify $(\text{Robust.pp}, j, (Y_j, \tilde{\Phi}_{j,X}), vk_j, X)$: If $\text{HS.HVerify}(vk_j, C_X, Y_j, \tilde{\Phi}_{j,X}) = 1$, return 1. Otherwise, return 0.

The correctness of the **Robust.Verify** algorithm follows from the correctness of the homomorphic signature scheme.

We remark that, in the partial evaluation algorithm **Robust.PEval**, the evaluated circuit C_X only needs to have small depth when the above construction is applied to our scheme of Section 4.1. Indeed, since the matrix $\mathbf{A}(x)$ is public, it can be hard-wired in C_X , which only takes as input the bits of a secret key share in order to compute a matrix-vector product followed by a rounding operation.

The robustness of the DPRF is implied by the selective unforgeability of the homomorphic signature. From a robustness adversary, we can easily construct a selective forger in the sense of Definition 5.3 as the only way for the adversary to trick the verification algorithm into accepting an incorrect partial evaluation is to break the unforgeability of the homomorphic signature. Moreover, the latter only needs to provide selective unforgeability since, in the reduction, the adversary only needs to see one verification key vk after the generation of signatures on secret key shares $\{SK_j\}_{j \in [N]}$.

Theorem 5.5. *Assuming that the homomorphic signature scheme provides selective unforgeability, the above construction of the Robust.DPRF achieves robustness.*

Proof. Towards a contradiction, suppose that there exists a polynomial-time adversary $\mathcal{A}$ that breaks the robustness for the index $j \in [N]$ (in the sense of Definition 5.1) of the Robust.DPRF. We build a polynomial-time adversary $\mathcal{B}$ that can forge signatures, thus breaking the selective unforgeability of the homomorphic signature HS.

At step 1 of the selective unforgeability game, the reduction $\mathcal{B}$ runs the underlying $\text{DPRF.Setup}(1^\lambda, 1^\ell, 1^t, 1^N)$ on input of parameters $(1^\lambda, 1^\ell, 1^t, 1^N)$ such that any key in $\mathcal{K}$ can be represented using K bits and, for any input $X \in \mathcal{X}$, the Boolean circuit $C_X(\cdot) := \text{DPRF.PEval}(\cdot, X)$ has depth at most d . Then, $\mathcal{B}$ is able to compute **Robust.pp** from the public parameters of the underlying DPRF before running $\mathcal{A}$. Before that, $\mathcal{B}$ samples a secret key $SK_0 \leftarrow \mathcal{K}$ and runs the underlying $\text{DPRF.Share}(SK_0)$. To finalize step 2, $\mathcal{B}$ sends $\mathbf{m}^* := SK_j$ to the challenger to sign. The challenger runs $(vk_j, sk_j) \leftarrow \text{HS.KeyGen}(1^\lambda, 1^d, 1^K)$, with sufficiently large parameters d, K and provides $\mathcal{B}$ with the public verification key vk_j .

together with $\Phi_j \leftarrow \text{HS.Sign}(sk_j, \mathbf{m}^*)$. Moreover, $\mathcal{B}$ runs $(sk_i, vk_i) \leftarrow \text{HS.KeyGen}(1^\lambda, 1^d, 1^K)$ for all $i \in [N] \setminus \{j\}$ and computes the rest of the signatures $\Phi_i \leftarrow \text{HE.Sign}(sk_i, SK_i)$, for all $i \in [N] \setminus \{j\}$. In turn, $\mathcal{B}$ gives $\mathcal{A}$ the vector $\mathbf{SK} := ((SK_1, \Phi_1), \dots, (SK_N, \Phi_N))$ which corresponds to all the secret shares of the Robust.DPRF scheme.

Since $\mathcal{A}$ can break robustness, by hypothesis, it has non-negligible probability of outputting $(\tilde{Y}_j, \tilde{\Phi}_j^*, \tilde{X})$ such that $\text{Robust.Verify}(\text{pp}, j, (\tilde{Y}_j, \tilde{\Phi}_j^*), vk_j, \tilde{X}) = 1$ and $\tilde{Y}_j$ is not obtained by running $\text{Robust.PEval}((SK_j, \phi_j), \tilde{X})$. By construction, this translates for the homomorphic signature scheme into:

$$\text{HS.Verify}(vk_j, C_{\tilde{X}}, \tilde{Y}_j, \tilde{\Phi}_j^*) = 1 \quad \wedge \quad \tilde{Y}_j \neq \text{DPRF.PEval}(SK_j, \tilde{X}),$$

where the second condition is equivalent to $\tilde{Y}_j \neq C_{\tilde{X}}(SK_j)$.

At step 4 of the selective unforgeability game, algorithm $\mathcal{B}$ outputs a triple $(C_{\tilde{X}}, \tilde{Y}_j, \tilde{\Phi}_j^*)$, which constitutes a forgery for the signature scheme since the verification algorithm accepts even though $C_{\tilde{X}}(\mathbf{m}^*) \neq \tilde{Y}_j$. The probability that $\mathcal{B}$ outputs a forgery is the same as the probability of $\mathcal{A}$ in breaking the robustness of the Robust.DPRF. $\square$

We still have to prove that above construction preserves the security of the underlying DPRF under adaptive corruptions. For this purpose, we need to assume that the servers can reliably erase the random coins of the Hide procedure after the computation of $\tilde{\Phi}_{j,X}$ at step 3 of PEval.

Theorem 5.6. *The above construction is an adaptively secure robust DPRF assuming that:*
(i) *The underlying DPRF is adaptively secure;* (ii) *The homomorphic signature scheme is context-hiding;* (iii) *Each server can reliably erase its random coins.*

Proof. To prove the result, we first consider two indistinguishable games where **Game**₀ is the real adaptive security experiment and **Game**₁ makes it easier to give a reduction from the security of the underlying DPRF.

Game₀: This is the real adaptive security experiment of Definition 2.22.

Game₁: This game is identical to **Game**₀ with a modification in the partial evaluation oracle at each partial evaluation query (j, X) . At step 3 of PEval, instead of computing $\tilde{\Phi}_{j,X} \leftarrow \text{HS.Hide}(vk_j, Y_j, \Phi_{j,X})$, the challenger computes $\tilde{\Phi}_{j,X} \leftarrow \text{Sim}(sk_j, C_X, Y_j)$. As long as the servers can erase their local randomness after each execution of PEval,¹⁰ we can meaningfully rely on the context-hiding security of HS to show that **Game**₁ and **Game**₀ are statistically indistinguishable. Note that the derived signature $\Phi_{j,X}$ obtained at step 2 of PEval is now independent of the adversary's view.

Now, assuming that there exists a PPT adversary $\mathcal{A}$ with noticeable advantage in **Game**₁, we build an efficient adversary $\mathcal{B}$ with the same advantage in the adaptive security game for the underlying (non-robust) DPRF.

¹⁰ Without this assumption, an adaptive adversary could distinguish the two games by corrupting a server after having obtained a partial evaluation from it.

Algorithm $\mathcal{B}$ runs $\mathcal{A}$ on input of the public parameters $\mathbf{pp}$ that it receives from its own challenger. Before answering $\mathcal{A}$'s queries, $\mathcal{B}$ generates N homomorphic signature key pairs $(sk_i, vk_i) \leftarrow \text{HS.KeyGen}(1^\lambda, 1^d, 1^K)$, for $i \in [N]$, where d is the maximum depth of the Boolean circuits $C_X(\cdot) = \text{PEval}(\cdot, X)$, for any input $X \in \mathcal{X}$. (The input space and the key space are known from the public parameters $\mathbf{pp}$). The verification keys vk_i , are also given to $\mathcal{A}$.

At each corruption query $j \in [N]$ made by $\mathcal{A}$, the reduction $\mathcal{B}$ relays the query to its DPRF challenger. Upon receiving SK_j from the latter, $\mathcal{B}$ computes $\Phi_j := \text{HS.Sign}(sk_j, SK_j)$ and hands (SK_j, Φ_j) to $\mathcal{A}$. At each partial evaluation query $(i, X) \in [N] \times \{0, 1\}^\ell$ made by $\mathcal{A}$, $\mathcal{B}$ sends the same query to its challenger and, upon receiving the response Y_i , it computes $\tilde{\Phi}_i^* = \text{Sim}(sk, C_X, Y_i)$ and returns the pair $(Y_i, \tilde{\Phi}_i^*)$ to $\mathcal{A}$.

For the challenge query X^* made by $\mathcal{A}$, $\mathcal{B}$ sends the same query to its own DPRF challenger and forwards the latter's challenge Y^* to $\mathcal{A}$. At the end of the game, $\mathcal{B}$ outputs whatever $\mathcal{A}$ outputs. It is easy to see that, by construction, $\mathcal{B}$ succeeds whenever $\mathcal{A}$ does. $\square$

Acknowledgements

We thank Javier Herranz for his suggestion to use linear integer secret sharing schemes. Part of this research was funded by the French ANR ALAMBIC project (ANR-16-CE39-0006) and by BPI-France in the context of the national project RISQ (P141580). This work was also supported in part by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701). The second author was supported by ERC Starting Grant ERC-2013-StG-335086-LATTAC.

References

1. M. Abe and S. Fehr. Adaptively secure Feldman VSS and applications to universally-composable threshold cryptography. In *Crypto*, 2004.
2. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Eurocrypt*, 2010.
3. S. Agrawal, X. Boyen, V. Vaikuntanathan, P. Voulgaris, and H. Wee. Functional encryption for threshold functions (or fuzzy ibe) from lattices. In *PKC*, 2012.
4. S. Agrawal, B. Libert, and D. Stehlé. Fully secure functional encryption for inner products from standard assumptions. In *Crypto*, 2016.
5. S. Agrawal, P. Mohassel, P. Mukherjee, and P. Rindal. DiSE: Distributed Symmetric-key Encryption. In *ACM-CCS*, 2018.
6. I. Almansa, I. Damgård, and J.-B. Nielsen. Simplified threshold RSA with adaptive and proactive security. In *Eurocrypt*, 2006.
7. J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs. Learning with rounding, revisited - new reduction, properties and applications. In *Crypto*, 2013.
8. A. Banerjee and C. Peikert. New and improved key-homomorphic pseudo-random functions. In *Crypto*, 2014.
9. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Eurocrypt*, 2012.
10. M. Bellare, E. Kiltz, C. Peikert, and B. Waters. Identity-based (lossy) trapdoor functions and applications. In *Eurocrypt*, 2012.
11. J. Benaloh and J. Leichter. Generalized secret sharing and monotone functions. In *Crypto*, 1988.
12. R. Bendlin and I. Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In *TCC*, 2010.
13. R. Bendlin, S. Krehbiel, and C. Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In *ACNS*, 2013.
14. D. Boneh and X. Boyen. Secure identity based encryption without random oracles. In *Crypto*, 2004.

15. D. Boneh, R. Gennaro, S. Goldfeder, A. Jain, S. Kim, P. Rasmussen, and A. Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In *Crypto*, 2018.
16. D. Boneh, R. Gennaro, S. Goldfeder, and S. Kim. A lattice-based universal thresholdizer for cryptographic systems. Cryptology ePrint Archive: Report 2017/251, Sept. 2017.
17. D. Boneh, K. Lewi, H. Montgomery, and A. Raghunathan. Key-homomorphic PRFs and their applications. In *Crypto*, 2013.
18. D. Boneh, H. Montgomery, and A. Raghunathan. Algebraic pseudorandom functions with improved efficiency from the augmented cascade. In *ACM-CCS*, 2010.
19. C. Cachin, K. Kursawe, and V. Shoup. Random oracles in constantinople: practical asynchronous byzantine agreement using cryptography. In *PODC*, 2000.
20. R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Adaptive security for threshold cryptosystems. In *Crypto*, 1999.
21. R. Canetti and S. Goldwasser. An efficient threshold public key cryptosystem secure against adaptive chosen-ciphertext attacks. In *Eurocrypt*, 1999.
22. M. Chase and S. Meiklejohn. Déjà Q: using dual systems to revisit q-type assumptions. In *Eurocrypt*, 2004.
23. J. Chotard, E. Dufour Sans, R. Gay, D.-H. Phan, and D. Pointcheval. Decentralized multi-client functional encryption for inner product. In *Asiacrypt*, 2018.
24. R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, and T. Rabin. Efficient multi-party computations secure against an adaptive adversary. In *Eurocrypt*, 1999.
25. R. Cramer and S. Fehr. Optimal black-box secret sharing over arbitrary abelian groups. In *Crypto*, 2002.
26. I. Damgård and R. Thorbek. Linear integer secret sharing and distributed exponentiation. In *PKC*, 2006.
27. Y. Desmedt and Y. Frankel. Threshold cryptosystems. In *Crypto*, 1989.
28. Y. Dodis. *Exposure-resilient cryptography*. PhD thesis, MIT, 2000.
29. Y. Dodis. Efficient construction of (distributed) verifiable random functions. In *PKC*, 2003.
30. Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. In *PKC*, 2005.
31. Y. Dodis, A. Yampolskiy, and M. Yung. Threshold and proactive pseudo-random permutations. In *TCC*, 2006.
32. Y. Frankel, P. MacKenzie, and M. Yung. Adaptively-secure distributed public-key systems. In *ESA*, 1999.
33. E. Freire, D. Hofheinz, K. Paterson, and C. Striecks. Programmable hash functions in the multilinear setting. In *Crypto*, 2013.
34. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
35. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Crypto*, number 8042 in LNCS, pages 75–92, 2013.
36. O. Goldreich. Valiant’s polynomial-size monotone formula for majority, 2014.
37. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. In *J. of ACM*, volume 33, 1986.
38. S. Goldwasser, S. Gordon, V. Goyal, A. Jain, J. Katz, F.-H. Liu, A. Sahai, E. Shi, and H.-S. Zhou. Multi-input functional encryption. In *Eurocrypt*, 2014.
39. S. Goldwasser, Y. Kalai, C. Peikert, and V. Vaikuntanathan. Robustness of the Learning with Errors assumption. In *ICS*, 2010.
40. S. Gorbunov, V. Vaikuntanathan, and D. Wichs. Leveled fully homomorphic signatures from standard lattices. In *STOC*, 2015.
41. R. Goyal, S. Hohenberger, V. Koppula, and B. Waters. A generic approach to constructing and proving verifiable random functions. In *TCC*, 2017.
42. J. Hastad, R. Impagliazzo, L. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 8(4):1364–1396, 1999.
43. D. Hofheinz and E. Kiltz. Programmable hash functions and their applications. In *Crypto*, 2008.
44. S. Hoory, A. Hager, and T. Pitassi. Monotone circuits for the majority function. In *APPROX-RANDOM*, 2006.
45. T. Jager. Verifiable random functions from weaker assumptions. In *TCC*, 2015.
46. S. Jarecki and A. Lysyanskaya. Adaptively secure threshold cryptography: Introducing concurrency, removing erasures. In *Eurocrypt*, 2000.
47. S. Katsumata and S. Yamada. Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In *Asiacrypt*, 2016.
48. A. Lewko and B. Waters. Efficient pseudorandom functions from the decisional linear assumption and weaker variants. In *ACM-CCS*, 2009.
49. B. Libert, M. Joye, and M. Yung. Born and raised distributively: Fully distributed non-interactive adaptively secure threshold signatures with short shares. In *PODC*, 2014.

50. B. Libert and R. Titu. Multi-client functional encryption for linear functions in the standard model from LWE. In *Asiacrypt*, 2019.
51. A. Lysyanskaya and C. Peikert. Adaptive security in the threshold setting: From cryptosystems to signature schemes. In *Asiacrypt*, 2001.
52. S. Micali and R. Sidney. A simple method for generating and sharing pseudo-random functions. In *Crypto*, 1995.
53. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proc. of EUROCRYPT*, pages 700–718. Springer, 2012.
54. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
55. M. Naor, B. Pinkas, and O. Reingold. Distributed pseudo-random functions and KDCs. In *Eurocrypt*, 1999.
56. M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS*, 1997.
57. M. Naor, O. Reingold, and A. Rosen. Pseudo-random functions and factoring. In *STOC*, 2000.
58. J.-B. Nielsen. A threshold pseudorandom function construction and its applications. In *Crypto*, 2002.
59. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196. ACM, 2008.
60. A. Raghunathan, G. Segev, and S. Vadhan. Deterministic public-key encryption for adaptively chosen plaintext distributions. In *Eurocrypt*, 2013.
61. A. Razborov and S. Rudich. Natural proofs. *J. of Computer and System Sciences*, 55(1):24–35, 1987.
62. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
63. R. Thorbek. *Linear Integer Secret Sharing*. PhD thesis, Department of Computer Science - University of Aarhus, 2009.
64. L. Trevisan and S. Vadhan. Extracting randomness from samplable distributions. In *FOCS*, 2000.
65. L. Valiant. Short monotone formulae for the majority function. *J. of Algorithms*, 3(5), 1984.
66. L. Valiant. A theory of the learnable. *Comm. of the ACM*, 27(11), 1984.
67. B. Waters. Efficient identity-based encryption without random oracles. In *Eurocrypt*, 2005.
68. S. Yamada. Asymptotically compact adaptively secure lattice IBEs and verifiable random functions via generalized partitioning techniques. In *Crypto*, 2017.